

Studies in Mechanobiology, Tissue Engineering and
Biomaterials 17

Liesbet Geris
David Gomez-Cabrero *Editors*

Uncertainty in Biology

A Computational Modeling Approach

 Springer

Studies in Mechanobiology, Tissue Engineering and Biomaterials

Volume 17

Series editor

Amit Gefen, Ramat Aviv, Israel

More information about this series at <http://www.springer.com/series/8415>

Liesbet Geris · David Gomez-Cabrero
Editors

Uncertainty in Biology

A Computational Modeling Approach

 Springer

Editors

Liesbet Geris
Biomechanics Research Unit
University of Liège
Liège
Belgium

David Gomez-Cabrero
Center for Molecular Medicine
Karolinska University Hospital
Stockholm
Sweden

ISSN 1868-2006 ISSN 1868-2014 (electronic)
Studies in Mechanobiology, Tissue Engineering and Biomaterials
ISBN 978-3-319-21295-1 ISBN 978-3-319-21296-8 (eBook)
DOI 10.1007/978-3-319-21296-8

Library of Congress Control Number: 2015945595

Springer Cham Heidelberg New York Dordrecht London
© Springer International Publishing Switzerland 2016

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

Springer International Publishing AG Switzerland is part of Springer Science+Business Media
(www.springer.com)

Preface

We are very proud to present the latest addition to the Springer series *Studies in Mechanobiology, Tissue Engineering and Biomaterials*. When we started this book, we wanted to create a book that could serve as a starting point for graduate students and researchers interested in the development of computational models of biological processes, with a specific focus on how to deal with the inherent uncertainty.

We have managed to get a great international set of authors together, each discussing on a particular aspect of the problem based on their own expertise and research background.

All chapters start with a detailed theoretical description that serves the dual purpose of introducing the technique and providing sufficient details (in the text or by means of references to the literature) for all researchers to start using it themselves. Subsequently one or more examples illustrate how the technique can be used in a practical setting. Chapters are ordered according to the order in which the technique they describe appears in the development and implementation of new models. Reading the book from start to finish will therefore provide new researchers with a quite extensive tool set to get started for themselves. More experienced researchers will find for specific techniques the latest developments and a discussion of future developments.

This book is the end product of a lengthy process which has suffered from some unforeseen delays. Yet the vision and drive always remained present amongst the editors and authors. We are very happy with the end result and hope that readers will enjoy the book as much as we've enjoyed putting it together.

Prof. Liesbet Geris
Prof. David Gomez-Cabrero

Contents

Part I Introduction

- 1 An Introduction to Uncertainty in the Development of Computational Models of Biological Processes 3**
Liesbet Geris and David Gomez-Cabrero

Part II Modeling Establishment Under Uncertainty

- 2 Reverse Engineering Under Uncertainty 15**
Paul Kirk, Daniel Silk and Michael P.H. Stumpf
- 3 Probabilistic Computational Causal Discovery for Systems Biology 33**
Vincenzo Lagani, Sofia Triantafillou, Gordon Ball, Jesper Tegnér and Ioannis Tsamardinos
- 4 Stochastic Modeling and Simulation Methods for Biological Processes: Overview 75**
Annelies Lejon and Giovanni Samaey

Part III Model Selection and Parameter Fitting

- 5 The Experimental Side of Parameter Estimation 127**
Monica Schliemann-Bullinger, Dirk Fey, Thierry Bastogne, Rolf Findeisen, Peter Scheurich and Eric Bullinger

6	Statistical Data Analysis and Modeling	155
	Millie Shah, Zeinab Chitforoushzadeh and Kevin A. Janes	
7	Optimization in Biology Parameter Estimation and the Associated Optimization Problem	177
	Gunnar Cedersund, Oscar Samuelsson, Gordon Ball, Jesper Tegnér and David Gomez-Cabrero	
8	Interval Methods	199
	Warwick Tucker	
9	Model Extension and Model Selection	213
	Mikael Sunnåker and Joerg Stelling	
10	Bayesian Model Selection Methods and Their Application to Biological ODE Systems	243
	Sabine Hug, Daniel Schmidl, Wei Bo Li, Matthias B. Greiter and Fabian J. Theis	
Part IV Sensitivity Analysis and Model Adaptation		
11	Sloppiness and the Geometry of Parameter Space	271
	Brian K. Mannakee, Aaron P. Ragsdale, Mark K. Transtrum and Ryan N. Gutenkunst	
12	Modeling and Model Simplification to Facilitate Biological Insights and Predictions	301
	Olivia Eriksson and Jesper Tegnér	
13	Sensitivity Analysis by Design of Experiments	327
	An Van Schepdael, Aurélie Carlier and Liesbet Geris	
14	Waves in Spatially-Disordered Neural Fields: A Case Study in Uncertainty Quantification	367
	Carlo R. Laing	
15	In-Silico Models of Trabecular Bone: A Sensitivity Analysis Perspective	393
	Marlène Mengoni, Sebastien Sikora, Vinciane d’Otreppe, Ruth Karen Wilcox and Alison Claire Jones	

Part V Model Predictions Under Uncertainty

16 Neuroswarm: A Methodology to Explore the Constraints that Function Imposes on Simulation Parameters in Large-Scale Networks of Biological Neurons 427
David Gomez-Cabrero, Salva Ardid, Maria Cano-Colino, Jesper Tegnér and Albert Compte

17 Prediction Uncertainty Estimation Despite Unidentifiability: An Overview of Recent Developments 449
Gunnar Cedersund

18 Computational Modeling Under Uncertainty: Challenges and Opportunities. 467
David Gomez-Cabrero, Jesper Tegnér and Liesbet Geris

Author Index 477

Part I
Introduction

Chapter 1

An Introduction to Uncertainty in the Development of Computational Models of Biological Processes

Liesbet Geris and David Gomez-Cabrero

Abstract This chapter aims to provide an introduction to the different ways in which uncertainty can be dealt with computational modelling of biological processes. The first step is model establishment under uncertainty. Once models have been established, data can further be used to select which of the proposed models best meets the predefined criteria. Subsequently, parameter values can be optimized for a specific model configuration. Sensitivity analyses allow to assess the influence of the previous choices on the model output. Additionally, model adaptation permits to focus on specific aspects of the model without losing its global predictive capacity. Finally, predictions with the established models should also consider the effect of uncertainty in the model development process.

1.1 Introduction

Computational modelling of biological processes is becoming a standard tool used in biomedical research groups. The amount of examples showing the added value of the computational modelling approach is increasing by the day [4, 10, 11], more so

L. Geris (✉)

Biomechanics Research Unit, University of Liège, Chemin des Chevreuils 1 B52/3,
4000 Liège, Belgium
e-mail: liesbet.geris@ulg.ac.be

L. Geris

Prometheus, Division of Skeletal Tissue Engineering Leuven, KU Leuven,
Herestraat 49, Box 813, 3000 Leuven, Belgium

L. Geris

Biomechanics Section, KU Leuven, Celestijnenlaan 300 C, Box 2419,
3001 Leuven, Belgium

D. Gomez-Cabrero

Unit of Computational Medicine, Department of Medicine, Karolinska Institutet,
Solna, Sweden
e-mail: david.gomezcabrero@ki.se

D. Gomez-Cabrero

Center for Molecular Medicine, Stockholm, Sweden

© Springer International Publishing Switzerland 2016

L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_1

if we include all Systems Medicine approaches [1]. One of the biggest challenges in creating useful models is the way in which they deal with uncertainty—uncertainty related to the experimental data but also to the modelling choices.

Uncertainty is defined in the Oxford dictionary as ‘the state of being uncertain’. Uncertain in turn is defined as ‘not able to be relied on; not known or definite’. Wikipedia defines Uncertainty as ‘a term used in subtly different ways in a number of fields, including philosophy, physics, statistics, economics, finance, insurance, psychology, sociology, engineering, and information science. It applies to predictions of future events, to physical measurements that are already made, or to the unknown. Uncertainty arises in partially observable and/or stochastic environments, as well as due to ignorance and/or indolence’ [5]. Uncertainty in computational biomedicine

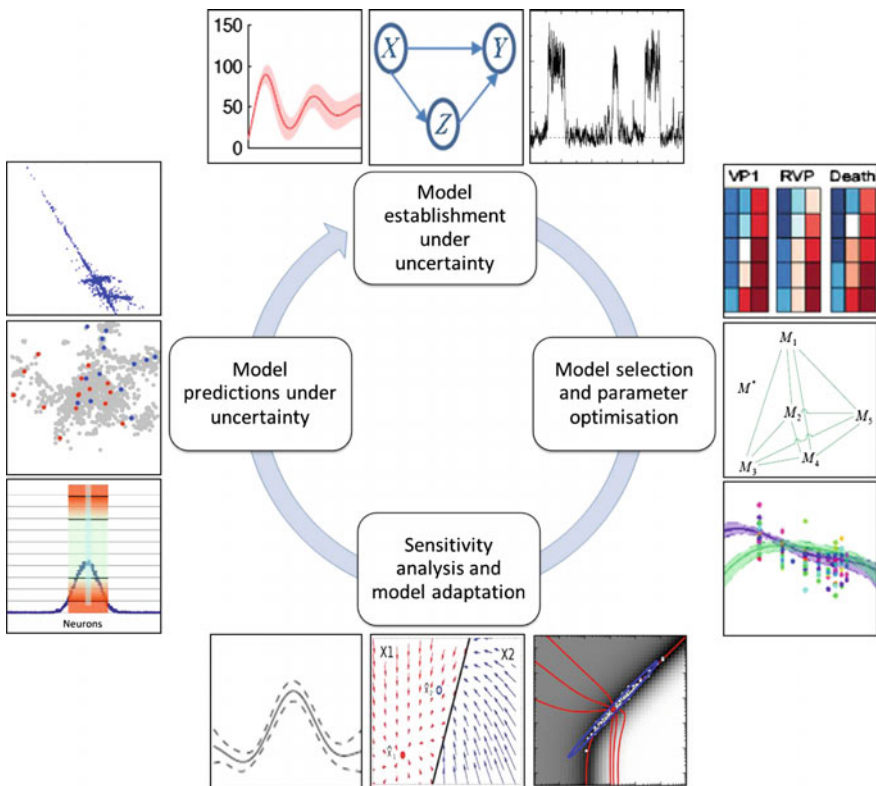


Fig. 1.1 Schematic overview of the model development life cycle with a specific attention to the uncertainty in various forms: (i) establishment of a model under uncertainty, (ii) model selection and parameter optimization, (iii) sensitivity analysis and model adaptation, (iv) model predictions under uncertainty. Snapshot images have been taken from the following chapters in this volume (starting upper left corner going clockwise; C: chapter; F: figure): C2, F1 [12]; C3, F2 [14]; C4, F1 [15]; C6, F4 [19]; C9, F3 [20]; C10, F5 [9]; C11, F5 [16]; C12, F1 [6]; C13, F16 [22]; C16, F3 [8]; C16, F3 [8]; C17, F3 [2]

can come from the experimental observations but can also be connected to the model itself either intrinsically (noise due variation in identically-regulated quantities within a single cell) or extrinsically (noise due variation in identically-regulated quantities between different cell) [12]. Throughout this book various ways of dealing with uncertainty are discussed. The structure of this chapter (and the whole book) follows that of the model development life cycle, starting with model establishment over parameter optimisation and model adaptation and ending with model prediction (Fig. 1.1).

1.2 Model Establishment Under Uncertainty

In order to *turn the ever increasingly available experimental (big) data into actionable knowledge*, mechanistic models are indispensable tools. They provide a conceptual and computational framework that allows for the interpretation and investigation of the experimentally observed behaviour and overcomes some of the limitations of classical statistical models in managing non-linear relations. Setting up these models however poses several challenges related to both the experimental and the modelling side.

Many experimental data sets tend to be noisy or incomplete and most often have not been collected with the specific intention of creating a model. Additionally, oversimplifications in model systems can obscure specific behaviours that have been observed experimentally. These uncertainties have an influence on the establishment of models. The chapter by Kirk et al. [12] provides an overview of the most prevalent problems in model establishment related to uncertainty in data and models, and proposes a number of strategies to tackle these problems. It furthermore discuss various techniques that have been put forward over the years to reverse engineer mechanistic models based on experimental data, each with their advantages and disadvantages. The chapter by Kirk et al. [12] additionally discusses some of the most common inverse techniques for this reverse engineering, including a more elaborate view on statistical inference techniques (additional discussion can be found in the chapter by Sunnåker and Stelling [20]). In the next chapter, Lagani et al. [14] go into detail on a particular kind of statistical predictive models, namely the causal modelling. Computational Causal Discovery allows discovering causal relations with a limited set of interventions or manipulations. Causal modelling goes beyond traditional statistical predictive modelling as it provides the capacity to predict the effects of actions and interventions on a system, e.g., the effects of drug treatment, gene knock-out, or the induction of a mutation in the genome. This is in contrast to non-causal predictive modelling which is only valid when the system under study is observed under the same experimental conditions it was derived from and is not otherwise manipulated. Lagani et al. [14] review the definition of causality and the basic concepts and principles of causal discovery, the nature of the underlying assumptions—particularly in relation to the uncertainty in the available data, potential pitfalls when applying the method, the most recent advances in the field and future directions. Both Kirk et al. [12] and Lagani et al. [20] are data-driven approaches to modelling under uncertainty.

A radically different way to establish models under uncertainty is the use of stochastic modelling and simulation techniques—which is particularly relevant when fluctuations become important. In Lejon and Samaey [15], the authors give a high-level overview of stochastic modelling techniques used for biological problems. They show the effect of stochasticity at different scales and different levels of description and provide computationally reasonable solutions and algorithms for various problem types. They pay particular attention to the equivalence between the stochastic process that governs the evolution of individual agents and the deterministic behaviour of the related probability distributions.

1.3 Model Selection and Parameter Optimisation

Once a number of potential models is established a choice needs to be made on which model and which parameters are the most appropriate to use given the experimental data that is available and the context in which the model will be used. This means that we need to understand the nature of the experimental data that is available to feed the models. After data acquisition, the use of data-driven modelling approaches allows to do a first processing of the data. Once a mechanistic model has been established, parameter estimation and optimisation can be performed in a variety of ways. Ultimately, when different modelling scenarios remain possible, specific tools can be used to determine which of these models is the most suitable, given the available data, the context and the preference of the modeller.

Western blotting, flow cytometry, protein mass spectrometry, DNA microarrays are just a few examples of a wide variety of experimental techniques frequently used in wet labs to gather data. In order to be useful for quantitative dynamic models, data needs to have a dimension of time as well as several perturbation experiments. The chapter by Bullinger-Schliemann [18] provides an introduction to various experimental techniques that are frequently used in the model development life cycle, paying particular to the significance of single-cell versus population measurements. With the increase in availability of large and structured datasets, there has been a need to develop efficient data analysis techniques. Data-driven approaches, in contrast to mechanistic approaches, do not make assumptions on the underlying mechanisms. They are often used to process data to a more useful format and are particularly helpful in identifying biomarkers in large datasets. Shah et al. [19] discuss a particular type of empirical models, namely the eigenvalue-based approaches. These approaches, including singular value decomposition, principle component analysis, and partial least squares regression, can identify important characteristics of big datasets through decomposition and dimensionality reduction. The chapter further discusses to way to deal with upscaling of these methods for understanding higher-order datasets (through tensor decomposition).

In the previous sections we described methodologies aimed to generate models from data; however data can also be gathered to define parameters in a model. When dealing with mechanistic models, the assignment of values to the parameter in a

model, i.e. the parameter estimation, is a crucial step. Depending on the type and amount of data available, this can also be particularly time-consuming task. Over the last years, building on computational and algorithmic developments, many new tools have been developed to facilitate the parameter estimation step. A specific aspect of the estimation process is the optimisation where the parameter space is determined that provides the most interesting results. An overview of these estimation and meta-heuristic optimisation techniques (simulated annealing, genetic algorithms, particle swarm optimisation and others) can be found in Samuelson et al. [3]. As an alternative to these statistical-type parameter estimation methods, Tucker [21] describes parameter estimation via set inversion and constraint propagation techniques (interval methods). These techniques, based on set-valued computations combined with branch and bound steps, allow to examine entire sets of parameters and thus complete the global search within a finite number of steps. As the potential downside of interval methods is their relatively low speed, the author additionally shows how the method can be accelerated by set-valued constraint propagation, allowing for a considerable improvement of its efficiency. Samuelson et al. [3] furthermore provide concepts and tools that allow the modellers to select the appropriate methodology for the specific scenario they are confronted with.

Once several models have been established, model assessment can help in identifying which model is the most appropriate for a given situation. Sunnåker and Stelling [20] discuss the most commonly used methods for model assessment of dynamical models, along with the underlying concepts and ideas. These methods include the information theoretic (e.g. the Akaike and deviance information criteria) and Bayesian approaches (e.g. posterior ratios for relative model probabilities from Bayes factors and the approximate Bayesian information criterion) as well as techniques such as cross-validation and bootstrapping. Bayesian model selection for biological dynamical systems is further elaborated by Hug et al. [9], working with the Bayes factor computed by Thermodynamic Integration. Fundamentally different approaches to model selection (as compared to Bayesian approaches) are also treated, e.g. the minimum description length. All techniques are illustrated with examples ranging from simple and sometimes analytically tractable problems, to medium sized models composed of ordinary differential equations. Information on how the most important results can be derived is provided in [20], alongside with a discussion on differences between methods [9, 20] and how these methods can be employed in practice as there is no generally applicable method for model assessment that is valid in all situations.

1.4 Sensitivity Analysis and Model Adaptation

Despite the techniques identified in the previous section, leading to the selection of the optimal model populated with the optimal parameter set, the uncertainty in the available data is often such that additional analyses of the parameter space and even model adaptation might necessary. Again, a variety of techniques is available

to study the parameter space. Some techniques focus on the general character of the parameter space (e.g. sloppiness) for specific model types. Other techniques focus on specific pre-defined ranges in parameter space assessing the importance of their influence on the model results, i.e. sensitivity analyses. For over-parametrized or vary complex models, various simplification and reduction techniques have been developed to enable the understanding of the model's underlying core dynamics and a subsequent simplification of the model whilst maintaining its capability of capturing those core dynamics.

Exploring the parameter space can be a very challenging task due to its high dimension and complex structure. Mannakee et al. [16] have shown that there exists a universal structure in the parameter space of models for nonlinear systems. More specifically, these models are often sloppy, with strong parameter correlations and an exponential range of parameter sensitivities all leading to good model behaviour. In their chapter [16], the authors review the evidence for universal sloppiness and its implications on parameter fitting and model prediction. They discuss how careful experimental design can lead to optimisation of parameter inference or general model behaviour (depending on the goals of the model and the modeller). They furthermore discuss the potential of transforming parameters to alleviate sloppiness. However, even when taking an information geometry perspective in order to have a parametrization-independent perspective on modelling, sloppiness arises and a deeper universal structure is revealed.

Rather than looking at the global parameter space, some methods specifically focus on a well-defined area, starting from specific intervals for all parameters present in the model (capturing the uncertainty of the parameters). 'Design of Experiments' (discussed by Van Schepdael et al. [22]) is a technique originally developed to optimize physical experiments allowing to comprehensively determine the effect of parameters settings (individual parameters and their interactions) on the process with a minimal amount of experimental runs. For computational models, the limits on the specific parameter values that can be tested and the amount of runs that can be executed are generally less stringent but the amount of parameters (and especially their interactions) might be considerably higher than for physical experiments. The design of experiment approach for computational models allows choosing a minimum amount of parameter combinations that will result in a maximum amount of information about the computational model. The chapter by Van Schepdael et al. [22] explains several designs and analysis methodologies.

The aforementioned methods all start from the fully developed model as derived in Sects. 1.2 and 1.3. However, with computational models of biological processes continuously increasing in size and model complexity (in part due to the data explosion in biology) it is increasingly difficult to obtain insights into what parts of a model generate a specific read-out. This hampers the correct interpretation of the model result and their use in e.g. the design of personalised therapies. The uncertainty in model structure and model parameters is a further complication. A solution to this dual problem of complexity and uncertainty is the systematic construction of simplified models from complex models. In their chapter, Eriksson et al. [6] review different methods for simplification and reduction of models with particular focus on recent

developments such as the iterative “tearing, zooming and simplifying” approach. This approach allows utilizing specific biological features such as modularity and robustness.

To wrap up the part on sensitivity analysis and model adaptation, two elaborate case studies are provided that investigate the sensitivity and effect of uncertainty on model outcome in the context of neural fields and bone mechanics. Laing et al. [13] discuss the introduction of randomly chosen “frozen” spatial noise to their modelling system. The effect of inclusion of said noise on particular model outcomes, such as the occurrence of specific activity in a particular neural field model, is investigated and discussed. The second example is that of Mengoni et al. [17] who provide an overview of computational mechanical modelling of trabecular bone from a sensitivity analysis perspective. The effect of model development choices on model results is reviewed and analysed at different scales (from micro up to organ). As the focus is on models generated starting from Computed Tomography images, particular attention goes to the image processing effects, the mesh-related aspects and the computational representation of the boundary conditions.

1.5 Model Predictions Under Uncertainty

With the model and its parameters all set, model predictions can be made that feed-back to the experiments, closing the modelling life cycle. Model predictions will assist in advancing knowledge of the system under study in various ways. One such type of predictions is the identification of alternative explanations for and interpretations of the existing experimental data. Another type is the discovery of specific mechanisms in the simulation data. Both will lead to the formulation of additional experiments that need to be executed in order to validate (or falsify) the model’s observations.

Gomez-Cabrero et al. [8] start from a very specific pre-frontal cortex working memory model and discuss issues related to non-uniqueness of parameter sets and the existence of various alternative solutions that can explain one particular experimental phenomenon. Using optimization techniques, they uncovered compensatory mechanisms in a subset of the parameters in the model, leading to the identification of hypothesis to be validated in dedicated experiments. On a more general note, Cedersund [2] provides an overview of various types of predictions that can be made—core predictions allowing to test the quality of the model or poorly determined predictions allowing to improve the overall well-determination of the model parameters. Even predictions that will not be tested experimentally can provide interesting insights into the studied model. In a medical context, reliability and accuracy of the predictions is important. Noteworthy is that this low degree of model uncertainty does not necessarily imply a similarly low degree of uncertainty on the model parameters. Such well-determined predictions are then also amenable to incorporation in

larger supermodels (e.g. models of individual organs connected into a multi-organ model). Cedersund [2] subsequently provides an overview of the recent developments in the methods dealing with prediction uncertainty and discusses the price that needs to be paid when bothering with prediction uncertainty.

1.6 Conclusion

This chapter has provided a brief overview of the model development life cycle with a specific focus on uncertainty in the various stages: (i) establishment of a model under uncertainty, (ii) model selection and parameter optimization, (iii) sensitivity analysis and model adaptation, (iv) model predictions under uncertainty. Each of the following chapters in this book elaborates in a detailed way one or more facets of this development life cycle, with a specific attention to the incorporation of uncertainty in data and modelling. Taken together, the information provided in this book should allow modellers to start from experimental data, work through the different modelling life cycle steps and finally make predictions that can be verified experimentally. The last chapter, Gomez-Cabrero and Geris [7], provides also an overview of nowadays open challenges.

Acknowledgments LG gratefully acknowledges funding the Belgian National Fund for Scientific Research (FNRS) grants FRFC 2.4564.12 and T.0254.13, as well as the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement n°279100.

References

1. Bousquet, J.I., Anto, J.M., Sterk, P.J., Adcock, I.M., Chung, K.F., Roca, J., Agusti, A., Brightling, C., Cambron-Thomsen, A., Cesario, A., Abdelhak, S., Antonarakis, S.E., Avignon, A., Ballabio, A., Baraldi, E., Baranov, A., Bieber, T., Bockaert, J., Brahmachari, S., Brambilla, C., Bringer, J., Dauzat, M., Ernberg, I., Fabbri, L., Froguel, P., Galas, D., Gojobori, T., Hunter, P., Jorgensen, C., Kauffmann, F., Kourilsky, P., Kowalski, M.L., Lancet, D., Pen, C.L., Mallet, J., Mayosi, B., Mercier, J., Metspalu, A., Nadeau, J.H., Ninot, G., Noble, D., Oztürk, M., Palkonen, S., Préfaut, C., Rabe, K., Renard, E., Roberts, R.G., Samolinski, B., Schüнемann, H.J., Simon, H.U., Soares, M.B., Superti-Furga, G., Tegnér, J., Verjovski-Almeida, S., Wellstead, P., Wolkenhauer, O., Wouters, E., Balling, R., Brookes, A.J., Charron, D., Pison, C., Chen, Z., Hood, L., Auffray, C.: Systems medicine and integrated care to combat chronic noncommunicable diseases. *Genome Med.* **3**(7), 43 (2011). doi:[10.1186/gm259](https://doi.org/10.1186/gm259)
2. Cedersund, G.: Prediction uncertainty estimation despite unidentifiability: an overview of recent developments. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
3. Cedersund, G., Sameulsson, O., Ball, G., Tegnér, J., Gomez-Cabrero, D.: Optimization in biology parameter estimation and the associated optimization problem. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
4. Clermont, G.I., Auffray, C., Moreau, Y., Rocke, D.M., Dalevi, D., Dubhashi, D., Marshall, D.R., Raasch, P., Dehne, F., Provero, P., Tegnér, J., Aronow, B.J., Langston, M.A., Benson, M.:

- Bridging the gap between systems biology and medicine. *Genome Med.* **1**(9), 88 (2009). doi:[10.1186/gm88](https://doi.org/10.1186/gm88)
5. Consulted March 30th 2015. <http://en.wikipedia.org/wiki/Uncertainty>
 6. Eriksson, O., Tegnér, J.: Modeling and model simplification to facilitate biological insights and predictions. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 7. Gomez-Cabrero, D., Tegnér, J., Geris, L.: Computational modeling under uncertainty: challenges and opportunities. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 8. Gomez-Cabrero, D., Ardid, S., Cano-Colino, M., Tegnér, J., Compte, A.: Neuroswarm: a methodology to explore the constraints that function imposes on simulation parameters in large-scale networks of biological neurons. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 9. Hug, S., Schmidl, D., Li, W.B., Greiter, M.B., Theis, F.J.: Bayesian model selection methods and their application to biological ODE systems. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 10. Hunter, P., Coveney, P.V., de Bono, B., Diaz, V., Fenner, J., Frangi, A.F., Harris, P., Hose, R., Kohl, P., Lawford, P., McCormack, K., Mendes, M., Omholt, S., Quarteroni, A., Skår, J., Tegnér, J., Randall Thomas, S., Tollis, I., Tsamardinos, I., van Beek, J.H., Viceconti, M.: A vision and strategy for the virtual physiological human in 2010 and beyond. *Philos. Trans. A Math. Phys. Eng. Sci.* **368**(1920), 2595–2614 (2010). doi:[10.1098/rsta.2010.0048](https://doi.org/10.1098/rsta.2010.0048)
 11. Hunter, P., Chapman, T., Coveney, P.V., de Bono, B., Diaz, V., Fenner, J., Frangi, A.F., Harris, P., Hose, R., Kohl, P., Lawford, P., McCormack, K., Mendes, M., Omholt, S., Quarteroni, A., Shublag, N., Skår, J., Stroetmann, K., Tegnér, J., Thomas, S.R., Tollis, I., Tsamardinos, I., van Beek, J.H., Viceconti, M.: A vision and strategy for the virtual physiological human: 2012 update. *Interf. Focus.* **3**(2), 20130004 (2013). doi:[10.1098/rsfs.2013.0004](https://doi.org/10.1098/rsfs.2013.0004)
 12. Kirk, P., Silk, D., Stumpf, M.P.H.: Reverse engineering under uncertainty. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 13. Laing, C.R.: Waves in spatially-disordered neural fields: a case study in uncertainty quantification. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 14. Lagani, V., Triantafyllou, S., Ball, G., Tegnér, J., Tsamardinos, I.: Probabilistic computational causal discovery for systems biology. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 15. Lejon, A., Samaey, G.: Stochastic modeling and simulation methods for biological processes: overview. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 16. Mannakee, B.K., Ragsdale, A.P., Transtrum, M.K., Gutenkunst, R.N.: Sloppiness and the geometry of parameter space. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 17. Megnoni, M., Sikora, S., d’Ottreppe, V., Wilcox, R.K., Jones, A.C.: In-silico models of trabecular bone: a sensitivity analysis perspective. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 18. Schliemann-Bullinger, M., Fey, D., Bastogne, T., Findeisen, R., Scheurich, P., Bullinger, E.: The experimental side of parameter estimation. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 19. Shah, M., Chitforoushzadeh, Z., Janes, K.A.: Statistical data analysis and modeling. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 20. Sunnåker, M., Stelling, J.: Model extension and model selection. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 21. Tucker, W.: Interval methods. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
 22. Van Schepdael, A., Carlier, A., Geris, L.: Sensitivity analysis by design of experiments. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)

Part II
Modeling Establishment Under
Uncertainty

Chapter 2

Reverse Engineering Under Uncertainty

Paul Kirk, Daniel Silk and Michael P.H. Stumpf

Abstract The increased availability of experimental data in systems biology and systems medicine can only lead to better understanding of biological and disease related processes, if we can place them in the context of mechanistic models. Such models can serve as conceptual, but also computational frameworks in which we can reason about, or predict the behaviour of e.g. molecular networks, or cellular processes. Constructing such models, however, remains a formidable challenge: not only are the data noisy and incomplete, but the models that are currently available are hopelessly oversimplified. In this chapter we set out the problems and a list of potential ways of tackling them. The essential premise is always to be aware of the uncertainties inherent in the data and our models.

Keywords Inverse problems · Model selection · Extrinsic versus intrinsic noise · Model misspecification

2.1 Introduction

Reverse engineering the processes that govern the behaviour of biological systems is one of the principal aims of systems biology [46]. From experimental data, we seek to elucidate key aspects of the underlying mechanisms that give rise to observed complex behaviour. We may initially have only very vague, perhaps even wrong, ideas regarding these mechanisms, in which case our first aim may be to use the data in order to generate testable hypotheses. Alternatively, we may have already expressed our existing hypotheses as one or more mathematical models, in which case we may wish to use the data in order to tune their parameters, or to choose between them.

A defining feature of reverse engineering in a biological context is the variety of ways in which we encounter uncertainty [11]. In addition to the usual challenges

P. Kirk · D. Silk · M.P.H. Stumpf (✉)
London, UK
e-mail: m.stumpf@imperial.ac.uk

© Springer International Publishing Switzerland 2016
L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_2

presented by measurement noise, we must also contend with the inherently stochastic nature of biochemical and biophysical processes. Moreover, given the complexity and interconnectedness of biological processes, we are currently only able to probe incomplete portions of the systems of interest, which has obvious consequences for the analysis [12]. While *in vitro* studies typically provide us with more control and might even enable us to isolate a particular process, we are still faced with the problem of establishing whether this idealised environment can be representative of the much more complex one that exists *in vivo* [33]. This combination of measurement noise, incomplete observations, and inherently nonlinear and stochastic underlying processes makes reverse engineering biological systems a particularly difficult task.

In this chapter, we discuss some of the challenges presented by reverse engineering under uncertainty in a biological context. In Sect. 2.2, we provide a broad overview of the *inverse problem* in systems biology, and consider the various ways in which this problem is encountered in practice. We then consider manifestations of uncertainty in Sect. 2.3, and ways in which we can try to cope with them when addressing the inverse problem. In Sect. 2.4, we consider the consequences of uncertainty in the context of modelling, and the potential limitations that uncertainty imposes on what we are able to learn. We offer some final conclusions and advice in Sect. 2.5.

2.2 The Inverse Problem in Systems Biology

An *inverse problem* is one in which we seek to reverse engineer details of a system (or data-generating mechanism) from experimental observations or measurements [49]. Typically, this will involve inferring a model or its parameters from experimental data. In contrast, a *forward problem* is one in which we have a fully specified model and we use it to make predictions or draw conclusions about its behaviour. There is clearly an interplay between inverse and forward problems: a reverse engineered model can subsequently be used for prediction, while a model whose predictions disagree with novel experimental observations might form the basis for a new model. The inverse problem has gained particular prominence in systems biology [22, 55, 58, 61], where we often have access to large quantities of high-throughput data, but may initially lack a deep understanding of how these measurements relate to one another, or what they can tell us about the underlying biological processes [36].

The difficulty of the inverse problem is hard to overstate. Even for simple systems (in terms of the model) it presents formidable challenges and is vastly more complicated than any associated forward problem.

2.2.1 The Different Types of Inverse Problems

We can consider three different, yet closely related, types of inverse problem: (i) we do not have a model and need to reverse engineer one from the data; (ii) we have

a model, the parameters of which need to be estimated/inferred from the data; and (iii) we have a number of distinct candidate models (for which we may or may not know the parameters) and we need to choose between them.

The first type of inverse problem has attracted a lot of attention in systems biology, particularly in the context of *network inference* [22, 38, 39, 50, 57]. Network inference approaches often proceed by first calculating measures of statistical dependence between different biological entities (which form the nodes of the network), and then identifying the pairs of entities between which there is a significant statistical dependence (these define the edges of the network). Some approaches take pains to try to identify direct, causal relationships by eliminating conditional dependencies. Network inference techniques typically have the advantage of being applicable to large-scale problems (e.g. finding dependencies between the expression levels of genes). The resulting network representations tend to be *descriptive* rather than *predictive*, and hence network inference is often seen as a method for hypothesis generation, which may be a first step toward developing more detailed mechanistic models.

The second type of inverse problem describes the problem of estimating the parameters of a known (or assumed) model, which is sometimes known as *model calibration*. In addition to more heuristic methods [6], approaches such as maximum likelihood estimation [56] and Bayesian inference have gained traction in recent years as ways in which to tackle model calibration problems. We consider these methods in more detail in Sect. 2.2.2.

The third type of inverse problem refers to *model selection*. In this case, we wish to choose the ‘best’ model(s) from a collection (and/or may wish to reject the ‘worst’). Usually, our assessment of a model requires us to strike a balance between two criteria: (i) quality of fit; and (ii) complexity. In the interests of parsimony (also known as *Occam’s razor*), we ideally wish to maximise the former while minimising the latter, and numerous approaches exist that seek to address this problem. Measures such as the Akaike information criterion (AIC) and the Bayesian information criterion (BIC) do this by combining an assessment of quality of fit with a penalty on the number of parameters (which is taken as a proxy for model complexity). Alternatively, Bayesian approaches usually focus on estimating the *evidence* (or *marginal likelihood*) for different models, and then compare these quantities via the calculation of *Bayes factors*. Marginal likelihood estimation is typically challenging and computationally costly; however, Bayesian approaches have the advantage of naturally embodying the principle of Occam’s razor. These and other procedures for model selection are discussed in more detail in [28, 47].

2.2.2 Statistical Inference Approaches

The general problem of fitting a model to data is often approached by considering some function that quantifies the discrepancy (or, alternatively, agreement) between the model’s predictions and the observed data, and then tuning the model’s parameters

in order to obtain a good fit. Examples of the kind of discrepancy function that might be employed include quantitative distances such as sum of squares and sum of absolute difference errors, or qualitative measures such as the eigenvalues or Lyapunov spectrum of a dynamical system [3, 41]. The choice of discrepancy function is usually based upon heuristic arguments, but is often important, affecting (for example) the degree to which outliers influence the fit. A key problem when adopting such a fitting approach is how to find the minimum of the discrepancy function, and numerous optimisation strategies exist that can be applied for this purpose [7, 53]. Two further important considerations are: (i) the problem of local minima; and (ii) over fitting. The first of these refers to the common problem of the optimisation algorithm getting “stuck” in a local minimum, rather than identifying the parameters that yield the true, global minimum. The second refers to the challenge of how to avoid fitting the experimental noise [40], which will typically result in poor predictive performance.

If our model is probabilistic, we will often be able to define a *likelihood function* [10], $L(\theta) = p(D|\theta, \mathcal{M})$, which scores parameters by assessing how likely the observed data, D , would be under the assumption that those parameters θ (and our model, \mathcal{M}) are correct. In *maximum likelihood* (ML) estimation, we seek the parameters that maximise this likelihood function. In order to improve numerical stability, in practice we often work with the *log* likelihood function. Moreover, due to the way in which optimisation routines are typically implemented, we often think in terms of *minimising* the *negative* log likelihood, which we can consider as a particular kind of discrepancy function that happens to have the advantage of having a formal probabilistic grounding. The challenges of escaping local minima and avoiding over fitting remain.

The Bayesian formalism [16, 45], provides a framework for performing parameter inference, in which assessments of fit (as quantified by the likelihood function) are combined with our prior belief regarding the parameter values. Here, “prior belief” refers to the belief we have before observing the current dataset, and may have been obtained on the basis of previous experiments (e.g. on related biological systems, or in similar conditions). The Bayes rule provides us with a formal mathematical means by which to update our prior belief in light of the observed data, in order to obtain the *posterior distribution*. The posterior quantifies the uncertainty remaining in the values of the parameters after having observed the data, and may be used to derive *credible regions* for the parameter vector. More precisely, we have,

$$p(\theta|D, \mathcal{M}) = \frac{p(\theta|\mathcal{M})p(D|\theta, \mathcal{M})}{\int_{\theta \in \Theta} p(\theta|\mathcal{M})p(D|\theta, \mathcal{M})d\theta}, \quad (2.1)$$

where D is the dataset, θ is the vector of parameters that is to be inferred, and \mathcal{M} represents the model. In words, we have,

$$\text{Posterior} = \frac{\text{Prior} \times \text{Likelihood}}{\text{Model evidence}}. \quad (2.2)$$

In practice, elucidation of the posterior distribution is rarely possible analytically, and hence we must resort to techniques for obtaining samples from the posterior, such as Markov chain Monte Carlo (MCMC), sequential Monte Carlo (SMC), or nested sampling [25, 42]. For some problems it may not even be possible to write down the likelihood, in which case *approximate likelihood* techniques and approaches such as approximate Bayesian computation might be appropriate [52, 60].

2.2.3 *Bypassing the Inverse Problem*

It is usually impossible to measure all of the parameters or all of the components of a biological system experimentally, and hence addressing the inverse problem is an unavoidable reality. However, even if we *were* able to measure all of these quantities, they would only be valid for the particular experimental and biological conditions under which the measurements were taken; molecular reaction rates, for example, depend on ambient temperature and pH values among many other things. Given that these conditions are themselves subject to random fluctuations, modelling the variation in these quantities is of vital importance if we wish to understand the sources of uncertainty and variability in the system and in our data.

2.3 Manifestations of Uncertainty

One of the most significant challenges to be overcome when trying to reverse engineer biological processes is the variety of sources of uncertainty that we must take into account. In this section, we describe the various sources of noise that might be important, and discuss strategies for coping with this noise when performing inference.

2.3.1 *Sources of Noise*

There are many different sources of noise that have an impact on if and how we can reverse engineer a given biological process. On the one hand, we have *experimental noise*, which arises from imprecision and or inaccuracy in the measurement process. On the other, we have the inherently stochastic nature of the underlying biological system [4, 29, 51], which is a component of what we seek to reverse engineer. In the context of cellular noise, this is often investigated in terms of *intrinsic* and *extrinsic* sources.

2.3.1.1 Experimental Noise

In the analysis of experimental error, a distinction is made between the precision and accuracy of an observation [37]. Precision refers to the inherent error distribution associated with a particular type of experiment, and accuracy to the existence of systematic errors in the experimental process. Contributions to the former will vary for repeat observations and include, for example, random fluctuations in the experimental conditions or behaviour of the experimental instruments. In contrast, systematic errors remain unchanged for repeated experiments, and are caused by, for example, imperfect calibration of experimental instruments. If the cause is known, systematic errors should be explicitly modelled in order to avoid bias in any inferred quantities. Otherwise, undetected systematic error can be viewed as a source model misspecification which will be discussed more generally in Sect. 2.4.2.

2.3.1.2 Intrinsic Noise

Cellular behaviour is governed by the biochemical reactions that occur between different molecular species within the cell. The timing of individual reactions is a random quantity, which gives rise to the source of cellular stochasticity known as intrinsic noise. Since each individual reaction only changes the numbers of molecules of the reacting species by one or two, the effects of intrinsic noise are particularly important when there are only low copy numbers of the molecular species of interest.

2.3.1.3 Extrinsic Noise

Extrinsic noise refers to variability in the physical and biological environment within which the intrinsically noisy interactions take place. For example, a collection of cells may vary in cellular volume, be at different stages of the cellular cycle, or have different abundance of RNA polymerase and ribosomes; all of which may contribute to variability in behaviour between cells and subsequent experimental measurements.

2.3.2 *Coping with Uncertainty in Inference*

Having identified a variety of sources of noise, we now discuss how we should address or capture these when performing parameter inference. The key question is how to model each type of noise, so that we can either derive a likelihood function (and hence adopt a maximum likelihood or Bayesian approach) or else find some other (possibly simulation-based) method for inferring parameters. We consider strategies for coping with each of the three types of noise identified in the previous section.

2.3.2.1 Coping with Experimental Noise

Each source of random experimental uncertainty may be categorised further as applying to either the inputs or outputs of an experiment. In the latter case the measurement error, ϵ , is typically assumed independent of both the parameters and known inputs, (θ, u) , and the true state of the system. The likelihood thus factorises into components describing the uncertainty generated by the system and parameters, and by the measurement process,

$$\begin{aligned} L(\theta) &= p(D|\theta, u) \\ &= p(D^*|\theta, u)p(\epsilon) \end{aligned}$$

where D^* is the error-free (i.e. absent of experimental noise) state of the system.

In the less commonly discussed case of uncertain inputs, the true state of the observable is no longer independent of the uncertainty in question, and the likelihood is obtained by integrating over possible values of u ,

$$L(\theta) = p(\epsilon) \int p(D^*|\theta, u)du. \quad (2.3)$$

The integral in Eq. 2.3 describes how the error propagates through the system for particular values of θ , and often may only be approximately evaluated. A variety of methods to do so exist, including Monte Carlo approaches [35], Sigma point methods [26], or Gaussian quadrature [43], the appropriateness of each of which is determined by both the complexity of the system model, and the distribution, $p(u)$.

Commonly the total experimental error is summarised as additive and Gaussian. Such an approximation may be justified (as a consequence of the Central Limit theorem) when the errors are the accumulation of large numbers of independent sources of uncertainty. The Gaussian assumption is certainly computationally convenient. For example, if all sources of uncertainty and the data itself are Gaussian distributed, then calculation of the integral in Eq. 2.3 may be undertaken with relative efficiency (e.g. by using the unscented transform [27]). However, it is important to note that the effects of input error (even when assumed Gaussian) on $p(y)$ will almost certainly not be Gaussian in the presence of any non-linearity. Further, care must be taken when measured quantities lie close to limiting boundaries (e.g. abundance or concentration is strictly positive), as this can induce non-Gaussian effects upon the error distribution. In these cases, more sophisticated and computationally expensive Monte-Carlo based approaches are necessary for evaluating the likelihood.

2.3.2.2 Coping with Intrinsic Noise

We assume that the available data comprise intrinsically noisy measurements obtained at discrete time points. While it is possible to derive exact Markov chain Monte Carlo schemes for inference in such situations, their computational cost is usually

prohibitively expensive. However, a number of approaches exist for *simulating* intrinsic fluctuations, and hence several *simulation-based* inference procedures have been proposed. We refer the reader to [19, 59] for examples. At the heart of all of these approaches is simulation using Gillespie’s stochastic simulation algorithm (SSA) [18] (see also the top plots in Fig. 2.1 for example realisations). Given a chemical reaction system with known rate constants and initial molecule numbers, the SSA proceeds by using Monte Carlo techniques to simulate both the time until the next reaction, and the next reaction to occur. A number of modifications exist in order to accelerate simulation using the SSA, including the Gibson-Bruck algorithm and the τ -leap method [19, 59]. All of these simulation methods have in common that they provide exact realisations from the underlying (discrete state, continuous time) stochastic kinetic model.

Alternative methods for parameter inference approximate the underlying stochastic kinetic model in order to derive *approximate likelihood* functions. A popular approach is to consider the continuous-state *diffusion approximation* of the true process, which yields a stochastic differential equation (SDE) known as the chemical Langevin equation (CLE). An alternative continuous approximation is given by the linear noise approximation (LNA) [20]. Additionally, several moment expansion and moment closure approaches have been proposed as ways of approximating the underlying model, some of which have also been used in order to allow parameter estimation to be performed.

2.3.2.3 Coping with Extrinsic Noise

Extrinsic noise may be modelled by specifying a probability distribution, $p(\theta, x_0)$, over the parameters and initial conditions [4]. In Fig. 2.1 (right column), we illustrate the effects of extrinsic noise on the oscillations in a model of p53 dynamics, where the extrinsic noise enters the model through fluctuations in just one of the parameters. In this example, we have both intrinsic and extrinsic effects (see also Sect. 2.3.2.4). In the absence of intrinsic stochasticity, extrinsic effects may be simulated in exactly the same way as propagating input uncertainty (discussed in Sect. 2.4.2)—by propagating $p(\theta, x_0)$ through the model. The parameters of the extrinsic noise distribution $p(\theta, x_0)$, may also be the subject of inference given suitable data, such as multiple measurements at single cell resolution.

2.3.2.4 Coping with Mixed Noise Sources

When intrinsic and extrinsic noise are both present, the modelling challenges are more substantial, both conceptually and computationally. The most common approach, originating from [48], is to derive a framework under which each source of noise may be considered separately, whilst other sources are held fixed. The theoretical justification is made via the following decomposition of the stochasticity of cellular products, x , as the direct sum of extrinsic and intrinsic (and experimental) contri-

butions. Defining the extrinsic and intrinsic variables (or parameters) as E and I respectively, the total law of variance gives us,

$$\sigma_x^2 = \underbrace{\sigma_{\langle x|E \rangle}^2}_{\text{Extrinsic}} + \underbrace{\langle \sigma_{\langle x|E,I \rangle|E}^2 \rangle}_{\text{Intrinsic}} + \underbrace{\langle \sigma_{x|E,I}^2 \rangle}_{\text{Experimental}} \quad (2.4)$$

where the angular brackets represent the expectation. The first term is the variance of the mean values of x with E held fixed, and describes the portion of the total uncertainty arising from extrinsic variability. The second term describes the intrinsic contribution—the mean variance of x when sources of uncertainty other than E and I are averaged out, and E is held fixed. The final term is that part of the total variance




Fig. 2.1 We consider a model of oscillatory p53 dynamics [17]. The model comprises three protein species (p53, precursor of Mdm2 and Mdm2) connected through a nonlinear feedback loop. We take the parameters of the system (see [1] for details) to be $[k_1, k_2, k_3, k_4, k_5, k_6, k_7] = [90, 0.002, 1.7, 1.1, 0.93, 0.96, 0.01]$, with initial conditions fixed at $[p53, \text{pre-mdm2}, \text{mdm2}] = [10, 20, 30]$ at time $t = 0$. In the *top left*, we show individual realisations of the number of p53 molecules over time, obtained using Gillespie’s stochastic simulation algorithm (SSA). Below this, we have 3 plots showing the population mean (*solid line*) and 1 standard deviation shaded region for the 3 protein species (as indicated), obtained by averaging over many SSA runs. On the *right*, we show the same 4 plots, but this time we illustrate the effects of extrinsic fluctuations by assuming that the k_4 parameter is drawn from a Gamma(12, 0.1) distribution (so that the mode is at $k_4 = 1.1$). While it is difficult to discern any difference from the individual SSA simulations (top plots), it is clear from the plots of the population means that the effect of extrinsic noise in this case is stronger dampening of the oscillations

that is not explained by experimental or intrinsic sources, and which we attribute to uncertainty in the measurement process.

This noise decomposition suggests the innovative dual-reporter experiments—where the products of two genes, regulated by identical promoters are simultaneously measured—in order to quantify intrinsic and extrinsic contributions. Furthermore, it suggests that intrinsic, extrinsic and experimental uncertainty may be modelled jointly by combining their separate strategies in hierarchical fashion. This is demonstrated for intrinsic and extrinsic variability by Toni and Tidor [51], using the linear noise approximation and the unscented transform respectively.

The total law of variance based approach, however, is only accurate when changes in extrinsic variables with time are much slower than fluctuations in intrinsic variables [23]. It turns out that inferring the contributions to total variance from extrinsic and intrinsic sources is reliant upon the history of extrinsic fluctuations and not just their present state. Even if all extrinsic variables can be measured accurately, Eq. 2.4 will introduce errors if the extrinsic variables cannot be assumed constant in time.

2.3.3 Quantifying Information and Knowledge

Given the variety of noise sources that may exist in the underlying processes that generated the data, we may wonder exactly how much information can be extracted from a given dataset. In the context of reverse engineering, our principal concern is the degree to which we will be able to reconstruct the biological process of interest from the available experimental observations. It is therefore useful to be able to quantify the amount of information that our data contain about the parameters that we seek to infer. In the Bayesian formalism, this is conceptually simple. Before we conduct the experiment, the prior distribution describes the knowledge that we have regarding the values of the unknown model parameters. The posterior distribution serves the same role, but *after* observation of the data. The compression from prior to posterior provides an information theoretic measure of the information gain provided by the data. This compression can be quantified by calculating the *Kullback-Leibler divergence* [9] between posterior and prior,

$$d_{KL}(p(\theta|D, \mathcal{M}), p(\theta|\mathcal{M})) = \int_{\theta \in \Theta} p(\theta|D, \mathcal{M}) \log \left(\frac{p(\theta|D, \mathcal{M})}{p(\theta|\mathcal{M})} \right) d\theta. \quad (2.5)$$

Typically, it will not be possible to calculate this divergence analytically; however, there are Monte Carlo methods that permit its estimation.

2.4 Models in Biology and Confidence in Models

2.4.1 Data versus Reality

Despite the increasing range and power of experimental techniques, datasets continue to represent low-dimensional snapshots of the complex cellular environment. It is the task of reverse engineering to interpret the data and fill in the blanks—to explain observed, and allow the prediction of unobserved properties of the real system. It is clear that the quality, quantity, context and subject of experimental observations determines both the inferences that may be drawn and the confidence we associate with them. For example, larger datasets with higher signal to error ratios will in general lead to greater accuracy and precision. However, in many cases the relative utility of different experimental choices can be hard to foresee, e.g., which species should be measured or perturbed (illustrated in Fig. 2.2 and more generally by Liepe et al. [32]), and whether longitudinal datasets or time-point data should be generated in order to reduce the uncertainty in parameter estimates [30].

Here it can be useful to close the loop between experiment and model, by rationally seeking experiments that maximise the expected information available for the inference task at hand. This is known as *experimental design*, of which recent developments in the context of model calibration include the work of Liepe et al. [32] that builds upon existing methods [2, 8, 24, 31, 34, 54], by utilising a sequential approximate Bayesian computation framework to choose the experiment that maximises the expected mutual information between prior and posterior parameter distributions. In so doing, they are able to optimally narrow the resulting posterior parameter or predictive distributions, incorporate preliminary experimental data




Fig. 2.2 Some experiments are more informative than others. **a** Schematic of a three variable system of ordinary differential equations. *Arrows* represent interactions that are modelled by linear terms with coefficients shown. Inference for k is performed independently for two timeseries datasets that are generated by simulating the model with $k = -0.1$, and measuring the state of **b** the top variable and **c** the middle variable for times $t = 0.5, 1, 1.5, 2$. The broadness of the resulting marginal posterior distributions differ substantially, reflecting the different levels of information contained within the datasets

and provide sensitivity and robustness analyses. Design frameworks also exist for model selection (e.g. [15]), where experiments are sought that maximally distinguish the prior predictive distributions of the competing models.

Although experimental design offers a powerful auxiliary tool to statistical inference, care must be taken in interpreting the confidence associated with inferred models and parameters. For example, it is unsurprising that we assign high confidence to the outcome of a model selection analysis given data from an optimally designed experiment. When each of the models is subject to some level of misspecification, such confidence may be misleading.

2.4.2 Models versus Reality

The complexity of cellular behaviour makes it inevitable that reverse (or forward) engineered systems models will be subject to misspecification errors (when they relate to the observation model, they are called systematic errors). These errors in the model may remain undetected, or they may be introduced knowingly via model reductions aimed at simplifying downstream analyses or at increasing interpretability. In either case, such model uncertainty affects predictions and the outcomes of statistical inferences. For example, inferred values for the physical parameters of a ‘wrong’ model will also be ‘wrong’ in order to compensate for misspecification (for example, see Fig. 2.3). Indeed, strictly speaking, Bayesian inference is valid only when a ‘true’ model is considered.

The effects of parameter and input uncertainty may be quantified by assessing their effect on the likelihood and posterior model predictions. For some classes of model

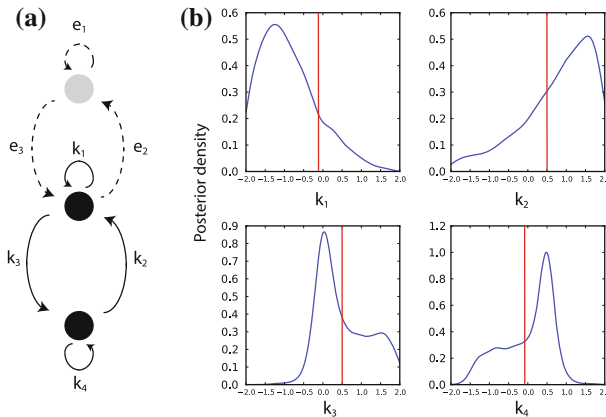


Fig. 2.3 Inference using a ‘wrong’ model. **a** A misspecified model of a ‘true’ data generating system are considered. The grey circle and dotted arrows represent a true variable and its interactions that are absent from the wrong model. Fixing $(e_1, e_2, e_3) = (-0.1, 0.1, 0.5)$, and $(k_1, k_2, k_3, k_4) = (-0.1, 0.5, 0.5, -0.1)$ a timeseries of 10 data points is simulated from which the k_i are inferred using the incomplete model. **b** Marginal posterior densities for the k_i . Maximum a posteriori (MAP) estimates do not coincide with the true parameter values (shown in red)

a similar treatment of model uncertainty may be undertaken by capturing the range of possible errors through parametric expansions, and examining the importance of each (e.g. [44]). However more generally, and certainly for mechanistic models, such an approach is undermined by the conceptual and computational difficulties of specifying the complete space of model errors. However, a consideration of the possible sources of model uncertainty may still suggest a collection of possible models that are in reasonable agreement with the data. In this case, the propagation of misspecification may be managed, to an extent, by conditioning upon the whole collection, rather than just on the single best model. This is the basis of the model averaging framework, where the best estimate of the state or parameter of the system, θ , along with confidence intervals may be calculated from the averaged probability under the various models,

$$p(\theta|\mathcal{M}_1, \dots, \mathcal{M}_N) = \sum_{i=0}^N p(\mathcal{M}_i|D)p(\theta|\mathcal{M}_i)$$

where $p(\mathcal{M}_i|D)$ is the posterior probability of model, \mathcal{M}_i , given the data, D , and $p(\theta|\mathcal{M}_i)$ is the posterior distribution for θ under model \mathcal{M}_i . While each \mathcal{M}_i is still ‘wrong’, the averaged prediction of all ‘wrong’ models at least accounts for a portion of model uncertainty. However the major drawback of averaging, rather than selecting, is to diminish their physical interpretability.

2.4.3 What Can Be Learned from Data?

Frequently we find that parameters, or more often combinations of parameters, can be varied over orders of magnitude without changing the output of a system appreciably [13]. This has major implications for the inverse problem of estimating parameters from data, as large sub-regions in parameter-space may be commensurate with a given dataset. The dependence (or lack thereof) of parameters with respect to data is referred to as *inferability*, which in practice may be quantified as the variance about the Maximum *a posteriori* (MAP) estimate. More formally, the Cramer-Rao inequality [10] gives us a bound on the precision to which a parameter may be estimated in terms of the likelihood,

$$\sigma_{\theta}^2 \geq \mathcal{I}^{-1}(\theta),$$

with $\mathcal{I}(\theta)$ being the Fisher information matrix (FIM),

$$\mathcal{I}(\theta) = E_{\theta} \left[\left(\frac{\partial \log(p(D|\theta))}{\partial \theta} \right)^2 \right],$$

and σ_{θ}^2 , the covariance matrix of a vector-valued θ .

The FIM is at the heart of much of statistical inference and can be interpreted as the curvature of the likelihood surface around the maximal value of the likelihood function. It can also be used as a means to consider robustness and sensitivity of dynamical systems. The reason for this is that if a system is sensitive to variation in a parameter, or a combination of parameters, then this means that changing the parameter, e.g. from θ to $\theta + \delta$, will result in a noticeable change to the system output, which in turn means that the likelihood will also be altered appreciably. Notice, however, that inferability is a property of both system and data—it is possible that further observations will render previously ‘sloppy’ [21] parameters inferable with high certainty (see Fig. 2.2).

Often improved fits to data or better model predictions are interpreted as evidence that more about the true system is being captured. However, it is easy to construct counter-examples where improved data fitting and even predictive power (although desirable in their own right) can be achieved by including more inaccuracies into a misspecified model. It is crucial then not to interpret the physical meaning of any model too assuredly, but instead use them as tools to generate hypotheses for experimental testing (with the result, perhaps, of invalidating the model).

2.5 Conclusion

Reverse engineering is never easy, and probably even harder in biology than in the physical sciences, where sound physical principles can constrain the search space considerably. But once we accept that there is a point to applying quantitative methods and mathematical or computer models in biology, we have to face up to the challenges presented by inverse problems. There have been some arguments, perhaps most notably from Sydney Brenner [5], stating that the inverse problem in molecular and cellular biology is insurmountable and that we should use “the CELLMAP”; how this looks and where it would come from has thus far, sadly or unsurprisingly, been left unspecified.

In order to make progress with the topic of this chapter we have to consider two aspects of reverse engineering. First, problems where models can be tackled by existing methods of reverse engineering. Here we consider only those that make a meaningful and robust attempt at quantifying uncertainty as serious contenders, which restricts us essentially to methods based on statistical and sound probabilistic principles. For such systems it is easy to show that the inverse problem should be tackled in preference of solving sets of forward problems, which rely on experimentally measured parameter values, and which typically are associated with levels of uncertainty that are, it appears, rarely propagated in forward analyses. The best we can make out the elusive “CELLMAP” appears to be a fully parameterized model for the (cellular) system under consideration. Taking the predictions of such a system at face-value ignores uncertainty and does not appear a sound way of making progress.

There are statistical procedures which are provably consistent as the amount of data becomes infinite. This is clearly a situation far from reality but it seems advisable

to use these techniques also in situations where data are rare. The alternative would be to use an approach which is provably sub-optimal as data become perfect and abundant in the hope that it does a good job on poor data.

The second set of problems is more interesting, and probably more widespread: there are numerous systems (and models thereof) for which the inverse problem is indeed insurmountable. Here simple solutions simply do not exist (and a “CELLMAP” is sadly lacking). Two obvious attempts at addressing such problems—each with its own set of caveats—include *partial inference* and *model reduction*. While the details of their respective applicability depend crucially on the specific problem, we can make some general statements.

By *partial* or *composite inference* we mean a pragmatic approach that proceeds by either breaking up the problem into sub-systems for which satisfactory inferential solutions might exist, and then stitching the solutions for such subsystems together. This has the disadvantage that any correlations or interdependencies among subsystems are ignored. Nevertheless, techniques such as composite likelihood approaches [14] can help to make progress in inference problems that are not amenable to a comprehensive or holistic analysis. This will, we believe, continue to be a fruitful area for computational statistics.

Model reduction, on the other hand, requires more domain expertise about the system to be investigated. In the simplest case, it could be an effective model, which, for example, ignores some molecular species, if they exist only briefly and transiently. It could also be a model that looks at lower dimensional spatial problems (although this can be fraught with fundamental problems as mathematical solutions to problems in 1D and 2D can be qualitatively different from solutions in 3D).

Either approach, individually or in combination, may be worthwhile exploring in problems in systems biology (developmental biology seems to be replete with problems that pose challenges to inferential techniques), and is preferable to an analysis of corresponding forward problems for fixed parameters, which would mask uncertainty.

In summary, recent years have shown the fundamental new insights that can result from searching for or determining the origins of uncertainty in biological systems. In some cases, it will turn out that uncertainty is merely a nuisance (e.g. if it enters via the experimental procedure), whereas other types of uncertainty either reveal exciting new biological mechanisms (e.g. extrinsic variability typically points to aspects of a biological system that require further investigation), or are fundamental and inalienable aspects of biomolecular dynamics.

Failure to account for uncertainty in the analysis of biological systems (and in particular in reverse engineering tasks) will likely introduce bias and mask interesting biology. On the other hand, uncertainty becomes easier to deal with once we know where and how it arises.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

1. Ale, A., Kirk, P., Stumpf, M.P.H.: A general moment expansion method for stochastic kinetic models. *J. Chem. Phys.* **138**(17), 174101 (2013)
2. Apgar, J.F., Witmer, D.K., White, F.M., Tidor, B.: Sloppy models, parameter uncertainty, and the role of experimental design. *Mol. Biosyst.* **6**(10), 1890–1900 (2010)
3. Barnes, C.P., Silk, D., Stumpf, M.P.H.: Bayesian design strategies for synthetic biology. *Interf. Focus* **1**(6), 895–908 (2011)
4. Bowsler, G.O., Swain, P.S.: Identifying sources of variation and the flow of information in biochemical networks. *Proc. Natl. Acad. Sci. USA* **109**(20), E1320–E1328 (2012)
5. Brenner, S.: Sequences and consequences. *Philos. Trans. Royal Soc. Lond. Ser. B Biol. Sci.* **365**(1537), 207–212 (2010)
6. Camacho, D., Vera Licona, P., Mendes, P., Laubenbacher, R.: Comparison of reverse-engineering methods using an in silico network. *Ann. N Y Acad. Sci.* 1115:73–89 (2007)
7. Cedersund, G., Sameulsson, O., Ball, G., Tegnér, J., Gomez-Cabrero, D.: Optimization in biology parameter estimation and the associated optimization problem. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
8. Chu, Y., Hahn, J.: Integrating parameter selection with experimental design under uncertainty for nonlinear dynamic systems. *AIChE J.* **54**(9), 2310–2320 (2008)
9. Cover, T.M., Thomas, J.A.: *Elements of Information Theory*. Wiley-Interscience, New York (2006)
10. Cox, D.R., Hinkley, D.V.: *Theoretical Statistics*. Chapman&Hall/CRC, New York (1974)
11. Csete, M.E., Doyle, J.C.: Reverse engineering of biological complexity. *Science* **295**(5560), 1664–1669 (2002)
12. de Silva, E., Thorne, T., Ingram, P., Agrafioti, I., Swire, J., Wiuf, C., Stumpf, M.P.H.: The effects of incomplete protein interaction data on structural and evolutionary inferences. *BMC Biol.* **4**, 39 (2006)
13. Erguler, K., Stumpf, M.P.H.: Practical limits for reverse engineering of dynamical systems: a statistical analysis of sensitivity and parameter inferability in systems biology models. *Mol. Biosyst.* **7**(5), 1593–1602 (2011)
14. Fearnhead, P., Donnelly, P.: Approximate likelihood methods for estimating local recombination rates. *J. Royal Stat. Soc. Ser. B Stat. Methodol.* **64**(4), 657–680 (2002)
15. Flassig, R.J., Sundmacher, K.: Optimal design of stimulus experiments for robust discrimination of biochemical reaction networks. *Bioinformatics (Oxford, England)* **28**(23), 3089–3096 (2012)
16. Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B.L.: *Bayesian Data Analysis*, 2nd edn. Chapman & Hall/CRC, Boca Raton (2003)
17. Geva-Zatorsky, N., Rosenfeld, N., Itzkovitz, S., Milo, R., Sigal, A., Dekel, E., Yarnitzky, T., Liron, Y., Polak, P., Lahav, G., Alon, U.: Oscillations and variability in the p53 system. *Mol. Syst. Biol.* **2**, 2006.0033 (2006)
18. Gillespie, D.T.: A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J. Comput. Phys.* **22**(4), 403–434 (1976)
19. Gillespie, D.T.: Stochastic simulation of chemical kinetics. *Annu. Rev. Phys. Chem.* **58**(1), 35–55 (2007)
20. Grima, R.: A study of the accuracy of moment-closure approximations for stochastic chemical kinetics. *J. Chem. Phys.* **136**(15), 154105 (2012)

21. Gutenkunst, R.N., Waterfall, J.J., Casey, F.P., Brown, K.S., Myers, C.R., Sethna, J.P.: Universally sloppy parameter sensitivities in systems biology models. *PLoS Comput. Biol.* **3**(10), e189 (2007)
22. Hartemink, A.J.: Reverse engineering gene regulatory networks. *Nat. Biotechnol.* **23**(5), 554–555 (2005)
23. Hilfinger, A., Johan, P.: Separating intrinsic from extrinsic fluctuations in dynamic biological systems. *Proc. Natl. Acad. Sci.* **108**(29), 12167–12172 (2011)
24. Huan, X., Marzouk, Y.M.: Simulation-based optimal Bayesian experimental design for nonlinear systems. *J. Comput. Phys.* **232**, 288–317 (2013)
25. Johnson, R., Kirk, P., Stumpf, M.P.H.: SYSBIONS: nested sampling for systems biology. *Bioinform.* **31**, 604–605 (2014)
26. Julier, S., Uhlmann, J.: A general method for approximating nonlinear transformations of probability distributions. Department of Engineering Science (1996)
27. Julier, S., Uhlmann, J., Durrant-Whyte, H.F.: A new method for the nonlinear transformation of means and covariances in filters and estimators. *IEEE Trans. Autom. Control* **45**(3), 477–482 (2000)
28. Kirk, P., Thorne, T., Stumpf, M.P.: Model selection in systems and synthetic biology. *Curr. Opin. Biotechnol.* **24**(4), 767–774 (2013)
29. Komorowski, M., Miekisz, J., Stumpf, M.P.H.: Decomposing noise in biochemical signaling systems highlights the role of protein degradation. *Biophys. J.* **104**(8), 1783–1793 (2013)
30. Komorowski, M., Costa, M.J., Rand, D.A., Stumpf, M.P.H.: Sensitivity, robustness, and identifiability in stochastic chemical kinetics models. *Proc. Natl. Acad. Sci.* **108**(21), 8645–8650 (2011)
31. Kutalik, Z., Cho, K.-H., Wolkenhauer, O.: Optimal sampling time selection for parameter estimation in dynamic pathway modeling. *Biosystems* **75**(1–3), 43–55 (2004)
32. Liepe, J., Filippi, S., Komorowski, M., Stumpf, M.P.H.: Maximizing the information content of experiments in systems biology. *PLoS Comput. Biol.* **9**(1), e1002888 (2013)
33. Liepe, J., Taylor, H., Barnes, C.P., Huvet, M., Bugeon, L., Thorne, T., Lamb, J.R., Dallman, M.J., Stumpf, M.P.H.: Calibrating spatio-temporal models of leukocyte dynamics against in vivo live-imaging data using approximate Bayesian computation. *Integr. Biol* **4**(3), 335–345 (2012)
34. Lindley, D.V.: On a measure of the information provided by an experiment. *Ann. Math. Stat.* 986–1005 (1956)
35. Liu, J.S.: Monte Carlo strategies in scientific computing. Springer, Berlin (2008)
36. May, R.M.: Uses and abuses of mathematics in biology. *Science* **303**(5659), 790–793 (2004)
37. Pugh, E.M., Winslow, G.H.: The Analysis of Physical Measurements. Addison-Wesley series in physics. Addison-Wesley (1966). <http://books.google.co.uk/books?id=vREAAAIAAJ>
38. Sachs, K., Perez, O., Pe'er, D., Lauffenburger, D.A., Nolan, G.P.: Causal protein-signaling networks derived from multiparameter single-cell data. *Science* **308**(5721), 523–529 (2005)
39. Schäfer, J., Strimmer, K.: An empirical Bayes approach to inferring large-scale gene association networks. *Bioinformatics* **21**(6), 754–764 (2005)
40. Schliemann-Bullinger, M., Fey, D., Bastogne, T., Findeisen, R., Scheurich, P., Bullinger, E.: The experimental side of parameter estimation. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
41. Silk, D., Kirk, P.D.W., Barnes, C.P., Toni, T., Rose, A., Moon, S., Dallman, M.J., Stumpf, M.P.H.: Designing attractive models via automated identification of chaotic and oscillatory dynamical regimes. *Nat. Commun.* **2**, 489 (2011)
42. Skilling, J.: Nested sampling for general Bayesian computation. *Bayesian Anal.* **1**, 833–860 (2006)
43. Stoer, J., Bulirsch, R., Bartels, R., Gautschi, W., Witzgall, C.: *Introduction to Numerical Analysis*, vol. 2. Springer, New York (1993)
44. Strong, M., Oakley, J.E., Chilcott, J.: Managing structural uncertainty in health economic decision models: a discrepancy approach. *J. Royal Stat. Soc. Ser. C (Applied Statistics)* **61**(1), 25–45 (2012)

45. Stuart, A.M.: Inverse problems: a Bayesian perspective. *Acta Numerica* **19**, 451–559 (2010)
46. Stumpf, M.P.H., Balding, D.J., Girolami, M.: *Handbook of Statistical Systems Biology*. Wiley, Chichester (2011)
47. Sunnåker, M., Stelling, J.: Model extension and model selection. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
48. Swain, P.S., Elowitz, M.B., Siggia, E.D.: Intrinsic and extrinsic contributions to stochasticity in gene expression. *Proc. Natl. Acad. Sci.* **99**(20), 12795–12800 (2002)
49. Tarantola, A.: *Inverse Problem Theory and Methods for Model Selection*. SIAM, Philadelphia (2005)
50. Thorne, T., Stumpf, M.P.H.: Inference of temporally varying Bayesian networks. *Bioinformatics* **28**(24), 3298–3305 (2012)
51. Toni, T., Tidor, B.: Combined model of intrinsic and extrinsic variability for computational network design with application to synthetic biology. *PLoS Comput. Biol.* **9**(3), e1002960 (2013)
52. Toni, T., Welch, D., Strelkowa, N., Ipsen, A., Stumpf, M.P.H.: Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *J. Royal Soc. Interf.* **6**(31), 187–202 (2009)
53. Tucker, W.: Interval methods. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
54. Vanlier, J., Tiemann, C.A., Hilbers, P.A.J., van Riel, N.A.W.: A Bayesian approach to targeted experiment design. *Bioinformatics (Oxford, England)* **28**(8), 1136–1142 (2012)
55. Waldherr, S., Haasdonk, B.: Efficient parametric analysis of the chemical master equation through model order reduction. *BMC Syst. Biol.* **6**, 81 (2012)
56. Wang, Y., Christley, S., Mjolsness, E., Xie, X.: Parameter inference for discretely observed stochastic kinetic models using stochastic gradient descent. *BMC Syst. Biol.* **4**, 99 (2010)
57. Werhli, A.V., Grzegorzczak, M., Husmeier, D.: Comparative evaluation of reverse engineering gene regulatory networks with relevance networks, graphical gaussian models and bayesian networks. *Bioinformatics* **22**(20), 2523–2531 (2006)
58. Wilkinson, D.J.: Stochastic modelling for quantitative description of heterogeneous biological systems. *Nat. Rev. Genet.* **10**(2), 122–133 (2009)
59. Wilkinson, D.J.: *Stochastic Modelling for Systems Biology*. CRC Press I Llc, Boca Raton (2011)
60. Wilkinson, R.D.: Approximate Bayesian computation (ABC) gives exact results under the assumption of model error. *Stat. Appl. Genet. Mol. Biol.* **12**(2), 129–141 (2013)
61. Xu, T.-R., Vyshemirsky, V., Gormand, A., von Kriegsheim, A., Girolami, M., Baillie, Ketley, G.S.D., Dunlop, A.J., Milligan, G., Houslay, M.D., Kolch, W.: Inferring signaling pathway topologies from multiple perturbation measurements of specific biochemical species. *Sci. Signal.* **3**(113), ra20 (2010)

Chapter 3

Probabilistic Computational Causal Discovery for Systems Biology

Vincenzo Lagani, Sofia Triantafillou, Gordon Ball,
Jesper Tegnér and Ioannis Tsamardinos

Abstract Discovering the causal mechanisms of biological systems is necessary to design new drugs and therapies. Computational Causal Discovery (CD) is a field that offers the potential to discover causal relations and causal models under certain conditions with a limited set of interventions/manipulations. This chapter reviews the basic concepts and principles of CD, the nature of the assumptions to enable it, potential pitfalls in its application, and recent advances and directions. Importantly, several success stories in molecular and systems biology are discussed in detail.

Keywords Causality · Causal graphical models · Bayesian networks · Systems biology · Biological networks

3.1 Introduction

The winner of the 2011 ACM Turing Award—the Nobel Prize equivalent in Computing—was Prof. Judea Pearl, a pioneer in probabilistic and causal reasoning. Among many other contributions, the theory of Causal Bayesian Networks that he co-developed is now a standard tool for modeling, inducing, and reasoning with probabilistic causality. Bayesian Networks are at the heart of numerous decision support and expert systems as well as the basis for machine learning algorithms. After

Vincenzo Lagani, Sofia Triantafillou and Gordon Ball have contributed equally to this work.

V. Lagani · S. Triantafillou · I. Tsamardinos
Institute of Computer Science, Foundation for Research and Technology - Hellas,
N. Plastira 100, Vassilika Vouton, GR-700 13 Heraklion, Crete, Greece

G. Ball · J. Tegnér
Unit of Computational Medicine, Karolinska Institutet, Center for Molecular Medicine,
Karolinska University Hospital, L8:05, SE-171 76 Stockholm, Sweden

S. Triantafillou · I. Tsamardinos (✉)
Computer Science Department, University of Crete,
Voutes Campus, GR-700 13 Heraklion, Crete, Greece
e-mail: tsamard.it@gmail.com

© Springer International Publishing Switzerland 2016
L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_3

several decades of heated debate about the possibility of causal discovery without—or with a limited number of—controlled experiments, it seems that consensus converges towards an affirmative answer.

Knowledge of causal relations is paramount in systems biology. Causal modelling goes beyond traditional statistical predictive modelling by allowing one to *predict the effects of actions and interventions* on a system, e.g., the effects of treating with a drug, knocking out a gene, or inducing a mutation in the genome. In contrast, non-causal, predictive modelling is only valid when the system under study is *observed* under the same experimental conditions and not otherwise manipulated. For example, gene expressions A and B may be correlated: observing the expression levels of A allow us to better predict the observed expression levels of B . But, it does not assure us that A regulates B or the opposite. The difference between observing and intervening on a system is essential for understanding causal modelling. If A is the *only* regulator of B , then the two genes are still correlated in a controlled experiment where A is activated or suppressed; in contrast, the correlation disappears in a control experiment where B is activated and suppressed at will by the experimenter, since the effect of A now becomes irrelevant.

To establish causality, one traditionally needs to perform a manipulation (perturbation, intervention) on the system [29]. In contrast, computational Causal Discovery (CD) methods argue that given certain assumptions about the nature of causality one can sometimes induce causal relations from observational data alone or a limited number of manipulations/interventions. One can then analyse archived data, forgoing expensive, time-consuming, or even impossible experiments, and determine certain aspects of the causal mechanisms. Exactly which aspects of the causal structure can be induced depends on the system under study and the available data. Given the complexity of the cell, performing all the possible experiments to establish all relations among every subset of molecular quantities, under all possible experimental conditions, is impractical. CD may provide an alternative.

Causal models (not necessarily induced through Causal Discovery) are already heavily employed in systems biology: biological pathways are a form of causal models that are indispensable in biological research. Pathways are manually assembled from the literature, where relations are established by performing interventions. However, for the most part, such models are informal and have ambiguous semantics for the edges: an edge may imply a direct or indirect causation; a missing edge may imply lack of direct causation or a yet-to-be established relation. In addition, pathways are largely qualitative; the strength and functional form of the causal relations is not represented (some exceptions exist, such as well-characterized metabolic pathways annotated with flux equations [100]). In contrast, models induced with CD methods have specific formal causal semantics as well as quantitative information that enables quantitative predictions.

In the rest of this chapter we present the basic concepts of CD, focusing on the fundamental underlying assumptions and discussing its limitations and potential pitfalls. We also present selected applications of CD in systems biology, demonstrating the potential of this exciting field.

3.2 The Nature of Causality

3.2.1 Definition of Causality

We use the notation $A \rightarrow B$ to denote our belief that “ A causally affects B ” (or, “ A causes B ” for brevity). But, what exactly does this mean and how should it be interpreted? Most of CD employs a probabilistic notion of causality. A and B should denote two well-specified variables (interchangeably: measurable quantities, features) which are measured on a population of objects, such as two protein concentrations in human T-cells. We consider simultaneous measurements of these variables in a random sample of the population, from which we can estimate their joint probability distribution. Thus, for the purposes of this chapter, the data are assumed cross-sectional: snapshots of the state of a cell without regard for the time of measurement.

$A \rightarrow B$ denotes the fact that if an experimenter intervenes and changes the values of A , the distribution of B will also change. This statement is inherently probabilistic: Average-Cigarettes-Smoked-Per-Day causally affects Presence-Of-Cancer-by-Age-60 because the distribution of Presence-Of-Cancer-by-Age-60 changes and the people with value “Yes” become more prevalent. To a single individual, that means that the probability of her getting cancer increases. Yet, causality as presently defined is still deceptively simplistic. A may be causally affecting B only in a given context, e.g., in the presence of another protein C . Thus, a better definition is probably that $A \rightarrow B$ if there is conceivable intervention involving *only* A , and a context of some other variables that are held constant, such that the distribution of B changes (relative to the distribution of B when the context is the same but A is not intervened upon). The “intervention” may be just a thought experiment, technically impossible with present technology. Yet, it has to be theoretically plausible. For example, the statement $Cancer \rightarrow Protein$ is arguably undefined: we cannot intervene on the state of the cell to make it cancerous without affecting anything else in the cell. Such semantically vacuous statements often arise when variables that refer to different abstraction levels are modeled together. In this case Cancer, a quantity that refers to the cell as a whole, and the concentration of a protein are defined on a different time and spatial scale.

Finally, notice that the concept of causation is required to define “intervention”, used in the definition of causation; our definition is recursive! To break the vicious cycle, notice that intervention requires defining causality from outside the system (the experimenter) to within the system; causality as defined regards causal effects within the system. In other words, given that we understand what it means for an experimenter to intervene in a population of cells, we can define the causal relations among molecular cell quantities. We can proceed with using causality in an operational way, the same way humanity is doing statistics while still arguing about the philosophical issues of the semantics of probabilities.

3.2.2 Direct Causation

We'll need to distinguish between direct and indirect causality. We'll say A is *directly* causally affecting B relative to variables in the set O , if A remains a cause of B even when all other modeled variables' values are held fixed. Direct causation is relative to the observed variables. A hormone may directly causally affect a gene when nothing else is observed, but indirectly affect it when the status of cell membrane receptors are observed.

3.2.3 Quantitative Causality

Relations $A \rightarrow B$ are qualitative and useful for human inspection and visualization in the form of networks. But, quantitative relations are necessary to make quantitative predictions. If there are two or more direct causes of B ($A \rightarrow B$ and $C \rightarrow B$), then in general we cannot consider the relations independently. This is necessary because A and C jointly determine the values of B . In general, we can model the values of B with the *structural equation*:

$$B = f(Pa(B), U)$$

where $Pa(B)$ are the direct causes of B (or *parents* of B) and U represents all other non-modeled causes. If A and C are the only parents, then $B = f(A, C, U)$. The difference from a non-structural equation is the special role of the left-hand-side: the value of B is set (determined) by the values of PA_B and U and not vice versa: B cannot be moved to the right-hand-side. This special role of the left-hand-side is equivalent to dictating that if we intervene on the values of the right-hand-side, the left-hand-side may change, but not the other way round. The structural equation is not symmetrical. Also notice that the equation is deterministic! However, the presence of unknown values of U introduces uncertainty into the equation and induces a probability distribution of the values of B . The form of function f is important. A few examples follow, where $I(\bullet)$ is the indicator function taking values 1 when the argument holds and 0 otherwise:

- $B = a \cdot A + b + \varepsilon$, B 's concentration always increases with the same rate as A increases. This is an example of a linear relation (strictly speaking, if $b \neq 0$ it is called an affine function). The term $\varepsilon = \sum_{i \in U} U_i$ is the effect of all unmeasured causes of B ; it is not measurement noise.
- $B = a \cdot I(A > 100 \text{ and } C > 100) + b + \varepsilon$, B 's concentration follows a baseline of b , and level $a + b$ when both A and C are larger than 100. Thus, in order to discover this relation one must observe or impose values of A and C larger than 100.
- $B = a \cdot (A - 100)^2 + b + \varepsilon$, B 's concentration decreases as A increases if $A < 100$, and increases as A increases if $A > 100$. The rate of increase or decrease of B (as

A takes different values) is $2aA$ (the derivative of the equation) and thus it is not constant but depends on the values of A .

Linear relations are perhaps the easiest to discover, step functions (as in the second example) require observing the system in a suitable range of the parameters, and non-linear functions (as in the last example) require more sophisticated modeling approaches.

3.2.4 Necessary/Sufficient/Contributory Causes

Similarly to the distinction between necessary and sufficient conditions in logic, causes are also distinguished among necessary and sufficient, with the additional category of contributory causes:

- *Necessary*: a necessary cause is such that the effect will always imply the cause, but the cause does not imply the effect. For example, passing some course implies that you sat the examination, but sitting the examination does not imply that you will pass.
- *Sufficient*: a sufficient cause is such that the cause always implies the effect, but the effect does not imply the cause. For example, burning hydrogen and oxygen will always result in water, but the presence of water does not imply combustion.
- *Contributory*: a contributory cause is any other cause which may result in an effect, but of itself is neither necessary nor sufficient. For example, an intoxicated driver may result in a crash, but intoxication does not imply a certain crash, and neither does a crash always imply the driver was intoxicated.

The majority of cases for which causal analysis is useful concern contributory causes. Single necessary causes are usually relatively easy to identify: these are the “low-hanging fruit” for which experiment and intuition will readily recover causality without recourse to causal analysis. Conversely, a collection of mildly contributory causes is a harder problem to identify, and one for which causal methods applied to large datasets prove useful.

3.3 Basics of Causal Discovery Algorithms

3.3.1 Causal Graphical Models

A graphical representation is a useful way of quickly seeing the structure of a complex system. Intuitively, a set of causal relationships $A \rightarrow B$ can be readily represented as a graph where nodes represent quantities and directed edges represent causal relationships. Particularly, the formalism of Probabilistic Graphical Models (PGM) helps us

define, in a more principled way, the mathematical characteristics of causal models. We will mainly consider the Causal Bayesian Network (CBN) and Bayesian Network (BN) frameworks, as these are some of the most known and widely employed PGMs.

3.3.2 Causal Bayesian Networks

A CBN consists of a graph $G = \{V, E\}$ and a parameterization θ . The set V of nodes represents the observed (modeled) quantities (a.k.a. variables), while E is a set of directed edges $A \rightarrow B$ indicating *direct*¹ causal relationships (where A is the “cause” and B is the “effect”). The graph must consist only of directed edges and contain **no cycles** (Directed Acyclic Graph, DAG). For any node A , any node that can be reached by following directed edges is a *descendant* (*effect*), and any node from which A can be reached is an *ancestor* (*cause*). The direct causes of a node are named *parents*, while its directed effects are named *children*. A directed path consists of a sequence of nodes where each node, except the first one, is the direct effect of its predecessor in the sequence. An undirected path is a sequence of nodes where each pair of subsequent nodes is connected by an edge without regard to the direction of the edge. Whenever an undirected path $\{A \rightarrow C \leftarrow B\}$ exists with two incoming edges into C , the node C is called *collider on this path*.

The parameterization θ defines the joint probability distribution of data generated by a system with the causal structure of the network. The parameterization quantifies the functional form of the causal relations among the variables. Adding a parameterization allows us to express whether relationships are linear or not, the effect size of each interaction, and in general to make quantitative inferences. For a discrete joint distribution (all variables being discrete) there is one parameter θ_i for each combination of values of the variables:

$$P(V_1 = v_{i_1}, \dots, V_n = v_{i_n}) = \theta_i \quad (3.1)$$

A major assumption of the CBN framework is the **Causal Markov Condition**: *each node of V is independent of its non-descendants (non-effects) given its parents (direct causes)*. In other words, the Causal Markov Condition asserts that indirect causes or confounded quantities do not provide additional information for a variable, once the values of the direct causes are known. Notice that, *effects* of V may provide additional information, even when all direct causes of V are known. The Causal Markov Condition allows us to connect the causal structure (network) with the distribution parameters. By the chain rule in probabilities we obtain:

$$P(V_1 = v_{i_1}, \dots, V_n = v_{i_n}) = \prod_j P(V_j = v_{i_j} | V_1 = v_{i_1}, \dots, V_{j-1} = v_{i_{j-1}}) \quad (3.2)$$

¹Direct causation is defined in the context of all other modeled variables, i.e., a causal relation mediated by none of the observed variables.

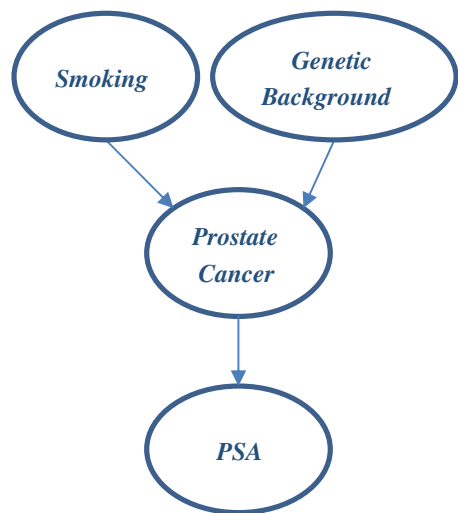
Let's assume without loss of generality that each variable in the equation above is listed after its parents, i.e., if $i < j$, then V_j cannot be a parent of V_i (in other words, we assume the variables are topologically sorted). Notice that this is always possible for DAGs. By using the Causal Markov Condition in the above equation, we obtain:

$$P(V_1 = v_{i_1}, \dots, V_n = v_{i_n}) = \prod_j P(V_j = v_{i_j} | Pa(V_j) = pa_{qj}) \quad (3.3)$$

That is, due to the Markov Condition for each variable V_j , all the variables in the conditioning part have disappeared except the parents of V_j , denoted with $Pa(V_j)$. The quantity pa_{qj} denotes the joint combination of values of the parents of variable V_j . The causal structure now dictates the form of the joint probability distribution, by entering the equation in the form of the parent sets $Pa(V)$ for each variable and imposing a factorization of the joint distribution. Employing Eq. (3.1) to represent an arbitrary distribution with n binary variables requires $2^n - 1$ parameters θ_i to be specified. However, using Eq. (3.3), we only need to represent the distributions $P(V_j = v_{i_j} | Pa(V_j) = pa_{qj})$ for each variable. If a causal system is sparse, e.g., each variable has at most 3 parents, then we need $(2 - 1) \cdot 2^3$ parameters for each such conditional distribution. So, in total, we need at most $n \cdot (k - 1) \cdot k^p$, where n is the number of variables, k the maximum number of values of each variables, and p the maximum number of parents of a variable: *knowledge of the structure of the causal network, assuming it is sparse, allows an exponential reduction in the number of distribution parameters required, and hence the dimension of the parameter space.*

CBNs also assume **Causal Sufficiency**, which corresponds to asserting that there are no external variables which are causes of two or more variables within the model. These common causes are called *confounders*. The Causal Sufficiency assumption implies that the following sub-graph is not present in the causal system under study: $X \leftarrow L \rightarrow Y$, where X and Y are modeled, and L is unobserved and not

Fig. 3.1 A simple graphical model depicting the (supposed) causal relationships among smoking, genetic background, cancer and Prostate-Specific Antigen (PSA). The parameterization of the distribution associated with this network is described in the text



modeled. A truly causally sufficient model is in practice hard to construct, especially in molecular biology where hundreds of thousands of molecular quantities may be confounding factors. Causal Sufficiency is one of the most restrictive assumptions of causal discovery. Fortunately, there exist more advanced generalizations of CBNs that admit latent confounding variables (see Sect. 3.6).

Let's employ an example in order to better explain the concepts above. Figure 3.1 portrays the DAG of a small CBN. In this example, the probability of developing *Prostate Cancer* is influenced by both *Genetic Background* [94] and *Smoking* [48], while the presence of prostate cancer increases the probability of deregulation in the expression of the Prostate-Specific Antigen (*PSA*, [50]). Let's suppose that all variables are binary, which means each variable can assume a value in the set $\{1, 0\}$. Regarding the *semantics* of the values, "1" means, respectively, *deregulated* for *PSA*, *harmful* for *Genetic Background*, *yes* for *Smoking* and *present* for *Prostate Cancer*. In each variable the value "0" negates the meaning of value "1". We can now parameterize this simple model as follows:

$$\begin{aligned} P(\text{Smoking} = 1) &= \pi_{\text{Smoking}} \\ P(\text{Genetic Background} = 1) &= \pi_{\text{Genetics}} \\ P(\text{Prostrate Cancer} = 1 | \text{Smoke}, \text{Genetics}) &= a_1 \cdot \text{Smoke} + a_2 \cdot \text{Genetics} + a_0 \\ P(\text{PSA} = 1 | \text{Prostate Cancer}) &= a_4 \cdot \text{prostrate Cancer} + a_3 \end{aligned}$$

In this parameterization, having a harmful genetic background and being a smoker are modelled as random events, whose respective probabilities are π_{Smoking} and π_{Genetics} . Coefficients a_0, a_1, a_2 quantify the extent to which *Smoking* and *Genetic Background* affect the probability of developing cancer, while a_3, a_4 quantify how *Prostate Cancer* changes the probability of *PSA* being deregulated.

Notice that all causal relationships are *probabilistic* (non-deterministic), i.e., *Smoking* and *Genetic Background* increase the *probability* of developing *Prostate Cancer*, while the presence of cancer *may* deregulate *PSA* expression. The probabilistic nature of the model is due to the existence of a number of factors $U_i, i \in U$ (e.g., physical activity, diet, medications, etc.) which influence the model's quantities but are not measured. However, recall that the *Causal Sufficiency* assumption requires that no external factor simultaneously influences two or more elements of the model; this means that each variable can be affected by multiple U_i , but each U_i can affect only one variable.

3.3.2.1 Inference in Causal Bayesian Networks

If the CBN is known (this includes both the structure and the parameterization), *any probabilistic inference is possible*. In particular, any predictive or diagnostic query of the form "what is the probability V_i will take/has taken value j given that we observed certain values for other variables" is possible. Without loss of generality let's assume we observed $V_1 = v_1, \dots, V_k = v_k$ and we would like to compute the conditional probability that $V_{k+1} = v_{k+1}$:

$$\begin{aligned}
P(V_{k+1} = v_{k+1} | V_1 = v_1, \dots, V_k = v_k) &= \frac{P(V_1 = v_1, \dots, V_k = v_k, V_{k+1} = v_{k+1})}{P(V_1 = v_1, \dots, V_k = v_k)} \\
&= \frac{\sum_{V_{k+2}, \dots, V_n} P(V_1 = v_1, \dots, V_k = v_k, V_{k+1} = v_{k+1}, \dots, V_n = v_n)}{\sum_{V_{k+1}, \dots, V_n} P(V_1 = v_1, \dots, V_k = v_k, V_{k+1} = v_{k+1}, \dots, V_n = v_n)} \quad (3.4)
\end{aligned}$$

where each index v_i runs on all values in the domain of variable V_i . Each term in the sums is computed by Eq. (3.3). Let's resume the example presented in Fig. 3.1, and assume that a specific patient has a deregulated *PSA*, is a smoker, and his *Genetic Background* is not harmful. A clinician may be interested in assessing the probability that the patient has *Prostate Cancer*, which can be evaluated by applying in turn Eqs. (3.3) and (3.4):

$$\begin{aligned}
&P(\text{Cancer} = 1 | \text{Smoke} = 1, \text{PSA} = 1, \text{Genetics} = 0) \\
&= \frac{P(\text{Cancer} = 1, \text{Smoke} = 1, \text{PSA} = 1, \text{Genetics} = 0)}{\sum_{pc=0,1} P(\text{Cancer} = pc, \text{Smoke} = 1, \text{PSA} = 1, \text{Genetics} = 0)} \\
&= \frac{P(\text{Genetics} = 0) \cdot P(\text{Smoke} = 1) \cdot P(\text{Cancer} = 1 | \text{Smoke} = 1, \text{Genetics} = 0) \cdot P(\text{PSA} = 1 | \text{Cancer} = 1)}{\sum_{pc=0,1} P(\text{Genetics} = 0) \cdot P(\text{Smoke} = 1) \cdot P(\text{Cancer} = pc | \text{Smoke} = 1, \text{Genetics} = 0) \cdot P(\text{PSA} = 1 | \text{Cancer} = pc)} \\
&= \frac{(a_1 \cdot 1 + a_2 \cdot 0 + a_0) \cdot (a_4 \cdot 1 + a_3)}{(a_1 \cdot 1 + a_2 \cdot 0 + a_0) \cdot (a_4 \cdot 1 + a_3) + [1 - (a_1 \cdot 1 + a_2 \cdot 0 + a_0)] \cdot a_3}
\end{aligned}$$

Assuming that *Smoke* sensibly increases the probability of cancer ($a_1 = 0.2$, $a_0 = 0.01$) and that *PSA* has a high sensitivity ($a_4 = 0.9$, $a_3 = 0.1$), the patient has a high probability (0.727) of having *Prostate Cancer*. A similar inference would have also been possible in the case information regarding *Genetic Background* was not available, though the sums would have contained more terms (the number of terms grows exponentially with the number of unobserved variables). In general, inference is in the worst case an NP-complete problem, however efficient exact or approximate algorithms do exist [71]. Thus, a CBN can predict/diagnose any variable or set of variables given the values of any other set of variables. It is like having trained an exponential number of predictive models, one for each variable subset as predictors. This is a key factor that has made CBNs popular in (clinical) decision support systems where one may have a varying and limited number of observations for each patient.

The graph of a CBN can also provide all the (conditional) independencies implied by the Causal Markov Condition. If *faithfulness* is assumed (see Sect. 3.4.1 for a definition of faithfulness) the graph can also provide all (conditional) dependencies. In other words, by examining the graph, one can determine which variables are conditionally or unconditionally correlated. The property that connects the topology of graphical/causal structure with the concept of conditional (in)dependence is called *d-separation*; two sets of variables A, B (such that $A \neq B$) are conditionally independent given a third set $C \subseteq V \setminus \{A, B\}$ if and only if they are d-separated by C in G . Formally, d-separation is defined as follows: A, B are said to be d-separated given a third set C if there is no undirected path U such that (i) every collider in U has a descendent in C and (ii) no other nodes in C is in U . Intuitively, we can think about d-separation as a criterion that tells us if the “flow of information” between two variables A and B is interrupted or not. For example, variables *Smoking* and *PSA* in Fig. 3.1 are d-connected when conditioned on the empty set (the “flow” of information can freely transfer from *Smoking* to *PSA* through the node *Prostate Cancer*), but they are

d-separated when we condition on *Prostate Cancer*, since knowing its value makes the information provided by *Smoking* superfluous in order to predict *PSA*. On the other hand, *Smoking* and *Genetic Background* are d-separated in absence of a conditioning set, but they become d-connected when conditioning on *Prostate Cancer*, *PSA* or both. In fact, in the latter case we condition on all the colliders (or their descendants) in the undirected path between *Smoking* and *Genetic Background*. Note that *Smoking* and *Genetic Background* are independent, but knowing both the values of *Smoking* and *Prostate Cancer* gives us some information on the value of *Genetic Background*, and thus the two variables are not independent anymore. Hereafter we will denote with $dep(A, B|C)$ the presence of a conditional dependency between variables A and B given the set C , while $indep(A, B|C)$ will denote independence.

Finally, CBNs can make inferences unique to causal models: *they can predict the effects of interventions/manipulations/changes in experimental conditions*. Given a CBN we can determine the effect of knocking out a gene, the effect of administering a drug, or the effect of changing any quantity modeled in the network. Conceptually, such inferences are straightforward. The effect of the experimenter on the system that sets the values of a variable V_k to v , removes the effect of any other variable to V_k . This is modeled by removing all incoming causal edges to V_k and setting $P(V_k = v) = 1$ and $P(V_k = v') = 0$, for $v \neq v'$ in the conditional probabilities associated with the graph. The edge removal is called *graph surgery*; in the resulting graph V_k will have no parents. The new joint probability distribution can now be computed with Eq. (3.4), and hence any probabilistic query about the effect of the intervention can also be computed. Interventions that deterministically set the values of some specific variables are called *hard interventions*. When interventions have a chance of not being effective, they are called *soft interventions*. In this case, the intervention does not completely remove the causal effect of all other quantities, and thus, a different treatment is necessary where the probability of effective intervention is also modeled. In addition, when an intervention is not specific to a quantity but may affect other quantities too, the intervention is called a *fat-hand intervention* and also requires different modeling techniques [25].

The main reason for causal modeling and discovery is exactly to enable the prediction of the effect of our actions onto the system. Causal models are the only types of models that enable such inferences. Statistical causal models, such as CBNs perform such inferences without modeling the underlying physical phenomena and mechanisms of causality, while other models such as Ordinary Differential Equations directly model these mechanisms.

3.3.2.2 Dropping the Causal Semantics: Bayesian Networks

It is rarely the case that a CBN can be constructed completely from prior knowledge. Typically, such models have to be learnt from data by algorithms that make numerous assumptions (see Sect. 3.4.1 for a discussion). In cases when the causal assumptions are not to be trusted, and the structure or parameters of the CBN is also not trusted, one may still use the Bayesian Network framework without the causal semantics. Similarly to CBNs, BNs consist of a DAG and a parameterization, but do not make any

causal claims or causal predictions. The Causal Markov Condition can be substituted with the Markov Condition: each node is independent of its non-descendant given its parents. Note the substitution of “direct causes” with “parents”, since in BNs the term “cause” does not make sense anymore.

An edge $X \rightarrow Y$ in a BN should be interpreted strictly from a probabilistic viewpoint: *the edge denotes that X provides unique information for Y (possibly given some other variables) and vice versa*. The direction of the edges should also not be interpreted causally; the direction is only employed to combine edges into paths and determine implied probabilistic dependencies and independencies from the network with the d-separation criterion. All probabilistic inferences possible with CBNs are also possible with BNs, except for causal inferences: if the causal semantics are dropped, one is not entitled to employ a BN to predict the effect of manipulations into the system. BNs may predict our future observations based on past observations, but not the effects of our actions.

Notice that, Bayesian Networks are only loosely related to some other Bayesian statistical approaches in this volume, for example Bayesian model selection (See Chapters [37, 49, 101]). Bayesian Networks treat probability in a “Bayesian” way, i.e., to represent measures of belief (in contrast to the frequentist interpretation of probabilities). They also make heavy use of the Bayesian Theorem to make inferences. Both of these characteristics justify the term Bayesian. Bayesian model selection also treats probabilities as measure of belief; in particular, Bayesian model selection uses probability distribution on the set of possible models to express the a priori belief on their validity (typically, favoring simpler models). However, Bayesian Networks serve to model and make inferences about joint distributions, while Bayesian model selection aims to select the statistical model that achieve the best trade-off between fitting the data and abiding to the prior beliefs.

3.4 Causal Discovery Approaches

The main goal of Causal Discovery algorithms is reconstructing the network of causal mechanisms underlying a given system, given a dataset D . The dataset D is usually composed of a set of n observations measured over m variables. Such causal learning algorithms have already proven useful to biologists as shown below in Sect. 3.5.

Unfortunately, reconstructing a Causal Model from data is not an easy task. Several algorithms have been proposed in the last few decades, and all of them consist of two stages: firstly, an appropriate causal graph is found, and secondly a parameterization is estimated in accordance with the graph structure. While the second stage is relatively straightforward (given a suitable assumption about the functional form of the causal relationships), identifying the correct causal graph has proven to be NP-hard [14]. So far, two main approaches have been developed for reconstructing the graphs of Causal Models, namely the *Constraint-based* and the *Score-based* (also known as *Search-and-Score*) approach. The basic principles of the two main approaches for learning CBN are the following:

- *Score-based*: first introduced in [78], these methods are based on a score function $S(G|D)$ measuring the fit of the graph to the data, while at the same time favoring simpler structures. A prototypical score function is the Bayesian Information Criterion (BIC), defined as $BIC(G|D) = |\theta| \cdot \ln(n) - 2 \cdot \log(L(G|D))$, where $\log(L(G|D))$ is the log-likelihood of the graph given the data, and $|\theta|$ is the number of model parameters [93]. The score function is usually combined with a search-heuristic that explores the space of possible graphs. A typical heuristic is the greedy one: start with the empty graph (no edges) and then add, reverse or delete the edge that maximally increases the score of the network (i.e., the fit to the data) at each step.
- *Constraint-based*: this approach relies on estimating some of the conditional (in)dependencies in the data distribution P from the data D through performing hypothesis tests of conditional independence. Typically, for discrete variables the X^2 or the G^2 tests are employed [63], while for continuous variables testing the partial linear correlation are employed based on the Fisher z-transformation [28]. The results of the hypothesis tests constrain the graph to reconstruct: in the resulting graph G , two variables X, Y should be d-connected given Z if and only if $indep(X; Y|Z)$ in the data. In fact, it can be proven (assuming faithfulness, as defined in Sect. 3.4.1) that two variables are connected by an edge if and only if there is no set of variables Z , such that $indep(X; Y|Z)$. Constraint-based methods usually start with a fully connected, undirected graph and progressively remove edges whenever a new conditional independency is discovered [98].

Typically for a given dataset there will be multiple solutions (i.e., networks) that are Markov equivalent, i.e., they imply the same set of conditional independencies and thus cannot be distinguished based on testing independencies on the data. Under typical scoring functions, these networks receive equivalent scores. Intuitively, each such network provides an equally good causal explanation for the data. *The issue of Markov Equivalence in learning causal structures is a point that an analyst should keep in mind.* The set of equivalent networks has some invariant characteristics, e.g., edges and directions upon which all solutions agree, and some variant characteristics upon which different solutions disagree. Even when all causal assumptions hold, *the analyst is warranted to make claims only about the invariant characteristics.* Fortunately, for CBNs the representation of the set of equivalent networks is compact: they can be represented with another type of network called the Completed Partially DAG (CPDAG) or essential graph [13] and the invariant characteristics can be identified from this graph. Particularly, CPDAGs contain two types of edges, directed and undirected. The first type represents arcs that are similarly (invariantly) oriented in all Markov Equivalent solutions, while the latter represents edges whose orientation varies among equivalent networks. In more complicated causal formalisms discussed in Sect. 3.6, the set of equivalent solutions cannot be compactly represented. See also Chapters [37, 49, 101] for further discussions on model identifiability and (Bayesian) model selection.

Causal Discovery algorithms can also be used for variable selection, i.e., identifying the smallest subset of quantities that can provide the optimal prediction or

diagnosis for an outcome variable of interest T (*this is equivalent to molecular signature identification*). Under certain conditions, the set of variables optimally predicting the value of an outcome or molecular quantity T is what is called the Markov Blanket of T : the set of direct causes, direct effects, and direct causes of direct effects [108]. Algorithms that can identify the Markov Blanket of T from data without knowledge of the underlying CBN exist and have proved to be some of the most effective signature identification algorithms from biological data [3]. *Importantly, these theoretical results connect molecular signatures for prediction or diagnosis with the causal structure of the system under study: the most predictive quantities have a specific causal interpretation.*


3.4.1 A Discussion on Some Typical Causal Discovery Assumptions and Practical Issues

We now focus in more detail on the most common assumptions of typical causal discovery algorithms and discuss their implication in the context of causal discovery in biological systems.

(Causal) Markov Condition: in a DAG G each node is independent from any non-descendant (non-effect) given its parent (direct causes). This condition formalizes our “common belief” about how Causality operates, i.e., indirect causes or confounded effects do not provide additional information, once the direct causes are known. For example, in the network $X \leftarrow W \rightarrow Y \rightarrow Q \rightarrow R$, we expect that once we know Y (the direct cause of Q), neither X (a confounded variable) nor W (an indirect cause) provide additional information for Q . Notice that, observing the effect R of Q still provides additional information for the value of Q . Interestingly, while the Causal Markov Condition is (explicitly or implicitly) accepted and employed “all the time in laboratory, medical and engineering setting” ([98], p. 38), whether it holds in the sub-atomic systems studied by quantum physics it is still under debate [70]. This assumption is what allows the algorithms to discover direct causal relations and drop edges from the causal network being reconstructed. While relatively uncontested in practice, *the Causal Markov Condition may appear to be violated due to measurement error (see below).*

Acyclicity: CBN and other PGMs assume that no node in the graph can be a cause of itself, either directly or through other variables. CBNs are not able to represent feedback loops, and in some biological applications this limitation can be quite restrictive. However, some approaches have been developed that do not require this assumption [41]. Typically, formalisms that admit the presence of feedback cycles assume only linear relations. In the presence of both non-linear relations and feedback chaotic phenomena may arise that significantly complicate the problem and the applicable algorithms. Thus, one must substitute one assumption for the other to make causal discovery possible.

Fig. 3.2 Simple example of a feed-forward network



Causal sufficiency: no pair of nodes shares a common, unmeasured cause. In statistical terms, we assume that there are no latent confounders that may introduce correlations that are not explained by the measured variables. Specifically, consider the network $X \leftarrow L \rightarrow Y$. Because L is a confounder of X and Y , we expect the latter pair of variables to be correlated (dependent). If L is not observed, it is not modeled in the network. There are only two networks with variables X and Y that entail a dependency and fit the data: $X \rightarrow Y$ and $X \leftarrow Y$. Both of them are Markov equivalent and correctly represent the data distribution. But, their causal semantics are wrong: X does not cause Y nor vice versa. The network with the correct causal semantics is “ X (no edge) Y ”. There is no way to correctly simultaneously represent both the probabilistic semantics and the causal semantics of the network without admitting new, unobserved variables in the network. Causal Sufficiency is one of the most restrictive assumptions in CBNs particularly for systems biology where there are millions of possible molecular quantities that may be confounding the observed quantities. For this reason some PGM frameworks have been recently developed (e.g., Maximal Ancestral Graphs [85]) that generalize CBNs to admit and reason over the presence of hidden confounders.

Faithfulness: a distribution P is faithful to a DAG G if it entails all and only the conditional independences implied by G . This assumption turns out to be important particularly for the efficiency of Causal Discovery algorithms, in order to search and identify all solution networks. One interpretation of faithfulness is that the set of conditional independencies is *stable* under infinitesimal perturbations of the data distribution [79]. For example, consider the following feed-forward gene network (Fig. 3.2):

X regulates Y both directly as well as indirectly through Z . The two paths for regulating Y may be competing with each other: X up-regulating Y directly, X up-regulating Z which in turn down-regulates Y . If the causal effects of each regulation are just so finely tuned it is possible that the association between X and Y completely disappears even though X causes Y . Such fine tuning of the parameters of the distribution seems unlikely (and it is infinitely unlikely under certain assumptions, see [98], p. 66) and leads to independences that are unstable: they become dependencies if the parameters of the distribution are slightly perturbed. Faithfulness dictates that this fine tuning is not present in the data distribution. Thus, whenever X causes Y in a network directly or indirectly or through multiple causal paths, we assume the variables are dependent.

Faithfulness seems innocent at first glance, but there are several pitfalls. First, in practice a distribution may be faithful but “close to unfaithfulness”; in the example above, the association between X and Y may not disappear completely but may be

too small to be detected with typical sample sizes. Second, while fine-tuning of the parameters seems unlikely to occur by chance, there is evidence that natural selection leads to systems which may be unfaithful; in particular, the presence of negative feedback cycles may lead to associations that disappear [22]. Deterministic relations also violate faithfulness! It seems that randomness (i.e., natural occurring perturbations) is required to allow causal discovery, which is philosophically intriguing to say the least. For example, consider the network $X \rightarrow Y \rightarrow Z$, where X and Y are deterministically related, e.g., they always have equal values. In that case, X provides for Z the same information as Y and so $\text{indep}(Y; Z|X)$ holds which is not entailed by the Markov Condition. There are algorithms that do not assume faithfulness for learning CBNs [54]. However, simultaneously dropping the acyclicity and faithfulness assumptions requires sophisticated theory and algorithms [41].

There are some additional assumptions that are often not declared explicitly, but that should be carefully taken into consideration:

No measurement error: the variables are measured without error. This is a subtle assumption that is required to learn CBNs, often not realized by practitioners who apply these techniques. In other words, to allow causal discovery we need to assume that the variance of the measurements of a variable X stems from our uncertainty about (marginalizing over) all other causes of X , and is not due to measurement error. Consider the effect of measurement error: let's assume we measure $X' = X + e_X$, $Y' = Y + e_Y$, $Z' = Z + e_Z$, where the last terms are the measurement noise terms. Let's assume the true structure is $X \rightarrow Y \rightarrow Z$. Thus, based on the Causal Markov Condition we expect that $\text{indep}(X; Z|Y)$. However, we observe the noisy versions of the variables, so what we test instead is $\text{indep}(X'; Z'|Y')$. If the variance of e_Y is larger than e_X , it may turn out that X' does provide additional information for Z' given Y' . This is equivalent to the Causal Markov Condition being violated. A more relaxed assumption is that all error terms have the same variance, which would lead to noisy versions of the variables that still maintain the same independencies as the true, underlying network involving only the original variables. *This observation is particularly important for measurements by biotechnologies that do have significant measurement error, such as micro-array gene expression data, where gene expression may have very different variance of measurement errors.*

Effect of data transformations (discretization, averaging): as above, this issue regards the connection between the actual quantities that we are modeling and the quantities measured and contained in the data. For example, it is common for a practitioner to discretize the data before applying a causal discovery method. However, depending on the discretization, the set of dependencies and independencies in the transformed data distribution may be changed compared to the original [61, 68]. Again, this may appear as a violation of the Causal Markov Condition on the transformed data. *Another important case of potentially harmful transformation is that of averaging. Averaging takes place in almost every mass-throughput technology.* For example, in micro-array gene expression data one tries to induce causal relations and networks among *gene expressions in a single cell*, e.g., that $X \rightarrow Y$. However, *what is measured in the data are the average expressions \bar{X} and \bar{Y} of X and Y in millions of cells.* The independencies of a network is $X \rightarrow Y \rightarrow Z$ defined on the single-cell

quantities X , Y , Z are not necessarily the same as the independencies on the averages \bar{X} , \bar{Y} , and \bar{Z} [15]. *This observation favors causal discovery from single cell data* (or in general, measurements that are not averaged) versus other biotechnologies.

No selection bias and case-control studies: A basic assumption for causal discovery is that the samples are not selected for inclusion in the data based on an effect of the modeled variables. Let's consider the case where two genes X , Y regulate the size of the cell Z : when both genes are high the cell is larger with high probability. In addition, we assume the two genes to be independent from each other. Thus, the true network is $X \rightarrow Z \leftarrow Y$. Now, let's imagine that a researcher measures these genes in a collection of cells including mostly in large cells (perhaps because small cells are harder to detect and isolate given the available equipment). In the selected population whenever X is high, Y is also high with large probability: the two gene expressions are correlated in the selected population. *This correlation is an artifact of the data sampling* and not present in the general cell population. *Cytometry data is a particular type of data with possible selection bias as an effect of the gating process and classification to different cell types. Another striking example of selection bias is case-control data.* In case-control studies, half the samples (cases) have been selected for inclusion based on the effect (disease) of the modeled variables. In the previous example, let us change the semantics of Z to being the presence or absence of a disease and X , Y two independent causes of disease. In all cases of disease, when X is high, Y is high with high probability, so they appear correlated even though they are not correlated in the general (unselected) population. Epidemiologists try to alleviate these spurious correlations by *matching* cases and controls based on some of the variables (age, gender, race, etc.). If cases and controls are matched in the example above, the spurious association between X and Y would disappear. However, matching cannot be achieved at a molecular level for every variable (e.g., gene expression) that is modeled and so one has to be particularly careful with causal discovery in case-control data. Some methods for learning causal networks [7] try to account for selection bias introduced by unmatched case-control study design.

It should also be noted that standard Causal Discovery algorithms assume that samples are independent and identically distributed (i.i.d.) and that they are all measured under the same experimental conditions and at same the point in time (cross-sectional data). Other algorithms exist for dealing with other types of data and information, e.g., data measured under different experimental conditions [19, 51], in different points in time [30] or for co-analyzing data in the context of prior causal knowledge [8].

Finally, *practical issues* also determine the success of causal discovery:

Statistical errors: statistical errors in the results of the conditional independence tests, or equivalently statistical fluctuations in the score of networks to the data may result in learning networks or relations that are wrong. In fact, in a large network it is almost certain that some parts of the network will be erroneously induced. Robustness against statistical errors and sample sizes depends on the learning method. Employing the most appropriate hypothesis testing procedure or scoring function for the given data is paramount. Inappropriate tests or score functions may introduce systematic reasoning errors. For example, if functional relations are non-linear but linear

hypothesis tests are used, some conditional (in)dependencies may not be detectable even with large sample sizes. Methods for assessing the reliability of each feature of the network (e.g., presence of an edge or an edge direction) do exist and should be employed. Some of them employ bootstrapping, i.e., learning with resampled data. However, notice that bootstrapping provides the confidence *given by the method for a given feature* (e.g. edge in the network); bootstrapping does not provide an absolute confidence for the feature. For example, if a method systematically reports a false edge because it employs inappropriate tests for the specific data, bootstrapping will also return high confidence on this edge.

Non-linear relations: non-linear relations in continuous data present particular problems to causal discovery. For example, consider the case when two quantities do not interact, unless a third quantity is present in a sufficient concentration. If the data do not contain samples where this third quantity is indeed in high concentration, the causal relation will be undetected. Equivalently, for discrete data, a correlation may be present only for specific values of the variables that never appear in the dataset and hence will not be detected by any algorithm.

3.5 Causal Discovery in Systems Biology: Success Stories

Despite the philosophical, theoretical, and algorithmic problems described above, CD can work when applied with care, and assumptions, technicalities and limitations are duly taken into consideration. The following success stories from *systems biology* provide evidence for this.

3.5.1 *Inferring Causal Relationships Among Genotype and Quantitative Traits*

In recent years, computational methods have been introduced for *identifying causal relationships among genetic characteristics and quantitative traits in observational data*. These methods were named differently by their respective authors, e.g., Likelihood-based Causality Model Selection (*LCMS*, [91]) or *Trigger* (Transcriptional Regulation Inference from Genetics of Gene ExpREssion, [12]). For simplicity, hereafter we will collectively refer to all these methods as Causal Quantitative Trait Loci (CQTL) algorithms.

Specifically, CQTLs methods attempt to reconstruct the causal interaction between a genome marker L and two quantitative traits, namely T_1 and T_2 , all measured in the same segregating population. Each quantitative trait can represent the expression value of a given gene, a quantitative phenotype, or any other continuous measurement on the population of interest. CQTL's cornerstone assumption is that *a statistical association between the genetic marker L and the traits of interest must denote a*




Fig. 3.3 Possible Causal Models among a genetic marker L and two quantitative traits T_1 and T_2 given that the Mendelian Randomization assumptions hold and that all quantities are associated with each other. The node H represents one or more hidden confounders

causal effect of L on the associated trait. This assumption is justified by the theory developed in the context of Mendelian Randomization [23, 47]. In a nutshell, Mendelian Randomization methods assume that the random re-composition of the genome during conception can be considered equivalent, from a statistical point of view, to the randomization procedures performed during Randomized Control Trials (RCTs). Consequently, any statistical association between the genetic information and the traits/phenotype of interested cannot be affected by latent confounders, i.e., must denote a causal association.² All Quantitative Trait Loci (QTL) studies [66] are based on Mendelian Randomization and its assumptions.

Thus, given that (a) the Mendelian Randomization assumptions hold (i.e., T_1 and T_2 cannot cause L), (b) the Causal Markov and Faithfulness conditions hold as well, and (c) L , T_1 and T_2 are found in the data all statistically associated with each other (i.e., the following dependencies hold: $dep(L, T_1 | \emptyset)$, $dep(L, T_2 | \emptyset)$, $dep(T_1, T_2 | \emptyset)$), then only a very restricted number of causal structures (see Fig. 3.3) are admissible. Each causal model is represented as a CBN, where the node H represents one or more unknown, latent confounders.

Can we further screen out the models presented in Fig. 3.3 and identify the unique, actual causal structure that generated the data at hand? Using the d-separation

²Linkage disequilibrium, pleiotropic effects and other factors can invalidate the Mendelian Randomization approach; these issues are better explained later in the text.

criterion above and assuming faithfulness, if L and T_2 are not statistically associated given T_1 (i.e., $\text{indep}(L, T_2|T_1)$ holds), then the true causal model must be $L \rightarrow T_1 \rightarrow T_2$.

In more detail, each model where L and T_2 are connected through a direct edge is incompatible with $\text{indep}(L, T_2|T_1)$, since T_1 cannot d-separate T_2 and L . This leaves models (1), (7) and (8) as the only possible candidates. In the two latter models, T_1 is a collider in the path $L \rightarrow T_1 \leftarrow H \rightarrow T_2$, and thus conditioning on T_1 makes L and T_2 dependent. Thus, the single causal model in agreement with all the assumptions and (in)dependencies encoded in the data is model (1).

Similarly, $\text{indep}(L, T_1|T_2)$ holds only when the true underlying causal model is $L \rightarrow T_2 \rightarrow T_1$.

Thus, the causal relationships among a genetic marker and two quantitative traits can be identified, in principle, by assessing whether a limited number of conditional (in)dependencies hold in the data. Particularly, studies focusing on large panels of genomics markers/quantitative traits (e.g., Genome Wide Association Studies) can opportunistically apply CQTL methods on each possible triplet of the form $\{L, T_1, T_2\}$, and potentially discover a large number of causal relationships.

The first theoretically-sound algorithm able to identify, under a well-defined set of assumptions, causal triplets $L \rightarrow T_1 \rightarrow T_2$ where L is known to be “uncaused” was introduced by Cooper in 1997 [18].³ The first applications of CQTL methods in biology appeared only a decade later: the work presented in Schadt et al. [91] was one of the first studies demonstrating CQTLs effectiveness on a specific biological problem.

Particularly, Schadt and co-authors investigated the causal relationships between a genome-wide panel of markers (L), transcript abundance levels in the liver (T_1) and obesity-related traits (T_2) in mice. They referred to model (1) and (2) in Fig. 3.3 as *Direct Causal model* and *Reactive Causal model*, respectively, while all other models were collectively indicated as the *Independent Causal model*. A model selection procedure, namely LCMS (Likelihood-based Causality Model Selection), was employed for identifying the most plausible causal model for each triplet {genomic marker, transcript abundance level, obesity related trait}. The LCMS procedure belongs to the class of Search-and-Score algorithms, and employs the Akaike Information Criterion (AIC, [1]) as the score metric: $AIC = 2k - 2\ln(\mathbf{L})$, where k is the number of parameters of each model and $\ln(\mathbf{L})$ its log-likelihood.

Chen and co-authors [12] developed a Constraint-based CQTL algorithm. Particularly, they demonstrated the *Causal Equivalence Theorem*, i.e., if the Faithfulness and Causal Markov Condition hold, then:

*The causal relationship $L \rightarrow T_1 \rightarrow T_2$ exists and there are no hidden variables causal for both T_1 and T_2 if and only if the following three conditions hold: $\text{dep}(L, T_1|\emptyset)$, $\text{dep}(L, T_2|\emptyset)$, and $\text{indep}(L, T_2|T_1)$.*⁴

³Statistical algorithms for identifying and quantifying mediation effects were known even earlier [58, 97]. However, these algorithms usually assume some particular (linear) distributional model and “fell short of providing a general, causally defensible measure of mediation” [80].

⁴Notably, the “Causal Equivalence Theorem” is identical to the LCD procedure presented in [18].

The authors employ the Causal Equivalence Theorem in order to derive a method, namely Trigger, which provides probability values $\hat{p}_{1,2}$ and $\hat{p}_{2,1}$ for the causal structures $L \rightarrow T_1 \rightarrow T_2$ and $L \rightarrow T_1 \rightarrow T_2$, respectively.

More recently, Millstein and co-authors have proposed another Constraint-based CQTL algorithm, the Causal Inference Test (CIT, [67]), which evaluates a larger set of (conditional) dependencies and independencies than Trigger. Particularly, the following conditions must be satisfied for accepting the Direct Causal Model:

- CIT Condition 1: L and T_2 are associated
- CIT Condition 2: L and T_1 are associated given T_2
- CIT Condition 3: T_1 is associated with T_2 given L
- CIT Condition 4: L is independent from T_2 given T_1

A p-value for each of the four CIT conditions can be calculated by applying a suitable statistical test of (conditional) dependency, while the maximum among the four p-values, namely p_{DCM} , is employed as a global statistic for assessing if the four conditions can be jointly accepted. A global p-value p_{RCM} for the Reactive Causal model $L \rightarrow T_2 \rightarrow T_1$ can be derived in a similar way.

Once p_{DCM} and p_{RCM} have been provided, the CIT procedure applies the following rules to distinguish among the possible causal models:

1. If $p_{DCM} < \alpha$ and $p_{RCM} > \alpha$, then the Direct Causal Model is accepted
2. If $p_{DCM} > \alpha$ and $p_{RCM} < \alpha$, then the Reactive Causal Model is accepted
3. If $p_{DCM} > \alpha$ and $p_{RCM} < \alpha$, then the Independent Causal Model is accepted
4. If $p_{DCM} < \alpha$ and $p_{RCM} > \alpha$, then no call is made

where α is a threshold for accepting statistical significance (e.g., $\alpha = 0.05$). Interestingly, CIT does not distinguish among the Independent Causal Model and the case when L is not associated with T_1 or T_2 .

CQTL methods have been applied in several studies in order to shade light on specific biological problems. The spread of CQTL methods has also been boosted by the availability of free, open source implementations, whose most notable examples are the R package *cit* (implementing the CIT method), the Network Edge Orienting (NEO) software [4], that implements a score-based CQTLs method, and the R package *qtlnet*, that implements a CQTL algorithm able to take in account and exploit complex correlation structures among multiple traits/phenotypes [72].

A recent example of a successful CQTL study has been presented by Gutierrez-Arcelus et al. [32]: the interaction between DNA sequence, DNA methylation and gene expression was investigated with the CIT method in fibroblasts, T-cells and lymphoblastoid cells extracted from the umbilical cord of 204 babies. This study showed that, when the two alleles of a gene are not equally expressed in a given type of cell, gene expression is mainly regulated by DNA sequence variation, with little or no influence by DNA methylation.

Liu et al. [55] employed the CIT method for disentangling the causal relationships among genome, DNA methylation and Rheumatoid Arthritis. By using the CIT algorithm, the authors found 535 genome—arthritis causal interactions that are mediated

by methylation, out of the initially 4016 initially considered associations between genome markers and the rheumatoid arthritis phenotype.

Some controversial CQTL results have been reported in another publication [44]. In this work the authors studied the genome characteristics and expression profile of leukocyte cells from 284 Moroccan individuals. By applying a basic CQTL method, it came out that the SNP rs11987927 seems to trans-regulate the expression of the ZNF71 gene which, in turn, regulates back the transcript abundance of the MYOM2 gene, i.e. the gene where rs11987927 is located. The authors were not able to show whether this counterintuitive result is trustworthy or is instead due to measurement errors [88] or to other causes (e.g. the presence of feedback cycles).

This last example reminds us that the CQTL approach has, obviously, some limitations. Particularly, the limitations affecting Mendelian Randomization [74] affect as well all CQTL studies. Mendelian Randomization assumes that the choice of the mating partner is not affected by the genome. Population stratification is another possible source of bias for Mendelian Randomization and CQTL studies. It can be the case that allelic frequencies and phenotype distributions vary similarly across different populations, even in absence of any causal relations. Consequently, artificial genome-phenotype associations could be detected if the population under study is composed by different sub-populations. Biological redundancy and adaptation to unfavorable genetically-determined phenotypes can hide genome-phenotype causal interactions. Markers that are physically close to each other on the genome tend to be highly associated (a phenomenon known as linkage-disequilibrium) and these associations can lead to the false identification of causal markers that are merely close to the real cause of the phenotype. Highly co-linear (associated, correlated) quantities are close to determinism and violations of Faithfulness (see Sect. 3.4.1 above). Genomic markers can have pleiotropic effects, i.e., simultaneously affecting several traits. If the effect of the pleiotropic marker on each trait is small, it may be necessary to jointly consider all the traits in order to detect the marker-traits causal associations. Furthermore, genomic modifications driven by reverse transcription [9] may ingenerate cases where the observed genomic profiles are actually influenced by the traits under study. Finally, to the best of our knowledge, all CQTL methods developed so far assume that all genome markers follow the same genomic model (usually the additive or co-dominant one), even if assuming the wrong genomic model can lead to a decrease of statistical power [5]. Methodological approaches have been proposed in order to mitigate the effect of some of these limitations, particularly in order to detect causal markers in condition of strong linkage disequilibrium [73] and pleiotropic effects [115].

Despite these limitations, CQTL studies have proven to be able to identify actual causal relationships in a number of different biological context. The main factors enabling CQTL effectiveness are:

Incorporation of prior, biological knowledge: the (apparently) innocuous information that “nothing causes L ” is actually pivotal in order to dramatically reduce the number of possible causal models. This means that CQTL methods explore a very small space of possible models thanks to the adoption of Mendelian Randomization assumptions.

Opportunistic approach: CQTL methods are usually applied on a large number of triplets, and whenever a Direct or Reversal causal model cannot be identified, they forgo making a decision. Therefore, the CQTL approach can be thought of as “explorative analysis”, useful for discovering novel causal associations which can be subsequently experimentally validated.

“Local” causal discovery: a number of difficulties arise when Causal Discovery methods are applied with the intent to learn a complete network of all direct causal relations, i.e., the CBN among all quantities in the data. Errors in statistical inferences can “propagate”, and erroneously orientate edges even in distant regions of the reconstructed network. Conversely, the CQTL approach focuses on a small system formed by solely three quantities, and thus they do not suffer of the issues arising when large networks are induced.

Causal Sufficiency is not assumed: the CQTL approach is “robust” with respect to the presence of latent confounders: no unmeasured variable can affect the association between L and any of the two traits (given the Mendelian Randomization assumption), while if the two traits are both affected by the same latent confounder then the CQTL algorithm will simply forgo making a decision.

Computational feasibility: CQTL algorithms require performing a relatively limited number of statistical (conditional) association tests. Efficient implementations of CQTL algorithm can be easily realized, and CQTL can be applied on hundreds of thousands of triplets in a reasonable time.

Future developments for CQTL methods seem to move in the direction of data integration for network reconstruction. The CIT algorithm was originally proposed as a method for reconstructing causal interaction networks. The QTLnet algorithm [72] tries to reconstruct the interaction network among genome markers and multiple traits. Cai et al. [10] have developed a Structural Equation Model method, namely the Sparsity-aware Maximum Likelihood (SML) algorithm, for reconstructing gene regulatory networks by exploiting genetic perturbations. Finally, in a recent review [90], the author points out that causal triplets provided by CQTL methods can be used for deriving priors for (Causal) Bayesian Network reconstruction algorithms.

3.5.2 *Reconstructing Protein Signaling Pathways*

Co-ordination of complex cellular activities requires a well-orchestrated propagation of information. In living organisms, this information is transmitted across cells through chemical signals which enter the cell and cause a cascade of chemical, spatial and physical modifications of intracellular compounds. This procedure is broadly described with the term **cell signaling**, and a cascade of responses to a certain extra-cellular stimulus is generally called a **signaling pathway**, though many argue that presenting a signaling pathway as an isolated set of responses to a specific stimulus may be too simplistic.

Such pathways are typically reconstructed by manually synthesizing pathway components. Each pathway component is discovered through the aggregation of several studies examining the relationship in question under different experimental designs.

Signaling pathways are usually represented as graphs, where the nodes represent participating compounds and the edges represent direct causal links. Different shapes and colors are used to denote different types of participating molecules, and different edges are used to discriminate different types of causal influence.

Bayesian networks, being able to capture both causal and probabilistic relations in multivariate systems, seem fitting to model and quantify signaling pathways. In a ground-breaking paper published in 2005, Sachs et al. [89] applied a Bayesian network learning algorithm to reconstruct a known signaling pathway in T-cells.

The authors used multi-parameter flow cytometry data measuring 11 phosphorylated proteins and phospholipids – all known participants in T-cell signaling – under 9 different experimental conditions in naïve cd4+ T-cells. A score-based algorithm for learning Bayesian networks from a mixture of observational and experimental data [19] was then employed to infer the causal structure and the joint probability distribution of the measured variables.

Each experimental condition included a general or target-specific stimulatory condition, sometimes coupled with a target-specific inhibitor. In total, 5 activators and 5 inhibitors were used. All perturbations were modeled as “ideal” interventions [79] (i.e., hard interventions, not fat hand interventions), where the concentrations of the target molecules are set solely by the manipulation procedure (i.e. the selected inhibitor/activator completely determines the value of the target variable).

The data were discretized into 3 bins, representing “low”, “basal” and “high” concentration values, using an algorithm designed to preserve the joint distribution of the variables [33] before being used with the BN learning algorithm. To ensure statistical robustness, the algorithmic process was repeated 500 times with random initial graphs. The output model included only edges present in more than 85 % of the resulting graphs.

The returned network consists of 17 edges and is impressively similar to a consensus signaling pathway manually curated from the literature. Out of the 17 edges identified, 15 edges represent causal links that are well-established in the literature and 2 represent causal links that are not well-established but have been reported at least once. The algorithm failed to discover 3 edges that were expected based on the literature review. However, were they included, these edges would create feedback cycles, which cannot be modeled with Bayesian networks. The causal direction of identified edges was correct, with the exception of a single arc that was found reversed (Fig. 3.4).

To further evaluate the validity of the predicted relations, the authors performed an experiment to test one of the causal links that was found by the algorithm but was not sufficiently backed up by the literature. Specifically, the model included a direct causal link from Erk to Akt, a connection previously reported only in colon cancer cells [31, 114]. The model entails that a perturbation of Erk will influence the abundance of Akt, while it will have no effect in the abundance of PKA.

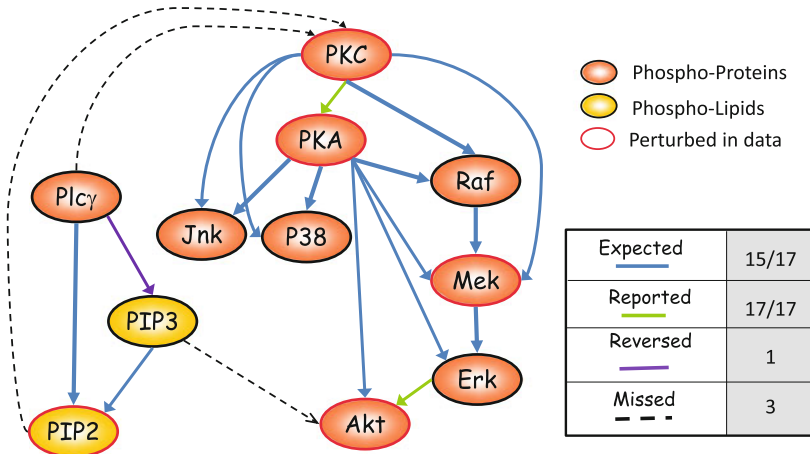


Fig. 3.4 Network inferred from flow cytometry data. The network is a consensus average of 500 high-scoring networks. Only edges present in more than 85% of the networks are included. Out of 17 edges, 15 are well established in the literature and 2 are reported but not well established. One of the edges is found reversed. The resulting network missed three edges that were expected based on the literature review. Figure from [89]

The authors validated this by inhibiting Erk with a suitable siRNA. True to the model's prediction, Akt activity was reduced ($p < 9.4 \times 10^{-5}$), while PKA activity remained uninfluenced ($p < 0.28$).

Despite the impressively accurate pathway reconstruction and the experimental validation of a previously unknown predicted arc, to the best of our knowledge, this paper remains the only case study of Bayesian network learning for automatic network reconstruction. To understand the reasons automatic causal discovery is still sparsely used in bioinformatics, let us discuss the main factors enabling causal discovery in flow cytometry data:

Network perturbations. An important factor in the success of this method is the inclusion of network perturbations, which are particularly important for correctly identifying the directionality of arcs. To test the significance of including experimental data sets, the authors test the algorithm on a data set consisting of 1200 samples measured without intervention. The resulting network contains only 8 out of the 18 expected edges (compared to 15 when the complete data set is used). In addition, all identified edges are undirected, demonstrating the significance of experiments in identifying causal relations. Nevertheless, we do note that the set of perturbations is still quite limited compared to the full set of experiments required to fully generate the structure without the use of CBN methodology.

Large sample size. Bayesian network learning methods require large sample sizes, while typical experimental designs in molecular biology are usually limited to producing just enough samples to ensure the technical soundness of the procedure. Flow cytometry, measuring the abundance of proteins in single cells, results in hundreds

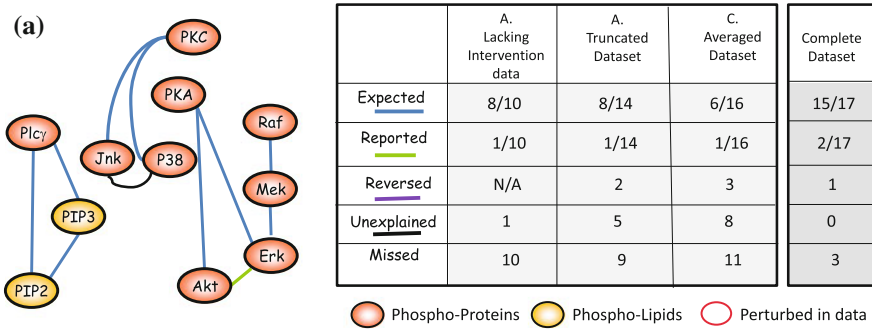


Fig. 3.5 Networks inferred from: **a** A data set consisting from observations alone. **b** A data set consisting of 420 randomly selected samples from the original data set. **c** A data set consisting of 420 data points, each of which is an average of 20 randomly selected samples from the original data set. In all three cases, the resulting network is far less accurate compared to the one resulting from the complete data set. We can therefore infer that the inclusion of experiments, the large sample size and the lack of averaging effects are crucial for accurate network reconstruction. Figure from [89]

of data points in each experiment, enabling the detection of causal relations in noisy multivariate data. The authors show the importance of large sample sizes by applying the same algorithmic procedure on a truncated version of the original data consisting of 420 randomly selected samples. The resulting network is shown in Fig. 3.5b. It consists of 14 edges, out of which only 8 are expected and only 1 is reported.

Single cell measurements. A key obstacle in applying Bayesian networks in molecular biology data is that the measurements are usually averages of quantities in cell tissues. Using averaged measurements for Bayesian network learning is known to be problematic [15], since the correlation structure of measured quantities may not be preserved. Flow cytometry measurements are single cell measurements, and are therefore suitable for this type of inference. To illustrate this point, the authors simulate a western blot data set over the same variables by selecting at random 20 data points at a time and averaging them, creating a data set of 420 samples in total. The resulting network, shown in Fig. 3.5c, displays a further decline in accuracy: Out of 16 edges, only 6 belong to the expected ones.

Causal sufficiency. In this work, the authors aim to provide a proof-of-concept of the use of Bayesian network in analyzing multivariate flow-cytometry data by reconstructing a well-studied pathway in mammalian T-cells. In doing so, the authors pick 11 compounds in the cell that are *not* confounded (in the context of the set of measured compounds), thus satisfying one of the most difficult assumptions of Bayesian networks, that of causal sufficiency. While the authors do not test how decisive this factor is for the success of the method, it is well known that violation of the causal sufficiency assumption causes errors that propagate through the network. Over the past few years, there has been a growing body of work on causal models for causally insufficient systems, some of which are discussed in Sect. 3.6. However, these models are for the most part developed and disseminated in the machine learning community, and remain fairly unknown in the field of molecular biology.

Overall, several attractive features of the flow cytometry technology render it an ideal test-bed for causal Bayesian network learning. Compared to other high-throughput molecular biology techniques, flow cytometry data have vast sample sizes, do not suffer the unwelcome effects of averaging, and samples can easily be perturbed with in-vitro, close-to-ideal interventions. Unfortunately, flow cytometry technology can only measure up to approximately 20 variables simultaneously, limited by the number of distinguishable fluorescents. This number prevents the measurement of all variables participating in known pathways, let alone the numerous cellular compounds for novel pathways. However, the recently developed technique of **mass cytometry**, where antibodies are tagged with rare isotopes instead of fluorescents, allows measuring up to 30 variables, with a theoretical limit of circa 60 variables [75]. Moreover, the demonstration of the problematic effects of using averaged data along with the development of novel technologies has resulted in growing availability of single-cell genomic data [83, 111], promising a bright future for automatic causal discovery in Bioinformatics.

3.5.3 Estimating Causal Effects in High-Dimensional, Observational Data: The Intervention Calculus when the DAG Is Absent Approach

Identifying cause-effect relationships is one of the main goals of Causal Discovery methods. However, in some cases assessing whether a causal relationship holds is not sufficient, and one also desires to quantify the size of the causal effect. For example, once it has been established that a gene regulates a particular protein, it may also be relevant to know what variation should be expected in the level of the protein's abundance (effect) for a given variation in the level of the expression of the gene (cause).

Estimating the size of a causal effect is not a trivial task, although it becomes feasible when the true causal structure is known. Pearl [78] proposed a technique,

named “do-calculus”,⁵ which, given a DAG and a suitable parameterization, allows estimating the magnitude of the causal effect between any pair of variables X, Y modeled in the DAG. Unfortunately, in almost all biology-related, real-world problems the actual underlying causal structure is not known, and its reconstruction is often prohibitive, as discussed in Sect. 3.4.

Recently, Maathuis and co-authors [56] proposed a method for estimating a lower bound on the size of the causal effect between two quantities by using a worst-case analysis. Their method, namely *IDA* (Intervention calculus when the DAG is Absent), has at least two appealing features: (a) it is able to estimate causal effects’ lower bounds solely on the basis of observational data, i.e., without requiring data from experimental perturbations, and (b) can scale up to high-dimensional settings involving thousands of variables.

The basic idea underlying the IDA algorithm is the following: first, let’s assume that the underlying causal mechanism that has generated the data can be represented as a DAG, and that no latent confounders are present (i.e., we assume causal sufficiency). Then, the size of the causal effect $X \rightarrow Y$ between any pair of quantities included in the data can be estimated with the following steps:

1. Identify the CPDAG P that best fits the distribution of the data at hand. Recall from Sect. 3.4 that a CPDAG is a compact representation of the set of DAGs that are Markov equivalent, i.e., the set of DAGs that cannot be distinguished among each other solely on the basis of the available (observational) data. P can be identified by applying any suitable Causal Discovery method, e.g., the PC algorithm [98].
2. Calculate the effect size $ES(X \rightarrow Y)$ for the causal relationship $X \rightarrow Y$ separately for each DAG represented by P . The minimum absolute value among these effect sizes is the lower bound for the effect size of the causal relationship $X \rightarrow Y$.

The apparent simplicity of the IDA algorithm hides an insidious technical issue: the number of DAGs included in P can become intractable even in the case of small systems (e.g., a few tens of measured quantities). For this reason, IDA exploits some sophisticated theoretical results in order to avoid a complete enumeration of the DAGs included in P , while ensuring the correctness of the final results. Moreover, IDA assumes that the data follow a multivariate normal distribution. This assumption is not strictly necessary for the general soundness of the algorithm, but leads to a great simplification of the calculations, since multivariate normality implies linearity of the causal effects. Under the multivariate normal distribution assumption the effect size ES of the causal relationship $X \rightarrow Y$ does not depend by the specific value of X and can be expressed as:

$$ES(X \rightarrow Y) = E(Y|do(X = x + 1)) - E(Y|do(X = x))$$

⁵Explaining the details of the do-calculus is beyond the scope of this chapter. Interested readers can refer to Pearl’s original publication.

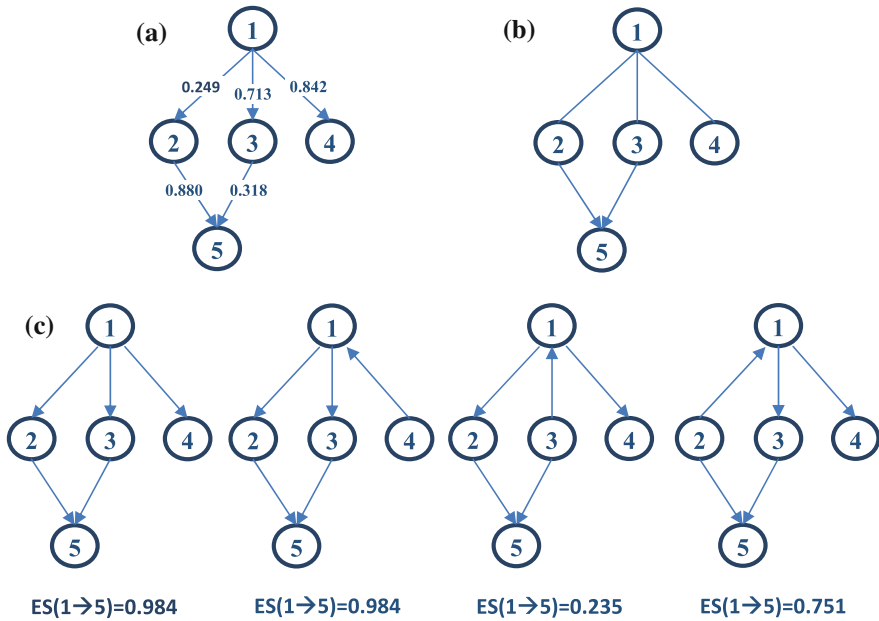


Fig. 3.6 Graphical representation of IDA operation. **a** Example causal network involving five nodes. Causal effects are assumed to be linear, with weights specified on each edge. **b** CPDAG reconstructed by the PC algorithm from 1000 samples simulated from the example causal network. Undirected edges denote arcs that are reversible. **c** DAGs corresponding to the reconstructed CPDAG. For each DAG the effect size of the causal relationship $1 \rightarrow 5$ is reported, as calculated with the do-calculus. The minimum among these values (0.235) is a lower bound of the real effect size. All simulations were performed with the R package *pcalg* [46]

where $E(Y|do(X = x))$, in the language of the do-calculus, represents the expected value of the random continuous variable Y if the value of X is forcefully set, through an external intervention, to a fixed value x over the whole population. If all quantities are scaled in order to have zero mean and unitary standard deviation, $ES(X \rightarrow Y)$ would represent the expected variation of Y for a variation of X equal to its standard deviation.

It should also be noted that IDA can be considered a conservative algorithm, performing a “worst case scenario” analysis, since it returns the minimum absolute value among the calculated size effects. Figure 3.6 shows a graphical representation of the operation of IDA.⁶

The main drawback of the IDA algorithm is that it is based on a set of assumptions that are unlikely to hold in real settings, particularly Causal Sufficiency and multivariate normality. Overall, it is not well understood how the results of the algorithm may change whenever one or more of these assumptions is violated.

⁶An implementation of the IDA algorithm is available in the R package *pcalg* [46].

Despite these limitations, IDA has proved to be effective in identifying and quantifying causal relationships from observational data. In a subsequent, ground-breaking publication [57], Maathuis and co-authors applied IDA on two different sets of observational data: a compendium of expression profiles of *Saccharomyces Cerevisiae*, and the set of simulated gene expression data from the DREAM4 competition *In Silico Network Challenge* [60]. For both sets of data, m “true” causal relationships were estimated and quantified through gene knock-out experiments, while q “predicted” causal relationships were obtained by applying the IDA algorithm *on the observational data*. For both the *Saccharomyces Cerevisiae* and the DREAM4 data the sets of true and predicted causal effects had an overlap statistically significantly larger than the one that can be expected by random guessing. The overlap was statistically significant for different values of q and m . Moreover, when contrasted against two state-of-the-art correlation-based algorithms, (the Lasso and Elastic Net regressions [116]), the IDA algorithm largely outperformed both methods in correctly ranking putative causal relationships; *in fact, the correlation-based algorithms’ predictions were only as good as random guessing*. The importance of these results was highlighted in an editorial in the same issue of *Nature Methods* [11].

An additional application of the IDA algorithm on another real-world problem was also reported [45]. In this work the researchers employed a slightly modified version of IDA (able to deal with binary variables) in order to identify the factors causally influencing the level of general health perception in a sample of spinal cord injury patients. The results of the study confirmed, once more, the capability of IDA in identifying and quantifying causal relationships from observational data.

The factors enabling effective causal discovery with the IDA approach are the following:

Worst case analysis: IDA provides a “worst-case” estimation of the causal effects. This means that only causal relationships strongly supported by the data will be retrieved.

Opportunistic approach: similarly to the CQTL algorithms, IDA is an explorative analysis whose main scope is identifying novel causal relationships, rather than confirming existing ones.

Ranking of putative causal associations: causal associations discovered by IDA are associated with their respective effect size. This means that researchers can rank the putative causal relationships provided by the IDA algorithm according to their estimated effect sizes, and eventually retain/experimentally validate only the top ones.

Finally, it is worth noting that some extensions of IDA were recently published. Le and co-authors presented a version of IDA modified to detect and quantify microRNA/mRNA causal relationships [53]. The Causal Stability Ranking (*CStaR*) method [99] employs the IDA algorithm and a re-sampling based stability selection method [65] to identify, out of a list of possible candidates, the factors that causally influence a given outcome.

3.6 Future Directions

In the previous sections, we presented some introductory concepts related to causality and causal discovery. We also briefly presented (Causal) Bayesian networks, which are one of the main tools for causal discovery without randomized control experiments. Finally, we discussed some of the most prominent and successful applications of causal discovery in the field of molecular biology. Despite years of research in the field of causal discovery and the increasing availability of public data, the applications remain limited and are often contrived examples rather than methods of global applicability. In this section, we explain some of the reasons thereof, and discuss recent developments in causal discovery that may help tackle some of the problems in applied causal discovery, and present some future directions for a unified, robust and integrative approach in causal discovery.

Admitting Latent Confounding Factors: The theory of Bayesian networks relies on the assumption of causal sufficiency, i.e. that no two variables included in the model shares an unobserved common cause (latent confounder). In most real scenarios, this assumption is somewhat arbitrary, since the possibility of a latent confounder can rarely be excluded [87]. The presence of latent confounders is a common source of error in the output of Bayesian network learning algorithms, and an even more common source of criticism and mistrust for causal discovery.

Over the past few years, however, several causal models that do not rely on the assumption of causal discovery have been developed. **Semi Markov causal models** (SMCMs, [103]) are causal models that implicitly model hidden confounders using bi-directed edges. Like Bayesian networks, SMCMs consist of a joint probability distribution over a set of variables and a causal graph over the same set of variables. The graph is an **acyclic directed mixed graph**, where nodes represent variables and edges represent causal relations: A directed edge (\rightarrow) denotes a direct causal relation (in the context of variables included in the model), while a bi-directed edge (\leftrightarrow) denotes that the variables in question share a latent common cause. Two variables can share both a directed and a bi-directed edge. Under the causal Markov condition and faithfulness, conditional (in)dependencies entailed in the distribution correspond to graph properties of the graph according to the criterion of m-separation, an extension of d-separation in BNs. While obtaining a parameterization of a mixed graph is possible for discrete variables [27, 86] there exists no algorithm that can reverse-engineer a semi-Markov causal model from data.

Maximal ancestral graphs (MAGs, [85]) constitute a different approach in modeling causality in causally insufficient systems. Maximal ancestral graphs are ancestral mixed graphs, meaning they contain no directed or almost directed cycles: An almost directed cycle occurs when $A \leftrightarrow B$ and there exists a directed path from A to B . Every pair of variables A, B in an ancestral graph is joined by at most one edge. The orientation of this edge represents (non) causal ancestry: A bi-directed edge $A \leftrightarrow B$ denotes that A does not cause B and B does not cause A , but (under the faithfulness assumption) the two share a latent confounder. A directed edge $A \rightarrow B$ denotes causal ancestry: A is a causal ancestor of B . Thus, if A causes B (not




Fig. 3.7 Causal insufficiency. **a** A causal Bayesian network over variables A, B, C, D, L . **b** The semi-Markov causal model over the (causally insufficient) set of variables, A, B, C, D . **c** The maximal ancestral graph over the same variables

necessarily directly in the context of causal variables) and the two are also confounded, there is an edge $A \rightarrow B$ in the corresponding MAG. Figure 3.7 illustrates an example of a marginal SMCM and MAG for the same underlying causal Bayesian network. Some features of a MAG are not identifiable from the joint probability distribution alone. Classes of MAGs that correspond to the same probability distribution form a Markov equivalence class. The FCI algorithm [98, 112] is a constraint-based algorithm that can learn all the invariant features of Markov equivalent MAGs from passive observational data. The algorithm is shown to be sound and complete.

Admitting Feedback Cycles: Another long debated assumption of causal Bayesian networks is acyclicity; i.e., the lack of feedback loops in the system under study. While some may argue that causal processes are acyclic *over time*, in many practical settings we only have cross-sectional, non-temporal data, hopefully having reached equilibrium. Particularly in molecular biology feedback is a well-known regulatory mechanism and thus, acyclicity a problematic assumption.

To address this shortcoming of causal Bayesian networks, several approaches have been introduced, most of which resort to the parametric assumption of linearity. Richardson and Spirtes are the authors of the first general constraint-based algorithm for learning linear cyclic models, the Cyclic Causal Discovery algorithm [84]. The algorithm however is not complete. Schmidt and Murphy present a method for learning discrete cyclic models [92], but their method heavily relies on experimental data. Moreover, the authors present no theoretical results for their algorithms completeness and identifiability status. Itani et al. introduce generalized Bayesian networks [96], an extension of Bayesian networks for cyclic systems with discrete variables, and present a learning algorithm. The method relies on experimental data to both identify data and to apply BN learning algorithms in data where the cycles are broken by perturbations. All of the methods above employ the assumption of causal sufficiency. Hyttinen et al. present a method for learning linear cyclic model from a series of experiments in causally insufficient systems [41], along with sufficient and necessary conditions for identifiability. Unfortunately, this method also relies on linearity, which is generally known not to hold in biological systems.

Local and Opportunistic Learning: Given the limitations, difficulties, and pitfalls of CD, learning complete large networks may degrade quality of learning and present large computational demands. *Local Causal Discovery* takes a different approach. There are at least two types of causal discovery. The first, pioneered by Cooper and colleagues attempts to identify (all) *marginal* graphs (i.e., representing the

distribution of a subset of the variables) of some special interest. For example, in [18] all triplets leading to a CBN of the form $L \rightarrow T_1 \rightarrow T_2$ when L is known not to have any causes within the system under study are identified. As discussed in Sect. 3.5.1 this work preceded the CQTL studies and was re-discovered independently later. This is the smallest graph that postulates a new causal relation $T_1 \rightarrow T_2$ *without assuming Causal Sufficiency*, due to the prior knowledge that nothing causes L (we do not consider $L \rightarrow T_1$ or $L \rightarrow T_2$ as new interesting causal postulates since if nothing causes L and L is correlated with T_1 or T_2 then the causal relation should hold trivially). When prior knowledge is not available, the smallest marginal graph that postulates a causal relationship without assuming Causal Sufficiency is called a Y-structure and is of the form $X \rightarrow Q \leftarrow Z, Q \rightarrow W$. If this CBN is induced from the data, then $Q \rightarrow W$ even if Causal Sufficiency is violated (the CBN of course also claims $X \rightarrow Q$ but this may not be the case if Causal Sufficiency is violated). Algorithms to identify Y-structures appeared in [59]. Another type of Local Causal Discovery is the reconstruction of focused regions of the underlying causal graph around a variable of interest, e.g., a specific gene, without the need to reconstruct the complete network. The first such method was [62], later receiving more attention in [81, 110]. Such local CD algorithms are closely related to variable selection as the Markov Blanket of a variable is the part of the network relevant for variable selection [3]. The difference between the two types of local causal discovery is that the first learns marginal networks, while the second learns sub-networks. For example, if the true network is $X \rightarrow Y \rightarrow Z \rightarrow W$, and nothing causes X , then the method by Cooper [18] will return 3 triplets: $X \rightarrow Y \rightarrow W, X \rightarrow Z \rightarrow W$, and $X \rightarrow Y \rightarrow W$ corresponding to marginalizing (treating as latent) one variable at a time. The method learning regions in [110] with target Z will return the network $Y \rightarrow Z \rightarrow W$ (if the region is restricted to be only the nodes adjacent to Z). The latter is a sub-graph of the original graph (in general, local discovery may not orient the same edges as global discovery). *Local Causal Discovery forgoes learning complete networks to save computational time or to make more robust inferences with fewer assumptions.* We also use the term *opportunistic learning* to denote all methods that perform a reliability, confidence, robustness estimation of findings and focus only on the findings for which the method is confident on. The CQTL methods presented above heavily use these ideas.

Integrative Causal Analysis: in recent years, the proliferation of publicly-available, on-line data repositories allow the possibility of co-analyzing large amounts of data and information. This is particularly evident in some fields, for example System Biology, where on-line data repositories are well-established [2, 6] and researchers are encouraged to share their raw data along with their results and findings. Typically, however, data from different studies cannot be pooled together naively and be jointly analyzed, even when all studies examine the same biological system. Any difference in recording conventions, study design or experimental procedures requires sophisticated statistical approaches in order to be addressed. A non-exhaustive list of approaches that attempt to address these issues includes Meta-Analysis [76], Transfer Learning [77], Statistical Matching (also called Data Fusion) [24] and Batch-Effect removal [52]. Each of these approaches is characterized by its own scope, advantages

and weaknesses. In general the integrative analysis of heterogeneous datasets is still an open problem and a field of active research.

Integrative analysis from a causal perspective takes a specific form. The key observation in this approach is that all data measuring the same biological system stem from a single causal mechanism. Each study maybe measuring different quantities, under different experimental conditions or sampling methodologies, yet there should exist a causal model that can produce all these datasets. Thus, to co-analyze a collection of datasets coming from heterogeneous studies one searches for a causal model (or all causal models) that simultaneously fit and can explain all data. Over the past few years, several methods for extending causal analysis to the integrative analysis of heterogeneous datasets have been introduced. We collectively refer to these methods as Integrative Causal Analysis (INCA). A major advantage of INCA is that it can model the effect of interventions, e.g. the knock-out of a gene in one dataset and treatment with a hormone in a second one, to enable the co-analysis of datasets over different experimental conditions.

INCA methods can address different types of heterogeneity. Several INCA works have focused on the problem of overlapping variable sets, i.e., co-analyzing data sets that have only a subset of the included variables in common. In this setting the scope of the analysis is usually to infer information regarding the causal mechanism defined over the union of all measured variables. A first pioneering work was published in 2002 by Danks [20], who proposed a two-stage approach consisting in separately learning a Bayesian Network from each study and then using a set of rules for extracting information about the underlying causal structure. Successive methods generally follow a similar two-stage approach, but use more expressive causal models in the first stage (e.g. MAGs) and employ more sophisticated rules that are able deal with conflicts arising from inconsistencies among the models [17, 21, 104, 106].

Studies often differ because they were conducted under different experimental conditions. In this setting, naively pooling data from different studies together can lead to the creation of spurious correlations or to the disappearance of present associations among the measured variables [51]. Several works propose modifications of Search-and-Score and Constraint-based algorithms able to deal with mixtures of observational and experimental data. Cooper and Yoo [19] propose a Bayesian score able to incorporate information about the different experimental settings, while Hauser and Bühlmann [34] investigate the concept of Markov Equivalence in the presence of experimental interventions and propose a learning algorithm on that basis. Eaton and Murphy [25] model interventions as special nodes of the network, and proposed an algorithm that attempts to infer the actual effects of each intervention directly from the data. Constraint-based algorithms for mixtures of experimental data are proposed in [16, 102], but they are limited to specific experimental settings. Sufficient conditions for checking conditional (in)dependencies in data coming from different experiments were proposed in [26, 51]. Other approaches assume specific functional forms for all interactions among variables [39, 41]. These approaches are even able to deal with hidden confounders, but their application is limited by their strict assumptions regarding functional forms among variables. Finally, some

algorithms first learn a provisional causal structure from observational data, and then employ experimental information in order to refine the learned model [35, 64].

A particular type of heterogeneity is obtained when the same information is recorded with different encodings, for example smoking information may be recorded as a binary (yes/no) or a continuous variable (number of packets a day). When a direct conversion is not possible, more sophisticated approaches must be employed [107].

Recent developments in Integrative Causal Analysis focus on co-addressing multiple sources of heterogeneity at the same time: several works attempt to integrate data measured over overlapping variable sets *and* in different experimental conditions [40, 42, 43, 105].

One of the main, unresolved issues in the INCA field is the efficacy of the current methods on real data. While significant efforts have been spent on laying down the theoretical foundations of this field, several algorithmic and methodological improvements are necessary before applying these methods on real data analysis tasks. A first attempt in applying INCA methods on real-world, large datasets has produced evidence that INCA methods can actually provide meaningful results and even outperform current statistical methods [109]. Bridging the gap between theory and practice is crucial for the future of integrative causal analysis.

CD Based on Functional-Form Analysis: So far we have mainly discussed causal discovery methods based on the analysis of conditional (in)dependencies. These methods query the joint probability distribution for (in)dependencies either directly (constraint-based methods) or indirectly (search-and-score methods) to identify all causal structure that fit the data. Recently, a different approach on causal discovery has been developed, one that is based on the exploiting possible asymmetries of causal relations. The methods assume Causal Sufficiency and acyclicity, thus if X and Y are correlated, either $X \rightarrow Y$ or $Y \rightarrow X$. Expressed as structural equations, either $X = f(Y, \varepsilon)$ or $Y = f(X, \varepsilon)$, where the disturbance term ε is the effect of all other factors. It turns out that one can distinguish between the two possibilities if either ε is non-Gaussian, or f is non-linear [36]. While the assumptions of linear relations and Gaussian residual term ε is probably the most common set of assumptions in statistics, it turns out that any departure from these assumptions allows the discovery of the directionality of causation!

A specific case follows. Assume that X and Y are variables and X causes Y in a linear manner, thus $Y = \beta_{YX}X + \varepsilon_Y$, where ε_Y follows a non-Gaussian distribution. Also assume that we have obtained a set of measurements of both X and Y and we want to identify the causal structure of the variables. By assuming linearity, additive disturbance terms ε , and causal sufficiency, we can fit both models using simple linear regression, and obtain estimates for both $\widehat{\beta_{XY}}$ and $\widehat{\beta_{YX}}$. Based on these estimates, we can then calculate the disturbances ε_X and ε_Y for both models. These disturbances will be independent with each other if the fitted model is the correct one, and dependent if the fitted model is the reverse one. A graphical depiction of this principle for uniform distributions of disturbances is shown in Fig. 3.8. LiNGAM [95] automates this procedure, inferring a unique causal model from observational data. LiNGAM is limited to linear relations, but this assumption has been relaxed in

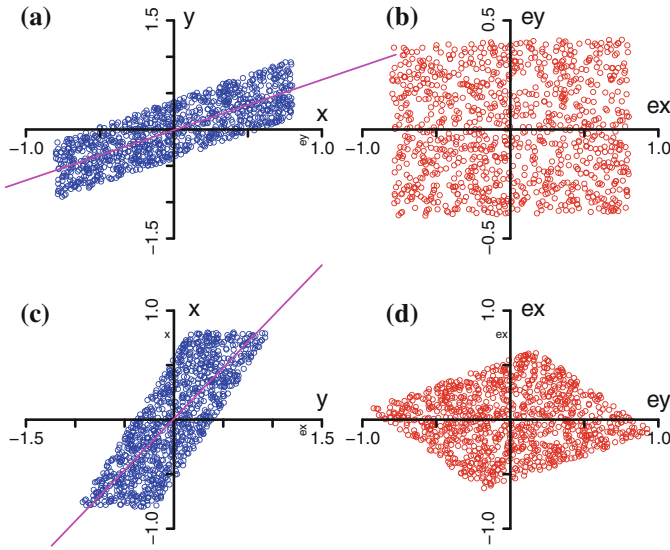


Fig. 3.8 The key idea for LiNGAM and similar algorithms: true structural equation $Y = \beta_{YX} X + \epsilon_Y$, where ϵ_Y follows a uniform distribution. **a** Regression with Y as the dependent variable (true model). **b** Regression with X as the dependent variable (reverse model). **c** Estimated $\hat{\epsilon}_Y$ versus $\hat{\epsilon}_X$ based on the model shown in **(a)**. The disturbances are independent. **d** Estimated $\hat{\epsilon}_Y$ versus $\hat{\epsilon}_X$ based on the model shown in **(b)**. The disturbances are dependent. Figure from [38]

a subsequent body of work [36, 82, 113] to include non-linear relations. However learning such relations requires non-linear optimization techniques and appropriate independence measures [69].

This class of methods is more powerful than traditional causal discovery methods, in the sense that with the functional form assumptions (e.g., linear relations, additive disturbances, non-Gaussian disturbances) causal models are fully identifiable (no statistical indistinguishability). Moreover, the methods also work under unfaithfulness. On the other hand, all methods in this category require large sample sizes and rely on some kind of parametric assumption, and have been shown to be unreliable when this assumption is violated. Nevertheless, these ideas add a new direction and dimension to the way we think about, treat, model, and induce causality and could soon lead to practical results.

3.7 Discussion

Inducing causal models or relations from data is necessary to fully understand biological mechanisms and design new drugs and therapies. Traditional means for such inferences rely on performing interventional experiments. Causal Discovery methods

attempt to make such inferences from observational data alone or with a limited set of such interventions by making assumptions that connect the notion of causality with quantities estimable from the data. The analyst should be aware and conscious of the explicit and implicit assumptions employed by the tools and algorithms that are used and whether they are appropriate for the type of biological data at hand. Despite the inherent theoretical and practical difficulties of the task, there are several successful applications of Causal Discovery methods in systems biology that demonstrate the potential of the field. In addition, recent theoretical and algorithmic breakthroughs promise to further improve the successful application of causal discovery on systems biology.

Acknowledgments This work was partially funded by STATegra EU FP7 project, No 306000 and EPILOGEAS GSRT ARISTEIA II project, No 3446.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

1. Akaike, H.: A new look at the statistical model identification. *IEEE Trans. Automat. Contr.* **19**(6), 716–723 (1974)
2. Akbani, R., Ng, P.K.S., Werner, H.M.J., Shahmoradgoli, M., Zhang, F., Ju, Z., Liu, W., Yang, J.-Y., Yoshihara, K., Li, J., Ling, S., Seviour, E.G., Ram, P.T., Minna, J.D., Diao, L., Tong, P., Heymach, J.V., Hill, S.M., Dondelinger, F., Städler, N., Byers, L., Meric-Bernstam, F., Weinstein, J.N., Broom, B.M., Verhaak, R.G.W., Liang, H., Mukherjee, S., Lu, Y., Mills, G.B.: A pan-cancer proteomic perspective on The Cancer Genome Atlas. *Nat. Commun.* **5**, 3887 (2014)
3. Aliferis, C.F.: Local causal and Markov blanket induction for causal discovery and feature selection for classification Part I?: algorithms and empirical evaluation. *J. Mach. Learn. Res.* **11**, 171–234 (2010)
4. Aten, J.E., Fuller, T.F., Lusi, A.J., Horvath, S.: Using genetic markers to orient the edges in quantitative trait networks: the NEO software. *BMC Syst. Biol.* **2**, 34 (2008)
5. Balding, D.J.: A tutorial on statistical methods for population association studies. *Nat. Rev. Genet.* **7**, 781–791 (2006)
6. Barrett, T., Wilhite, S.E., Ledoux, P., Evangelista, C., Kim, I.F., Tomashevsky, M., Marshall, K.A., Phillippy, K.H., Sherman, P.M., Holko, M., Yefanov, A., Lee, H., Zhang, N., Robertson, C.L., Serova, N., Davis, S., Soboleva, A.: NCBI GEO: Archive for functional genomics data sets—Update. *Nucleic Acids Res.* **41** (2013)
7. Borboudakis, G., Tsamardinos, I.: Bayesian network learning with discrete case-control data. In: *Uncertainty in Artificial Intelligence (UAI)*, 2015
8. Borboudakis, G., Tsamardinos, I.: Incorporating causal prior knowledge as path-constraints in Bayesian networks and maximal ancestral graphs. In: *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, 2012, pp. 1799–1806
9. Burns, M.B., Temiz, N., Harris, R.S.: Evidence for APOBEC3B mutagenesis in multiple human cancers. *Nat. Genet.* **45**(9), 977–983 (2013)

10. Cai, X., Bazerque, J.A., Giannakis, G.B.: Inference of gene regulatory networks with sparse structural equation models exploiting genetic perturbations. *Plos Comput. Biol.* **9**(5), e1003068 (2013)
11. Cause and effect. *Nat. Methods* **7**, 243 (2010)
12. Chen, L.S., Emmert-Streib, F., Storey, J.D.: Harnessing naturally randomized transcription to infer regulatory relationships among genes. *Genome Biol.* **8**(10), R219 (2007)
13. Chickering, D.M.: Learning equivalence classes of Bayesian-network structures. *J. Mach. Learn. Res.* **2**, 445–498 (2002)
14. Chickering, D.M., Heckerman, D., Meek, C.: Large-sample learning of Bayesian networks is NP-hard. *J. Mach. Learn. Res.* **5**, 1287–1330 (2004)
15. Chu, T., Glymour, C., Scheines, R., Spirtes, P.: A statistical problem for inference to regulatory structure from associations of gene expression measurements with microarrays. *Bioinformatics* **19**(9), 1147–1152 (2003)
16. Claassen, T., Heskes, T.: Causal discovery in multiple models from different experiments. In: *Advances in Neural Information Processing Systems (NIPS 2010)*, 2010, pp. 1–9
17. Claassen, T., Heskes, T.: Learning causal network structure from multiple (in) dependence models. In: *Proceedings of the Fifth European Workshop on Probabilistic Graphical Models (PGM)*, pp. 81–88
18. Cooper, G.F.: A simple constraint-based algorithm for efficiently mining observational databases for causal relationships. *Data Min. Knowl. Discov.* **1**, 203–224 (1997)
19. Cooper, G., Yoo, C.: Causal discovery from a mixture of experimental and observational data. In: *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence (UAI'99)*, 1999, pp. 116–125
20. Danks, D.: Learning the causal structure of overlapping variable sets. In: *Discovery Science: Proceedings of the 5th International Conference, 2002*, pp. 178–191
21. Danks, D., Glymour, C., Tillman, R.E.: Integrating locally learned causal structures with overlapping variables. In: *Advances in Neural Information Processing Systems*, pp. 1665–1672. MIT Press, Cambridge (2009)
22. Dash, D., Druzdzel, M.: Caveats for causal reasoning with equilibrium models. *ECSQUARU* **2143**, 192–203 (2001)
23. Davey Smith, G., Ebrahim, S.: Mendelian randomization: can genetic epidemiology contribute to understanding environmental determinants of disease?. *Int. J. Epidemiol.* **32**(1), 1–22 (2003)
24. D’Orazio, M., Di Zio, M., Scanu, M.: *Statistical Matching: Theory and Practice*, p. 268. Wiley, New York (2006)
25. Eaton, D., Murphy, K.: Exact Bayesian structure learning from uncertain interventions. In: *AISTATS* (2007)
26. Eberhardt, F.: Sufficient condition for pooling data from different distributions. *Error* (2006)
27. Evans, R.J., Richardson, T.S.: Marginal log-linear parameters for graphical Markov models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **75**(4), 743–768 (2013)
28. Fisher, R.A.: The distribution of the partial correlation coefficient. *Metron* **3**(3–4), 329–332 (1923)
29. Fisher, R.A.: *The Design of Experiments*. Hafner Publishing, New York (1935)
30. Friedman, N., Murphy, K., Russell, S.: Learning the structure of dynamic probabilistic networks. In: *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, 1998, pp. 139–147
31. Fukuda, R., Kelly, B., Semenza, G.L.: Vascular endothelial growth factor gene expression in colon cancer cells exposed to prostaglandin E2 is mediated by hypoxia-inducible factor 1. *Cancer Res.* **63**(9), 2330–2334 (2003)
32. Gutierrez-Arcelus, M., Lappalainen, T., Montgomery, S.B., Buil, A., Ongen, H., Yurovsky, A., Bryois, J., Giger, T., Romano, L., Planchon, A., Falconnet, E., Bielser, D., Gagnebin, M., Padiou, I., Borel, C., Letourneau, A., Makrythanasis, P., Guipponi, M., Gehrig, C., Antonarakis, S.E., Dermitzakis, E.T.: Passive and active DNA methylation and the interplay with genetic variation in gene regulation. *Elife* **2**, e00523 (2013)

33. Hartemink, A.J., Gifford, D.K., Jaakkola, T.S., Young, R.A.: Using graphical models and genomic expression data to statistically validate models of genetic regulatory networks. In: *The Pacific Symposium on Biocomputing*, 2001, pp. 422–433
34. Hauser, A., Bühlmann, P.: Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *J. Mach. Learn. Res.* **13**, 2409–2464 (2012)
35. He, Y.-B.: Active learning of causal networks with intervention experiments and optimal designs. *J. Mach. Learn. Res.* **9**, 2523–2547 (2008)
36. Hoyer, P.O., Janzing, D., Joris, M., Peters, J., Schölkopf, B.: Nonlinear causal discovery with additive noise models. In: *NIPS*, 2008
37. Hug, S., Schmidl, D., Li, W.B., Greiter, M.B., Theis, F.J.: Bayesian model selection methods and their application to biological ODE systems. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
38. Hyttinen, A.: *Discovering Causal Relations in the Presence of Latent Confounders*. University of Helsinki, Helsinki (2013)
39. Hyttinen, A., Eberhardt, F., Hoyer, P.O.: Noisy-OR models with latent confounding. In: *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence*, Barcelona, 2011
40. Hyttinen, A., Eberhardt, F., Hoyer, P.O.: Causal discovery of linear cyclic models from multiple experimental data sets with overlapping variables. In: *Proceedings of the Uncertainty in Artificial Intelligence*, 2012
41. Hyttinen, A., Eberhardt, F., Hoyer, P.O.: Learning linear cyclic causal models with latent variables. *J. Mach. Learn. Res.* **2013**(3387–3439), 3387–3439 (2012). Jan
42. Hyttinen, A., Eberhardt, F., Jarvisalo, M.: Constraint-Based Causal Discovery: Conflict Resolution with Answer Set Programming. In: *Proceedings of the Uncertainty in Artificial Intelligence*, 2014
43. Hyttinen, A., Hoyer, P.O., Eberhardt, F., Jarvisalo, M.: Discovering cyclic causal models with latent variables: a general sat-based procedure. In: *Proceedings of the Uncertainty in Artificial Intelligence*, 2013
44. Idaghdour, Y., Czika, W., Shianna, K.V., Lee, S.H., Visscher, P.M., Martin, H.C., Miclaus, K., Jadallah, S.J., Goldstein, D.B., Wolfinger, R.D., Gibson, G.: Geographical genomics of human leukocyte gene expression variation in southern Morocco. *Nat. Genet.* **42**(1), 62–67 (2010)
45. Kalisch, M., Fellinghauer, B.A.G., Grill, E., Maathuis, M.H., Mansmann, U., Bühlmann, P., Stucki, G.: Understanding human functioning using graphical models. *BMC Med. Res. Methodol.* **10**, 14 (2010)
46. Kalisch, M., Maechler, M., Colombo, D., Maathuis, M.H., Bühlmann, P.: Causal inference using graphical models with the R package pcalg. *J. Stat. Softw.* **47**(11), 1–26 (2012)
47. Katan, M.B.: Apolipoprotein E isoforms, serum cholesterol, and cancer. *Lancet* **1**(8479), 507–508 (1986)
48. Kenfield, S.A., Stampfer, M.J., Chan, J.M., Giovannucci, E.: Smoking and prostate cancer survival and recurrence. *JAMA* **305**(24), 2548–2555 (2011)
49. Kirk, P., Silk, D., Stumpf, M.P.H.: Reverse engineering under uncertainty. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
50. Labrie, F., Dupont, A., Suburu, R., Cusan, L., Tremblay, M., Gomez, J.L., Emond, J.: Serum prostate specific antigen as pre-screening test for prostate cancer. *J. Urol.* **147**(3 Pt 2), 846–851 (discussion 851–852) (1992)
51. Lagani, V., Tsamardinos, I., Triantafyllou, S.: Learning from mixture of experimental data: a constraint-based approach. In: *SETN'12 Proceedings of the 7th Hellenic Conference on Artificial Intelligence: Theories and Applications*, 2012, vol. 7297, pp. 124–131
52. Lazar, C., Meganck, S., Taminau, J., Steenhoff, D., Coletta, A., Molter, C., Weiss-Solís, D.Y., Duque, R., Bersini, H., Nowé, A.: Batch effect removal methods for microarray gene expression data integration: a survey. *Brief. Bioinform.* **14**(4), 469–490 (2013)
53. Le, T.D., Liu, L., Tsykin, A., Goodall, G.J., Liu, B., Sun, B.-Y., Li, J.: Inferring microRNA-mRNA causal regulatory relationships from expression data. *Bioinformatics* **29**(6), 765–771 (2013)

54. Lemeire, J., Janzing, D.: Replacing causal faithfulness with algorithmic independence of conditionals. *Minds Mach.* (2012)
55. Liu, Y., Aryee, M.J., Padyukov, L., Fallin, M.D., Hesselberg, E., Runarsson, A., Reinius, L., Acevedo, N., Taub, M., Ronninger, M., Shchetynsky, K., Scheynius, A., Kere, J., Alfredsson, L., Klareskog, L., Ekström, T.J., Feinberg, A.P.: Epigenome-wide association data implicate DNA methylation as an intermediary of genetic risk in rheumatoid arthritis. *Nat. Biotechnol.* **31**(2), 142–7 (2013)
56. Maathuis, M.H., Kalisch, M., Bühlmann, P.: Estimating high-dimensional intervention effects from observational data. *Ann. Stat.* **37**(6A), 3133–3164 (2009)
57. Maathuis, M.H., Colombo, D., Kalisch, M., Bühlmann, P.: Predicting causal effects in large-scale systems from observational data. *Nat. Methods* **7**(4), 247–248 (2010)
58. MacKinnon, D.P.: *Introduction to Statistical Mediation Analysis (Multivariate Applications Series)*, p. 488. Routledge, New York (2008)
59. Mani, S., Cooper, G.F.: Causal discovery using a Bayesian local causal discovery algorithm. *Stud. Health Technol. Inform.* **107**(Pt 1), 731–735 (2004)
60. Marbach, D., Schaffter, T., Mattiussi, C., Dario, F.: Generating realistic in silico gene networks for performance assessment of reverse engineering methods. *J. Comput. Biol.* **16**(2), 229–239 (2009)
61. Margaritis, D.: Distribution-free learning of Bayesian network structure in continuous domains. In: *AAAI'05 Proceedings of the 20th National Conference on Artificial Intelligence—Volume 2, 2005*, pp. 825–830
62. Margaritis, D., Thrun, S.: Bayesian network induction via local neighborhoods. *Adv. Neural Inf. Process. Syst.* **12**, 505–511 (2000)
63. McDonald, J.H.: *Handbook of Biological Statistics*, p. 291. Sparky House Publishing, Baltimore (2009)
64. Meganck, S., Maes, S., Leray, P., Manderick, B.: Learning semi-markovian models using experiments. In: *Third European Workshop on Probabilistic Graphical Models (PGM), 2006*
65. Meinshausen, N., Bühlmann, P.: Stability selection. *J. R. Stat. Soc. Ser. B* **72**(4), 417–473 (2010)
66. Miles, C., Wayne, M.: Quantitative trait locus (QTL) analysis. *Nat. Educ.* **1**(1) (2008)
67. Millstein, J., Zhang, B., Zhu, J., Schadt, E.E.: Disentangling molecular relationships with a causal inference test. *BMC Genet.* **10**, 23 (2009)
68. Monti, S., Cooper, G.F.: A multivariate discretization method for learning Bayesian networks from mixed data. In: *UAI'98 Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, 1998*, pp. 404–413
69. Mooij, J., Janzing, D., Peters, J., Schölkopf, B.: Regression by dependence minimization and its application to causal inference in additive noise models. In: *Proceedings of the 26th Annual International Conference on Machine Learning—ICML '09, 2009*, pp. 745–752
70. Näger, P.M.: Causal graphs for EPR experiments. In: *Foundations of Physics 2013: The 17th UK and European Meeting on the Foundations of Physics, 2013*
71. Neapolitan, R.E.: *Learning Bayesian Networks*. Pearson Prentice Hall, New York (2004)
72. Neto, E.C., Keller, M.P., Attie, A.D., Yandell, B.S.: Causal graphical models in systems genetics: a unified framework for joint inference of causal network and genetic architecture for correlated phenotypes. *Ann. Appl. Stat.* **4**(1), 320–339 (2010). Mar
73. Nica, A.C., Montgomery, S.B., Dimas, A.S., Stranger, B.E., Beazley, C., Barroso, I., Dermizakis, E.T.: Candidate causal regulatory effects by integration of expression QTLs with complex trait genetic associations. *PLoS Genet.* **6**(4), e1000895 (2010)
74. Nitsch, D., Molokhia, M., Smeeth, L., DeStavola, B.L., Whittaker, J.C., Leon, D.A.: Limits to causal inference based on Mendelian randomization: a comparison with randomized controlled trials. *Am. J. Epidemiol.* **163**(5), 397–403 (2006)
75. Ornatsky, O., Bandura, D., Baranov, V., Nitz, M., Winnik, M.A., Tanner, S.: Highly multiparametric analysis by mass cytometry. *J. Immunol. Methods* **361**(1–2), 1–20 (2010)
76. O'Rourke, K.: An historical perspective on meta-analysis: dealing quantitatively with varying study results. *J. R. Soc. Med.* **100**(12), 579–82 (2007)

77. Pan, S.J., Yang, Q.: A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* **22**, 1345–1359 (2010)
78. Pearl, J.: *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Francisco (1988)
79. Pearl, J.: *Causality: Models, Reasoning and Inference*. Cambridge University Press, Cambridge (2009)
80. Pearl, J.: *Interpretation and Identification of Causal Mediation*. University of California, Los Angeles (2013)
81. Peña, J.M., Björkegren, J., Tegnér, J.: Growing Bayesian network models of gene networks from seed genes. *Bioinformatics* **21**(Suppl 2), ii224–i229 (2005)
82. Peters, J., Mooij, J., Janzing, D., Schoelkopf, B.: Identifiability of Causal Graphs Using Functional Models. [arXiv.org](https://arxiv.org/abs/1212.4002) (2012)
83. Petretto, E.: Single cell expression quantitative trait loci and complex traits. *Genome Med.* **5**(8), 72 (2013)
84. Richardson, T., Spirtes, P.: Automated causal discovery in linear feedback models. In: Glymour, C., Cooper, G. (eds.) *Computation, Causation and Discovery*, pp. 253–302. AAAI press, Cambridge (1999)
85. Richardson, T., Spirtes, P.: Ancestral graph Markov models. *Ann. Stat.* **30**(4), 962–1030 (2002)
86. Richardson, T., Evans, R., Robins, J.: Transparent parametrizations of models for potential outcomes. *Bayesian Stat.* **9**, 569–610 (2011)
87. Robins, J.M., Wasserman, L.: On the impossibility of inferring causation from association without background knowledge. In: Glymour, C., Cooper, G.F. (eds.) *Computation, Causation, and Discovery*, pp. 305–321. AAAI Press/The MIT Press, Menlo Park, CA, Cambridge, MA (1999)
88. Rockman, M.V.: Reverse engineering the genotype-phenotype map with natural genetic variation. *Nature* **456**, 738–744 (2008)
89. Sachs, K., Perez, O., Pe’er, D., Lauffenburger, D.A., Nolan, G.P.: Causal protein-signaling networks derived from multiparameter single-cell data. *Science* **308**(5721), 523–529 (2005)
90. Schadt, E.E.: Causal inference and the construction of predictive network models in biology. In: *Handbook of Systems Biology Concept and Insights*, pp. 499–514. Elsevier Inc. (2013)
91. Schadt, E.E., Lamb, J., Yang, X., Zhu, J., Edwards, S., Guhathakurta, D., Sieberts, S.K., Monks, S., Reitman, M., Zhang, C., Lum, P.Y., Leonardson, A., Thieringer, R., Metzger, J.M., Yang, L., Castle, J., Zhu, H., Kash, S.F., Drake, T.A., Sachs, A., Lusk, A.J.: An integrative genomics approach to infer causal associations between gene expression and disease. *Nat. Genet.* **37**(7), 710–717 (2005)
92. Schmidt, M., Murphy, K.: Modeling discrete interventional data using directed cyclic graphical models. In: *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (UAI '09)*, 2009, pp. 487–495
93. Schwarz, G.: Estimating the dimension of a model. *Ann. Stat.* **6**(2), 461–464 (1978)
94. Shang, Z., Zhu, S., Zhang, H., Li, L., Niu, Y.: Germline homeobox B13 (HOXB13) G84E mutation and prostate cancer risk in European descendants: a meta-analysis of 24,213 cases and 73,631 controls. *Eur. Urol.* **64**(1), 173–176 (2013)
95. Shimizu, S., Hoyer, P.O., Hyvärinen, A., Kerminen, A.: A linear non-Gaussian acyclic model for causal discovery. *J. Mach. Learn. Res.* **7**(2), 2003–2030 (2006)
96. Sleiman Itani, B.S., Ohannessian, M., Sachs, K., Nolan, G.P., Dahleh, M.A., Guyon, I., Janzing, D.: Structure learning in causal cyclic networks. In: *NIPS*, 2008, pp. 165–176
97. Sobel, M.E.: Asymptotic confidence intervals for indirect effects in structural equation models. *Sociol. Methodol.* **13**(1982), 290–312 (1982)
98. Spirtes, P., Glymour, C., Scheines, R.: *Causation, Prediction, and Search*, vol. 81. Springer, New York (1993)
99. Stekhoven, D.J., Moraes, I., Sveinbjornsson, G., Hennig, L., Maathuis, M.H., Buhlmann, P.: Causal stability ranking. *Bioinformatics* **28**(21), 2819–2823 (2012)

100. Stelling, J., Klamt, S., Bettenbrock, K., Schuster, S., Gilles, E.D.: Metabolic network structure determines key aspects of functionality and regulation. *Nature* **420**(6912), 190–193 (2002)
101. Sunnåker, M., Stelling, J.: Model extension and model selection. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
102. Tian, J., Pearl, J.: Causal discovery from changes. In: *UAI'01 Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence, 2001*, pp. 512–521
103. Tian, J., Pearl, J.: On the identification of causal effects. Technical Report R-290-L, 2003
104. Tillman, R.E., Spirtes, P.: Learning equivalence classes of acyclic models with latent and selection variables from multiple datasets with overlapping variables. *J. Mach. Learn. Res. Proc. Track* **15**, 3–15 (2011)
105. Triantafillou, S., Tsamardinos, I.: Constraint-Based Causal Discovery from Multiple Interventions over Overlapping Variable Sets. *JMLR*, to appear
106. Triantafillou, S., Tsamardinos, I., Tollis, I.G.: Learning causal structure from overlapping variable sets. In: *Proceedings of Artificial Intelligence and Statistics, 2010*
107. Tsamardinos, I., Borboudakis, G.: Permutation testing improves Bayesian network learning. In: *ECML PKDD'10 Proceedings of the 2010 European Conference on Machine Learning and Knowledge Discovery in Databases: Part III, 2010*, pp. 322–337
108. Tsamardinos, I., Brown, L.E., Aliferis, C.F.: The max-min hill-climbing Bayesian network structure learning algorithm. *Mach. Learn.* **65**(1), 31–78 (2006)
109. Tsamardinos, I., Triantafillou, S., Lagani, V.: Towards integrative causal analysis of heterogeneous data sets and studies. *J. Mach. Learn. Res.* **13**(1), 1097–1157 (2012)
110. Tsamardinos, I., Aliferis, C.F., Statnikov, A., Brown, L.E.: Scaling-Up Bayesian Network Learning to Thousands of Variables Using Local Learning Techniques
111. Wills, Q.F., Livak, K.J., Tipping, A.J., Enver, T., Goldson, A.J., Sexton, D.W., Holmes, C.: Single-cell gene expression analysis reveals genetic associations masked in whole-tissue experiments. *Nat. Biotechnol.* **31**(8), 748–752 (2013)
112. Zhang, J.: On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artif. Intell.* **172**(16–17), 1873–1896 (2008)
113. Zhang, K., Hyvärinen, A.: On the identifiability of the post-nonlinear causal model. In: *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, 2009*, pp. 647–655
114. Zhang, W.M., Wong, T.M.: Suppression of cAMP by phosphoinositol/Ca²⁺ pathway in the cardiac kappa-opioid receptor. *Am. J. Physiol.* **274**(1 Pt 1), C82–C87 (1998)
115. Zhang, W., Zhu, J., Schadt, E.E., Liu, J.S.: A Bayesian partition method for detecting pleiotropic and epistatic eQTL modules. *PLoS Comput. Biol.* **6**(1), e1000642 (2010)
116. Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B* **67**(2), 301–320 (2005)

Chapter 4

Stochastic Modeling and Simulation Methods for Biological Processes: Overview

Annelies Lejon and Giovanni Samaey

Abstract The use of stochastic modeling and simulation techniques is widespread in computational biology when fluctuations become important. In this chapter, we give a high-level overview of stochastic modeling techniques for biological problems, focussing on some common individual-based modeling and simulation methods. We pay particular attention to the equivalence between the stochastic process that governs the evolution of individual agents and the deterministic behaviour of the involved probability distributions, and we discuss numerical methods that exploit this relation for variance reduction purposes. The discussion will be illustrated using examples involving intracellular chemical reactions, bacterial chemotaxis and tumor growth, showing the effects of stochasticity at different scales and different levels of description.

Keywords Stochasticity · Stochastic differential equations · Velocity-jump processes · Asymptotic variance reduction

4.1 Introduction

Stochastic effects are ubiquitous in biological systems, at multiple scales. At an intracellular level, for instance, gene regulatory networks often exhibit different metastable states. Since the number of molecules in a cell is not very high, significant fluctuations in concentrations can occur, triggering transitions between these metastable states [20, 28]. At a cellular level, individual cells can be modeled as agents that interact with each other and with their environment.

Stochasticity can then be introduced to account for differences between individual cells or to incorporate the coarse-grained effect of phenomena that occur at more microscopic scales directly at the cellular level. (An example of the latter would be the use of a Brownian motion to model the net effect of a large number of collisions of

A. Lejon · G. Samaey (✉)

Department of Computer Science, KU Leuven, Celestijnenlaan 200A, 3001 Leuven, Belgium
e-mail: Giovanni.Samaey@cs.kuleuven.be

© Springer International Publishing Switzerland 2016
L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_4

a large molecule with surrounding—but not explicitly modeled—solvent molecules.) Such an approach has been followed in many settings, including the applications that will be considered as illustrative examples in this chapter: bacterial chemotaxis (see, e.g., [5, 15] and references therein) and tumor growth (see, e.g., [43] for a recent review and references). At a population level, one usually models the evolution of cell densities via partial differential equations (PDEs) of reaction-advection-diffusion type. At this level, one can introduce, for instance, stochastic parameters and geometries to account for differences between individuals, resulting in PDEs with stochastic coefficients [22].

Despite the stochastic nature of the time evolution, one is usually interested in deterministic quantities, such as the mean switching time between metastable states or the expected behaviour of a large population of cells or individuals. Additionally, one may also be interested in deviations with respect to this mean behaviour, requiring information on higher order statistics or on the complete probability distribution of possible states of the system. While, in principle, the time evolution of these probability distributions can be modeled using *deterministic* evolution laws, the associated computational cost is usually prohibitive due to the high dimensionality of the resulting equations. One therefore needs to resort to some form of Monte Carlo simulation of the stochastic process [7].

In this chapter, we discuss stochastic individual-based modeling techniques and show the equivalence with deterministic techniques for modeling the involved probability distributions. In Sect. 4.2, we introduce stochastic models for chemically reacting systems with low number of molecules (as would occur in modeling intracellular dynamics). We discuss both the time-discrete and time-continuous case, and show how these models relate to classical mean-field equations for the evolution of concentrations. In Sect. 4.3, we consider advection-diffusion processes, as they occur, for instance, in agent-based models for bacterial chemotaxis and tumor growth. We briefly describe cellular automata and Markov jump processes, before giving a more detailed discussion of stochastic differential equations (SDEs). We show the relation between an SDE for an individual particle and an advection-diffusion equation (the Fokker–Planck equation) for the population density. In Sect. 4.4, we turn to more realistic microscopic processes, and relate these to kinetic theory and Boltzmann-type equations. For each of the described modeling techniques, we discuss the mathematical formulation of the stochastic process as well as practical simulation algorithms. In Sect. 4.5, we discuss Monte Carlo simulation using these stochastic processes. We show how to compute confidence intervals for the obtained results and introduce numerical algorithms that can yield results with significantly reduced variance.


We illustrate the introduced concepts for two specific applications: bacterial chemotaxis (Sect. 4.6) and tumor growth (Sect. 4.7). For each of these applications, we discuss in detail the modeled processes and introduce a dedicated simulation technique.

We conclude this introduction with a few remarks on topics that will not be treated in this chapter. First, we will not discuss uncertainty propagation resulting from uncertainty in parameters or geometries. For a mathematical and algorithmic introduction to this topic, we refer to [22]. Second, in many applications, one may

have, besides the mathematical model, some observation data available. Two questions can then be posed: (i) does the data provide sufficient support to validate the model; and (ii) can one use the data to estimate unknown parameters in the model? These topics form the subject of intense current research, and we refer to [51] for an introduction.

4.2 Stochastic Modeling of Chemical Reactions

In this section, we provide an individual-based description of chemically reacting systems. For concreteness, we start from the following system of chemical reactions, as introduced by Schlögl [46],



in which the rate constants k_i ($1 \leq i \leq 4$) indicate that the probability for two randomly chosen molecules of type A to react according to reaction i in any time interval $[t, t + dt)$ is given by $k_i dt$.

In Sect. 4.2.1, we discuss discrete-time simulation. It will appear that the corresponding algorithm introduces a time discretization error, and at the same time is also very inefficient. We therefore turn to continuous-time simulation algorithms in Sect. 4.2.2. Next, in Sect. 4.2.3, we introduce the chemical master equation, a deterministic system of equations for the probability of finding a given number of molecules at any given time. From this master equation, we can obtain information on the mean behaviour of the system and on fluctuations. We conclude in Sect. 4.2.4 with a numerical illustration on the Schlögl model. The exposition in this section is based on [17].

4.2.1 Discrete-Time Simulation

Let us denote by $A(t)$ the number of molecules of type A in the system at time $t \geq 0$, and assume as initial condition $A(0) = A_0$. We are interested in an approximate solution $A^k \approx A(t^k)$, with $t^k = k\Delta t$ and $\Delta t > 0$ a small time step. (The symbol k that indicates the discrete time instance is added as a superscript for notational consistency throughout the chapter.)

At time $t = t^k$, the probability of reaction 1 (with rate k_1) taking place in $(t, t + \Delta t)$ is approximately given by

$$\alpha_1^k \Delta t = A^k (A^k - 1) k_1 \Delta t, \quad (4.2)$$

with α_1^k the *propensity function*, since $k_1 \Delta t$ is the probability of two randomly chosen molecules of type A to react according to reaction 1, and $A^k(A^k - 1)$ is the number of pairs of molecules that can be randomly chosen. (Note that this is an approximation due to the fact that we have replaced the infinitesimal interval of length dt by a finite interval of length Δt .) Similarly, the propensity functions for the other reactions can be seen to be

$$\alpha_2^k = A^k(A^k - 1)(A^k - 2)k_2, \quad \alpha_3^k = k_3, \quad \alpha_4^k = A^k k_4. \quad (4.3)$$

To take a time step from time t^k to t^{k+1} , one needs to decide for each reaction whether or not it has occurred during the time step. This can be done using the following algorithm:

Algorithm 1 (*Time-discrete simulation of Schlögl's model*) Given the concentration A^k at time t^k , we compute A^{k+1} at time t^{k+1} as follows:

1. For each reaction i :
 - Compute the propensity α_i^k using (4.2) or (4.3);
 - Generate an independent random number r_i^k from the uniform distribution on the interval $(0, 1)$;
 - Decide that reaction i has occurred during the time interval if $r_i^k \leq \alpha_i^k \Delta t$ (else, the reaction has not occurred);
2. Compute A^{k+1} by applying all reactions that occurred during the time step (for instance, if only reaction 1 occurred, we have $A^{k+1} = A^k + 1$).

A simulation is then performed by repeating the above time step over the time interval of interest. The algorithm can easily be generalized to systems with multiple species and any number I of possible reactions.

There are two problems when using the above algorithm, which will both be dealt with when switching to continuous-time simulation in Sect. 4.2.2. First, Algorithm 1 introduces a time discretization error since it replaces an infinitesimal time interval of length dt by a finite interval of length Δt . This error manifests itself in two ways. First, we neglect the (non-zero) probability that two reactions of the same type occur within the time step of size Δt . Moreover, every reaction event influences the propensity function of other reactions, since the propensity functions depend on the state of the system. As a consequence, there will be an error in the used reaction probabilities when multiple reactions are performed during a single time step.

To limit the impact of these time discretization errors, Δt should be chosen fairly small. A common guideline is to choose Δt such that the probability of having a reaction is less than 1 % per time step. However, this implies that 99 % of the time steps are taken just to conclude that nothing happened! This is detrimental to computational efficiency, which is the second problem with the above algorithm.

4.2.2 Continuous-Time Simulation

Instead of taking discrete time steps of fixed size Δt from t to $t + \Delta t$, and calculating the probability that a reaction takes place in that time interval $[t, t + \Delta t)$, a continuous-time simulation computes the (random) time increment τ until the next reaction takes place by sampling from the corresponding probability distribution. Afterwards, a second random number is generated to decide *which* reaction occurred.

Let us first characterize the relevant probability distribution for τ . Consider a system with one reaction with propensity $\alpha_1(t)$. (Recall that the propensity is time-dependent due to the time-dependence of the concentrations of the present species.) Further, denote as $f_1(t, s)ds$ the probability, at time t , that the next reaction occurs in the time interval $[t + s, t + s + ds]$, and denote the probability that no reaction occurs in the interval $[t, t + s]$ as $g_1(t, s)$. We then have

$$f_1(t, s)ds = g_1(t, s)\alpha_1(t + s)ds = g_1(t, s)\alpha_1(t)ds, \quad (4.4)$$

in which we used the fact that $\alpha_1(t + s) = \alpha_1(t)$ in the absence of reactions, and independence of individual reaction events. (Then, the probability that the *first* reaction occurs in the time interval $[t + s, t + s + ds]$ is given by the product of the probability $\alpha_1(t + s)ds$ that there is a reaction in that time interval and the probability $g_1(t, s)$ that no reaction occurred earlier.)

The two quantities $f_1(t, s)$ and $g_1(t, s)$ are, as for now, unknown. We derive a differential equation for $g_1(t, s)$. We start by observing that the probability of not having a reaction in the time interval $[t, t + s + ds]$ can be written as the product of the probability of not having a reaction in the time interval $[t, t + s]$ and the probability of not having a reaction in the time interval $[t + s, t + s + ds]$, i.e.,

$$g_1(t, s + ds) = g_1(t, s)(1 - \alpha_1(t)ds), \quad (4.5)$$

from which we obtain

$$\frac{g_1(t, s + ds) - g_1(t, s)}{ds} = -\alpha_1(t)g_1(t, s). \quad (4.6)$$

Given that $g(t, 0) = 1$ (zero probability of having the reaction exactly at time t), we obtain

$$g_1(t, s) = \exp(-\alpha_1(t)s), \quad (4.7)$$

and hence, using (4.4), the probability that the first reaction occurs in the infinitesimal interval $t + s$ is given by

$$f_1(t, s)ds = \alpha_1(t) \exp(-\alpha_1(t)s)ds, \quad s \geq 0. \quad (4.8)$$

The corresponding cumulative distribution is given as

$$F_1(t, \tau) = \int_0^\tau f_1(t, s) ds = 1 - \exp(-\alpha_1(t)\tau), \quad \tau \geq 0. \quad (4.9)$$

(Note that we have $F_1(t, 0) = 0$ and $\lim_{\tau \rightarrow \infty} F_1(t, \tau) = 1$.)

Now that we know the probability distribution $f_1(t, s)$, we are ready to generate a random time increment $\tau_1(t)$, sampled from $f_1(t, s)$, which we denote as $\tau_1(t) \sim f_1(t, s)$. Using the transformation method for the generation of random numbers [7], the increment $\tau_1(t)$ until the next event of reaction 1 can be computed from a uniformly distributed random number $\tilde{\theta}_1$ in $(0, 1)$ via

$$\tau_1 = F^{-1}(t, \tilde{\theta}_1) = -\frac{1}{\alpha_1(t)} \log(1 - \tilde{\theta}_1), \quad \text{or} \quad \tau_1 = -\frac{1}{\alpha_1(t)} \log(\theta_1), \quad (4.10)$$

with θ_1 also a uniformly distributed random number in $(0, 1)$ and $F^{-1}(t, \theta)$ the inverse of the cumulative density $F(t, s)$ with respect to s , treating t as a parameter.

Remark 1 (Markov property) The probability distribution $f_1(t, s)$ is called the exponential distribution with rate $\alpha_1(t)$. Its main property is that it is memoryless, which implies that the evolution is completely determined by the current state (and no information from the past is required). In particular, this implies that, to determine when the next reaction will occur, it is irrelevant how long the system is already in its current state. Mathematically, this can be seen by checking that, for any positive T and τ , we have

$$\Pr(\tau_1(t) > T) = \Pr(\tau_1(t) > \tau + T | \tau_1(t) > \tau), \quad (4.11)$$

i.e., the probability that the next reaction will not occur within a time interval of length T from the current time, does not depend on the amount of time τ that the system is already in the current state.

If multiple reactions can occur in the system (as is the case for the Schlögl model (4.1)), a naive way of proceeding is to generate the next reaction time for each of the reactions and only select the reaction that occurs first, after which the system clock is updated and the procedure is repeated. With I possible reactions, this requires the generation of I exponentially distributed random numbers (independent of the number of species in the system) to choose a single reaction. This algorithm can be made much more efficient by making use of the following theorem:

Theorem 2 (Exponential distributions) *Consider I exponential distributions $f_i(t, s)$ with rates $\alpha_i(t)$, and consider an independent set of random numbers $\tau_i(t)$, each sampled from the corresponding distribution $f_i(t, s)$. Let $\tau(t)$ be the minimum of these, $\tau(t) = \min_i \tau_i(t)$. Then, the probability distribution for $\tau(t)$ is an exponential distribution with rate $\alpha(t) = \sum_i \alpha_i(t)$. Moreover, the probability that $\tau(t) = \tau_i(t)$ (the probability that the i -th reaction occurs first) is given by $\alpha_i(t)/\alpha(t)$.*

Using this theorem, only two random numbers need to be generated per time step: one to determine the time increment until the next reaction, and one to choose the next reaction. This gives rise to the following classical algorithm, due to Gillespie [23], which is immediately written for a system consisting of I reactions:

Algorithm 2 (*Stochastic simulation algorithm (SSA) for chemically reacting systems*) Given the concentration A^k at time t^k , we compute A^{k+1} at time t^{k+1} as follows:

1. For each reaction i , compute the propensity $\alpha_i^k = \alpha_i(t^k)$, and compute the total propensity $\alpha^k = \sum_i \alpha_i^k$;
2. Generate a uniformly distributed random number θ^k in $(0, 1)$, and compute the time increment until the next reaction occurs as $\tau^k = -\frac{1}{\alpha^k} \log(\theta^k)$,
3. Generate a uniformly distributed random number y^k in $(0, 1)$ and select the reaction i for which

$$\sum_{j \leq i-1} \alpha_j^k / \alpha^k \leq y < \sum_{j \leq i} \alpha_j^k / \alpha^k;$$

4. Compute A^{k+1} by applying the selected reaction i (for instance, if reaction 1 was selected, we have $A^{k+1} = A^k + 1$).

Again, the algorithm can easily be generalized to systems with multiple species. We refer to [26] for more details and variants. Here, we only state the main convergence result: the stochastic simulation algorithm (SSA) is *exact*, in the sense that it does not contain any time discretization error.

Remark 3 (Acceleration of SSA) While this method is significantly faster than the time-discrete Algorithm 1 without introducing a time discretization error, the resulting algorithm can still be computationally prohibitive, especially in situations with many chemical reactions with disparate time scales. Consider for instance a system with a fast, reversible reaction and an irreversible but very slow reaction. In such a system, most of the time steps will select the fast reversible reaction, resulting in very small time steps that go back and forth along the fast reaction. More sophisticated algorithms, tailored to these situations, have been developed. These include, for instance, the τ -leaping method [9, 25], the slow-scale stochastic simulation algorithm [8], and R-leaping [3]. Note, however, that such methods accelerate simulation at the expense of re-introducing a (small) time-discretization error.

4.2.3 Population-Level Dynamics and Mean-Field Approximation

In general, repeated simulation of a stochastic process for a chemically reacting system yields different results for each stochastic realization. The precise results depend on the generated sequence of random numbers. Usually, one is not interested in the

detailed behaviour of such an individual realization, but in quantitative statements on the mean behaviour and on fluctuations.

In this section, we first derive the *chemical master equation*, which gives a complete description of the (time-dependent) probability distribution of possible states for the system. Afterwards, we discuss the potential and limitations of using this equation to derive information on the statistics of the process.

In the previous sections, we denoted by the random variable $A(t)$ the number of molecules of type A at time t . Here, we define the *deterministic* quantity $p_n(t)$, which represents the *probability* that the system contains exactly n molecules of type A at time t , i.e.,

$$p_n(t) = \Pr(A(t) = n).$$

In the incremental time interval $[t, t + dt)$, the state can only change by ± 1 , since all reactions either create or destroy one molecule of type A . We can use the definition of the reactions (4.1) to compute $p_n(t + dt)$,

$$\begin{aligned} p_n(t + dt) &= p_n(t) + \underbrace{\{[k_1(n-1)(n-2) + k_3] p_{n-1}(t)dt + [k_2(n+1)n(n-1) + k_4(n+1)] p_{n+1}(t)dt\}}_{\text{gain}} \\ &\quad - \underbrace{\{[k_1n(n-1) + k_3] + [k_2n(n-1)(n-2) + k_4n]\} p_n(t)dt}_{\text{loss}}, \end{aligned} \quad (4.12)$$

in which we recognize

- a gain term that expresses the sum of (i) the probability that the system contains $n - 1$ molecules at time t and a molecule is produced during the time increment; and (ii) the probability that the system contains $n + 1$ molecules at time t and a molecule is destroyed during the time increment;
- a loss term that expresses the probability that the system contained n molecules at time t and a molecule is either created or destroyed during the time increment.

Reordering the terms, we get an (infinite-dimensional) system of ordinary differential equations, the chemical master equation:

$$\begin{aligned} \dot{p}_n(t) &= [k_1(n-1)(n-2) + k_3] p_{n-1}(t) + [k_2(n+1)n(n-1) + k_4(n+1)] p_{n+1}(t) \\ &\quad - [k_1n(n-1) + k_3 + k_2n(n-1)(n-2) + k_4n] p_n(t), \quad n \geq 0, \end{aligned} \quad (4.13)$$

in which, formally, $p_{-1}(t) = 0$.

The chemical master equation is equivalent to the stochastic description, but of limited practical use due to its infinite-dimensional nature. However, it can be used as a starting point to obtain information on the statistics of the stochastic process. In general, we denote the expectation of a function of the number of molecules, $f(n)$, by

$$F(t) := \mathbb{E}[f(A(t))] = \sum_{n=0}^{\infty} f(n)p_n(t), \quad (4.14)$$

which amounts to a weighted average of f , weighted by the probability density for n .

Let us consider the mean behaviour. The expected number of molecules of type A in the system at any given time is given by the (deterministic) quantity

$$M(t) := \mathbb{E}[A(t)] = \sum_{n=0}^{\infty} np_n(t). \quad (4.15)$$

To obtain an ordinary differential equation for the evolution of $M(t)$, we start from (4.13), and write

$$\dot{M}(t) = \sum_{n=0}^{\infty} n \dot{p}_n(t). \quad (4.16)$$

Using (4.13), we get

$$\begin{aligned} \dot{M}(t) &= \sum_{n=0}^{\infty} n \dot{p}_n(t) \\ &= \sum_{n=0}^{\infty} n \{ [k_1(n-1)(n-2) + k_3] p_{n-1}(t) + [k_2(n+1)n(n-1) + k_4(n+1)] p_{n+1}(t) \\ &\quad - [k_1n(n-1) + k_3 + k_2n(n-1)(n-2) + k_4n] p_n(t) \}. \end{aligned}$$

We regroup the terms per reaction. We first consider reaction 3, for which we have

$$\begin{aligned} \sum_{n=0}^{\infty} k_3 n (p_{n-1}(t) - p_n(t)) &= k_3 \sum_{n=0}^{\infty} ((n+1)p_n(t) - np_n(t)) \\ &= k_3 \sum_{n=0}^{\infty} p_n(t) = k_3. \end{aligned} \quad (4.17)$$

Next, consider reaction 4. Here, we have

$$\begin{aligned} \sum_{n=0}^{\infty} k_4 n ((n+1)p_{n+1}(t) - np_n(t)) &= k_4 \sum_{n=0}^{\infty} ((n-1)np_n(t) - n^2 p_n(t)) \\ &= -k_4 \sum_{n=0}^{\infty} np_n(t) = -k_4 M(t). \end{aligned} \quad (4.18)$$

Similarly, we obtain for the reactions 1 and 2,

$$\sum_{n=0}^{\infty} k_1 n ((n-1)(n-2)p_{n-1}(t) - n(n-1)p_n(t)) = k_1 \sum_{n=0}^{\infty} n(n-1)p_n(t), \quad (4.19)$$

$$\sum_{n=0}^{\infty} k_2 n ((n+1)n(n-1)p_{n+1}(t) - n(n-1)(n-2)p_n(t)) = -k_2 \sum_{n=0}^{\infty} n(n-1)(n-2)p_n(t), \quad (4.20)$$

resulting in the equation

$$\dot{M}(t) = k_3 - k_4 M(t) + k_1 \sum_{n=0}^{\infty} n(n-1)p_n(t) - k_2 \sum_{n=0}^{\infty} n(n-1)(n-2)p_n(t). \quad (4.21)$$

We notice immediately that the evolution of the mean $M(t)$ does not only depend on the mean itself, but also on higher-order statistics of the distribution, such as the second moment

$$\mathbb{E} \left[A(t)^2 \right] = \sum_{n=0}^{\infty} n^2 p_n(t).$$

We can proceed similarly to derive an evolution equation for the variance

$$V(t) = \mathbb{E} \left[(A(t) - M(t))^2 \right]. \quad (4.22)$$

However, we should expect even higher order moments to appear in the corresponding righthand side, leading to an infinite cascade. If we want to obtain an evolution law in terms of only the mean number of molecules $M(t)$, we will therefore be obliged to resort to an approximation.

Let us now take a more detailed look into the fluctuations around the mean, as measured by the variance $V(t)$. We are specifically interested in systems with large numbers of molecules, for which we assume a mean-field approximation to hold. We therefore introduce a characteristic number of molecules per unit volume N , and look at the *concentrations* $a(t) = A(t)/N$ and $\rho(t) = M(t)/N$. Then, the variance can be written as

$$V(t) = \frac{1}{N^2} \mathbb{E} \left[(a(t) - \rho(t))^2 \right]. \quad (4.23)$$

We conclude that fluctuations around the mean concentration become negligible as the number of molecules per unit volume N tends to infinity. In that limit, the quantized concentration n/N approaches a continuous variable a , and the probability distribution $(p_n(t))_{n=0}^{\infty}$ approaches a continuous probability distribution $p(a, t)$, $a \in [0, \infty)$. The fact that the fluctuations vanish in that limit implies that $p(a, t) = \delta_{M(t)}(a)$, i.e., the concentration $a(t) = M(t)$ almost surely. (We note that the above,




Fig. 4.1 Simulation of the Schlögl model (4.1): number of molecules as a function of time for a single realization of the stochastic simulation, using Algorithm 2 (solid) and the corresponding mean-field approximation (dashed). *Left* short time-scale; *right* long time-scale. Parameter values are in the text

rather heuristic, reasoning can be turned into a rigorous mathematical theory, see, e.g., [24].)

Using the above reasoning in the limit when N tends to infinity, one can derive from Eq. (4.21) the mean-field approximation for $\rho(t)$ as

$$\dot{\rho} = k_3 - k_4\rho + k_1\rho^2 - k_2\rho^3. \quad (4.24)$$

From the above derivation, one concludes that stochastic modeling of chemical reactions is mainly useful when the number of molecules present is not too large. In that case, results can deviate from the mean-field approximation for two reasons: (i) stochastic fluctuations around the mean can become important; and (ii) due to the present nonlinearities, the ensemble average of a large number of systems with a low number of molecules does not necessarily follow the mean-field behaviour (Fig. 4.1).

4.2.4 Numerical Simulations for the Schlögl Model

Let us now consider system (4.1) with (non-dimensionalized) reaction rates $k_1 = 0.18$, $k_2 = 2.5 \times 10^{-4}$, $k_3 = 2200$ and $k_4 = 37.5$. We simulate one realization of a stochastic simulation using Algorithm 2, as well as a forward Euler simulation of the mean field Eq. (4.24) with time step $\Delta t = 5 \times 10^{-3}$. As an initial condition, we take $A(0) = 0$. For the chosen parameter values, Eq. (4.24) has two stable steady state, $A_1 = 100$ and $A_2 = 400$, and an unstable steady state $A_u = 220$. Thus, the Schlögl model represents a bistable system. Figure 4.1 shows the results. The mean-field

equation converges to one of the two stable equilibria, depending solely on the initial condition. When $A(0) \in [0, A_u]$, the solution converges to A_1 ; when $A(0) > A_u$ the solution converges to A_2 . On short time-scales (left figure), the stochastic simulation fluctuates around the stable equilibrium of the mean-field equation. However, over long time scales (right figure), we notice that the fluctuations cause the system to occasionally switch between the two steady states. Such occasional switches are called *rare events*, and they occur when large fluctuations can occur; for instance, in gene regulatory networks [20, 28]. If one is interested in quantities such as mean switching times, one cannot use a mean-field approximation and needs to resort to stochastic simulation. We refer to [52] for theoretical and computational work related to stochastic simulation of rare events.

4.3 Stochastic Modeling of Advection-Diffusion Processes

In the previous section, all systems were assumed to be well mixed, such that only the temporal evolution of concentrations needed to be considered. In many processes, however, interesting dynamics arises from spatial heterogeneity. In a biological context, one can, for instance, think of bacterial chemotaxis, tumor growth, or bone tissue engineering. In this section, we give an overview of individual-based modeling techniques for biological systems consisting of moving individuals that are able to reproduce and die.

In Sect. 4.3.1, we consider the positions of the individuals to be discrete (on a lattice). We briefly discuss cellular automata and Markov jump processes, and give references to the corresponding literature. In Sect. 4.3.2, we introduce Brownian motion and stochastic differential equations (SDEs), which are used for space/time continuous modeling of random motion. Subsequently, in Sect. 4.3.3, we discuss the equivalence between the SDE for an individual particle and the (deterministic) Fokker–Planck equation that describes the evolution of the particle density.

4.3.1 Discrete-Space Modeling

Several techniques exist for the discrete stochastic modeling of biological particles. We briefly discuss cellular automata and Markov jump processes.

Cellular automata In cellular automata, one considers space to be discretized as a grid, say $\Pi(x) = \{x_n\}_{n=0}^N$ in one space dimension. The state is then given as the number of particles A_n^k at each grid location x_n at each discrete moment in time t^k . (Clearly, one can incorporate the presence of particles of multiple types.) The cellular automaton then defines an evolution law that determines the state $A^{k+1} = (A_n^{k+1})_{n=0}^N$ at time t^{k+1} from the state A^k . This evolution law can contain reactions with associated rates, as in the time-discrete schemes in Sect. 4.2.1. Movement on the grid can be modeled using *hops*, which can be seen as a reaction event (with an

associated reaction rate) in which a particle moves from one lattice site to another. Then, the algorithmic structure that resulted in Algorithm 1 can be reused to model advection-diffusion processes [17].

One particularly appealing feature of cellular automata is their modeling flexibility: in defining the evolution laws, any set of rules can be allowed. One can, for instance, change the rules depending on the number of neighbors, or let the evolution of individual particles depend on some internal state variable (see also Sect. 4.4). We refer to [21, 33, 39] for a number of cellular automata models in the context of tumor growth, and to [12, 31] for examples in bone tissue engineering.

Markov jump processes When keeping a discrete state space, but allowing time to be continuous, one ends up with a Markov jump process. Consider a particle with position $X(t)$ that is allowed to reside on any position in the lattice $\Pi(x)$. Given that the current position $X^k = x_n$ at time t^k , we can introduce a propensity $\alpha_{n,m}^k$, $1 \leq m \leq N$, such that $\alpha_{n,m}^k dt$ represents the probability that the particle moves from x_n to x_m in the infinitesimal time interval $[t, t + dt)$. Then, such a movement can be added to the table of reactions in Algorithm 2, and the same algorithm can be used.

4.3.2 Stochastic Differential Equations (SDEs)

When space and time are allowed to be continuous, the corresponding model becomes an SDE. In this section, we start from a definition of Brownian motion (Sect. 4.3.2.1). We then proceed to the construction of general SDEs in the Itô sense (Sect. 4.3.2.2) and discuss numerical methods (Sect. 4.3.2.3). We conclude in Sect. 4.3.2.4 with a numerical example that illustrates the results. The exposition is partly based on [29].

4.3.2.1 Brownian Motion

A scalar standard Brownian motion essentially describes an unbiased random walk in one space dimension. (Generalizations to multiple space dimensions are, of course, straightforward.) The description is phenomenological. We define a Brownian motion as a random variable $W(t)$, continuous in time $t \in [0, T]$, that satisfies the following conditions:

1. $W(0) = 0$ (with probability 1);
2. For $0 \leq s \leq t \leq T$, the *increment* $W(t) - W(s)$ is a normally distributed random variable with mean 0 and variance $t - s$, i.e., $W(t) - W(s) \sim \sqrt{t - s}N(0, 1)$, where $N(0, 1)$ denotes a standard normally distributed random variable (with mean 0 and variance 1);
3. For $0 \leq s \leq t \leq u \leq v \leq T$ are independent.

There are several ways of justifying this definition. One way is to start by defining a grid $\Pi(x) = \{-N\Delta x, \dots, 0, \dots, N\Delta x\}$ and letting a particle move one grid cell to the left or to the right (each with probability 1/2) in each time step of size Δt ,




Fig. 4.2 Realizations of Brownian motion: $W(t)$ as a function of time, using (4.25) with $\Delta t = 5 \times 10^{-2}$

starting at position $X(t_0) = 0$ at time $t_0 = 0$. When choosing $\Delta x^2 = \Delta t$ and taking the limit of $\Delta t \rightarrow 0$, we obtain the standard Brownian motion.

To visualize realizations of Brownian motion, we consider time-discretized Brownian paths, generated as $W^k \approx W(t^k)$, with $t^k = k\Delta t$ via time-stepping,

$$W^{k+1} = W^k + \sqrt{\Delta t}\xi^k, \quad \xi^k \sim N(0, 1), \quad (4.25)$$

where the random numbers ξ^k are independent and identically distributed (i.i.d.). In Fig. 4.2, we show 10 realizations of a Brownian motion with $\Delta t = 5 \times 10^{-2}$.

The figure illustrates some properties of Brownian motion. A proper definition of the probability spaces generated by Brownian motion is out of the scope of the present chapter. Let us just suffice by stating that, when writing $\mathbb{E}[\cdot]$, we imply the mean with respect to all possible Brownian paths $W(t)$. It can be proved (see, e.g., [11]) that the expected value of a Brownian motion $\mathbb{E}[W(t)] = 0$ for any $t \in [0, T]$. This is easily seen intuitively, as the Brownian increments do not have a preferred direction. Moreover, we observe that $\mathbb{E}[W(t)^2] = t$. Both properties can easily be proved in the time-discrete setting of Eq. (4.25) by using the basic rules of probability on the normal random variables ξ^k . A final property is that Brownian motion is nowhere differentiable with probability 1. This can be understood by realizing that

$$\text{Var} \left[\frac{W(t + \Delta t) - W(t)}{\Delta t} \right] = \frac{1}{\Delta t^2} \text{Var} [W(t + \Delta t) - W(t)] = \frac{1}{\Delta t}. \quad (4.26)$$

Note that, while $\mathbb{E}[W(t + \Delta t)] = \mathbb{E}[W(t)] = 0$, we have $\mathbb{E}[|W(t + \Delta t) - W(t)|] = \mathcal{O}(\Delta t^{1/2})$, which is proportional to the standard deviation.

4.3.2.2 Itô Stochastic Differential Equations

Often, the evolution of the position $X(t)$ of a particle is composed of a deterministic (mean) component, supplemented with a stochastic (fluctuating) part, modeled using a stochastic differential equation of the form

$$dX(t) = a(X(t))dt + b(X(t))dW(t), \quad X(0) = X_0, \quad (4.27)$$

in which a is called the *drift coefficient*, b is the diffusion coefficient, $W(t)$ is a Brownian motion and X_0 is the initial condition. For instance, one can consider the particle to model the position of a bacterium that is biasing its random motion to favor directions that are in line with the gradient of a chemoattractant. Then, the drift coefficient $a(X(t))$ models a preferred direction, whereas the second term takes into account the randomness of the motion. The diffusion coefficient $b(X(t))$ then defines, at the position $X(t)$, how strongly the evolution is affected by the Brownian motion $W(t)$.

Under mild assumptions on a and b , the stochastic differential equation (SDE) (4.27) has exactly one solution per Brownian path $W(t)$. To obtain this solution, one first needs to make sense of Eq. (4.27), something that will turn out to be nontrivial. Consider a classical solution,

$$X(t) = X_0 + \int_0^t a(X(s))ds + \int_0^t b(X(s))dW(s). \quad (4.28)$$

The first integral is a well-defined integral with respect to time; the second integral, however, is not well-defined and we need to be specific about its meaning.

Consider the integral to be defined via a Riemann sum using subintervals $[t^k, t^{k+1}]$, with $t^k = k\Delta t$ and $0 \leq k \leq K$, $K\Delta t = t$,

$$\int_0^t b(X(s))dW(s) = \lim_{\Delta \rightarrow 0} \sum_{k=1}^K b(X(s^k))\Delta W^k, \quad (4.29)$$

in which $\Delta W^k = W(t^{k+1}) - W(t^k)$ and $s^k \in [t^k, t^{k+1}]$. It is now easy to see, using only standard rules of probability, that the choice of s^k has a significant influence on the value of the integral. (This is due to the fact that $W(s)$ does not have bounded variations, which is related to (4.26).) The most common interpretation (the Itô interpretation) is obtained by choosing $s^k = t^k$, i.e., a left-point rule. In that case, we have

$$\begin{aligned} \mathbb{E} \left[\int_0^t b(X(s))dW(s) \right] &= \lim_{\Delta \rightarrow 0} \sum_{k=0}^{K-1} \mathbb{E} \left[b(X(t^k))\Delta W^k \right] \\ &= \lim_{\Delta \rightarrow 0} \sum_{k=0}^{K-1} \mathbb{E} \left[b(X(t^k)) \right] \mathbb{E} \left[\Delta W^k \right] = 0, \end{aligned}$$

where the first equality is due to the definition of the Itô integral and the linearity of the expectation, the second equality is due to independence of $b(X(t^k))$ and ΔW^k , and the last equality is due to the definition of Brownian motion.

A second popular interpretation (the Stratonovich interpretation) of (4.29) is obtained when choosing $s^k = (t^k + t^{k+1})/2$. It is clear that the above reasoning then can no longer be used, as $b(X(s^k))$ cannot be independent of ΔW^k . The resulting stochastic integral will then, in general, take a different value. We denote the Stratonovic integral as

$$\int_0^t b(X(s)) \circ dW(s)$$

and compute, as an example,

$$\mathbb{E} \left[\int_0^t W(s) \circ dW(s) \right] = \lim_{\Delta \rightarrow 0} \sum_{k=0}^{K-1} \mathbb{E} \left[W((t^k + t^{k+1})/2) \Delta W^k \right].$$

To compute this integral, we need to evaluate $W((t^k + t^{k+1})/2)$. It can be shown that this quantity is statistically equivalent to the quantity

$$\frac{W(t^k) + W(t^{k+1})}{2} + \Delta Z^k,$$

in which $\Delta Z^k \sim N(0, \Delta t/4)$ and independent of $W(t^k)$ and $W(t^{k+1})$. We can thus replace $W((t^k + t^{k+1})/2)$ by this alternative in computing the expectation, and we obtain

$$\mathbb{E} \left[\int_0^t W(s) \circ dW(s) \right] = \lim_{\Delta \rightarrow 0} \sum_{k=0}^{K-1} \left(\mathbb{E} \left[\left(\frac{W(t^k) + W(t^{k+1})}{2} \right) \Delta W_k \right] + \mathbb{E} \left[\Delta Z^k \Delta W^k \right] \right).$$

Since ΔZ^k and ΔW^k are independent, the second term is zero, and we continue with only the first term:

$$\begin{aligned} \mathbb{E} \left[\int_0^t W(s) \circ dW(s) \right] &= \lim_{\Delta \rightarrow 0} \sum_{k=0}^{K-1} \mathbb{E} \left[\left(\frac{W(t^k) + W(t^{k+1})}{2} \right) (W(t^{k+1}) - W(t^k)) \right] \\ &= \lim_{\Delta \rightarrow 0} \frac{1}{2} \sum_{k=0}^{K-1} \mathbb{E} \left[(W(t^{k+1}))^2 - (W(t^k))^2 \right] = \frac{1}{2} (W(t))^2, \end{aligned}$$

which is clearly different from the corresponding Itô integral (which is zero). For more details, we refer to [29] and references therein. In this chapter, we will always work with the Itô interpretation.

If $W(t)$ does not have bounded variations, then neither does $X(t)$. Consequently, also $X(t)$ will be nowhere differentiable (with probability 1). In particular, we have

$$\mathbb{E} [|X(t + \Delta t) - X(t)|] = \mathcal{O}(\Delta t^{1/2}).$$

Remark 4 (Further reading) We have deliberately kept the introduction on SDEs very brief. For more information, we refer the interested reader to the excellent books [19, 34].

Remark 5 (Reproduction and death) To model individuals that are also able to reproduce and die, one needs to add a (potentially stochastic) process that determines for each individual the moment at which it reproduces (and thus generates an additional individual) or dies (and is therefore removed from the system). To this end, one can, for instance, use the Markov processes that were discussed in Sect. 4.2 in the context of chemical reactions. Section 4.7 will contain additional modeling techniques that are of a more mechanistic nature.

4.3.2.3 Euler-Maruyama Method

Once an SDE model is obtained for a specific problem, a (numerical) solution needs to be computed. The most straightforward way to discretize an SDE of the type (4.27) is by using the stochastic extension of the forward Euler method, called Euler-Maruyama,

$$X^{k+1} = X^k + a(X^k)\Delta t + b(X^k)\Delta W^k, \quad (4.30)$$

in which Δt is the time step, and ΔW^k is sampled from a normal distribution with zero mean and variance Δt , i.e., $\Delta W^k \sim N(0, \Delta t)$. The scheme can equivalently be written as

$$X^{k+1} = X^k + a(X^k)\Delta t + b(X^k)\sqrt{\Delta t}\xi^k, \quad (4.31)$$

with $\xi^k \sim N(0, 1)$.

For a deterministic system, the convergence behaviour of the forward Euler method can be derived in a straightforward manner: given the numerical solution $X^K \approx X(t^K)$, the error $X^K - X(t^K)$ can be bounded as

$$\left| X^K - X(t^K) \right| \leq C \Delta t,$$

where K and Δt are varied simultaneously such that $K \Delta t = t^*$.

In the SDE case, this is no longer true. In fact, since both the numerical solution X^K and the exact solution $X(t^K)$ are random, only statistical statements can be made about the numerical error. One can immediately come up with two different definitions. We can define the *strong error* at time t^K as

$$e_{\Delta t}^K = \mathbb{E} \left[\left| X^K - X(t^K) \right| \right], \quad (4.32)$$

i.e., the expectation of the absolute value of the error on individual trajectories. Since it measures the “mean of the error”, averaged over all possible Brownian motions, the strong error gives an indication on the size of the error on an individual trajectory (defined by the Brownian path that generated it). Alternatively, one can define the *weak error* as

$$E_{\Delta t}^K = \left| \mathbb{E} \left[X^K \right] - \mathbb{E} \left[X(t^K) \right] \right|, \quad (4.33)$$

or, more generally, for an arbitrary function f in an appropriate function class,

$$E_{\Delta t}^K[f] = \left| \mathbb{E} \left[f(X^K) \right] - \mathbb{E} \left[f(X(t^K)) \right] \right|, \quad (4.34)$$

i.e., the error in the expectation of the function f when computed using time-discretized trajectories.

In general, these types of error are not the same, and also the order of convergence (as a function of Δt) differs. For the Euler-Maruyama method, we have

$$e_{\Delta t}^K \leq C \Delta t^{1/2}, \quad E_{\Delta t}^K \leq C \Delta t, \quad (4.35)$$

i.e., the Euler-Maruyama method has a strong order 1/2 and a weak order 1. Proving these orders would lead us too far. In this chapter, we simply illustrate this result numerically (see Sect. 4.3.2.4).

Remark 6 (Stability) The order of convergence of the Euler-Maruyama method only gives information on the asymptotic behavior of the error as Δt tends to zero. In practice, one will always take a finite time step. In that case, one needs the time step to be such that the Euler-Maruyama method is stable, i.e., loosely speaking, one needs to ensure that the numerical solution does not blow up for the chosen value of Δt when the exact solution remains bounded. While stability of time integration is a relatively straightforward concept for deterministic ODEs [27], this is no longer true for SDEs. As for convergence, multiple definitions of stability exist, see, e.g., [29, 32] for more details.

Remark 7 (Higher-order methods) Due to the presence of stochastic integrals, the definition of higher-order methods for SDEs is far more complicated than for ODEs. We refer to [32] for details.

4.3.2.4 Numerical Example: Geometric Brownian Motion

To illustrate the most important concepts in the previous sections, we perform some numerical experiments on a simple linear SDE, namely a geometric Brownian motion,

$$dX(t) = \lambda X(t)dt + \mu X(t)dW(t), \quad (4.36)$$




Fig. 4.3 Simulations of a geometric Brownian motion (4.36). *Top left* a single trajectory using the Euler-Maruyama method (4.31) with time step $\Delta t = 1/2^8$ (solid), compared to the exact solution for the same Brownian path (dashed); *Top right* four individual trajectories (dashed) and the empirical average of $M = 1000$ trajectories; *Bottom left* strong error of Euler-Maruyama as a function of Δt (solid), compared to the predicted strong order $1/2$ (dashed); *Bottom right* weak error of Euler-Maruyama as a function of Δt (solid), compared to the predicted weak order 1 (dashed). Remaining parameters are in the text

which we discretize with the Euler-Maruyama method (4.31) with time step Δt . For this SDE, an exact solution is known analytically (given the Brownian path $W(t)$), and given by

$$X(t) = X_0 \exp \left(\left(\lambda - \mu^2/2 \right) t + \mu W(t) \right). \tag{4.37}$$

In our experiments, we choose $\lambda = 2$, $\mu = 1$ and $X(0) = X_0 = 1$ with probability 1. The example is based on [29], in which also Matlab code can be found.

We first compare, for a single realization of $W(t)$, the exact solution $X(t)$ with the numerical solution $(X^k)_{k=0}^K$, with X^k obtained via Euler-Maruyama (4.31) with time step $\Delta t = 1/2^5$ on the time interval $t \in [0, 1]$ (hence $K \Delta t = 1$). The results are shown in Fig. 4.3, top left. We clearly see a discretization error.

To quantify this discretization error, we repeat the simulation for $M = 1000$ realizations of the Brownian path, $(W_m(t))_{m=1}^M$, resulting in M trajectories $(X_m(t))_{m=1}^M$, and compute an approximation to the strong and weak errors (4.32) and (4.33) as

$$\hat{e}_{\Delta t}^K = \hat{\mathbb{E}}_M \left[\left| X^K - X(t^K) \right| \right] = \frac{1}{M} \sum_{m=1}^M \left| X_m^K - X_m(t^K) \right|, \quad (4.38)$$

$$\hat{E}_{\Delta t}^K = \hat{\mathbb{E}}_M \left[X^K \right] - \hat{\mathbb{E}} \left[X(t^K) \right] = \frac{1}{M} \left| \sum_{m=1}^M \left(X_m^K - X_m(t^K) \right) \right|. \quad (4.39)$$

To achieve this, we first generate Brownian paths with time step $\Delta t = 1/2^9$, and subsequently use these Brownian paths to perform Euler-Maruyama simulations with time step $R\Delta t$, $R = 1, 2, 4, 8, 16$. The results are shown in Fig. 4.3, bottom. We clearly observe the predicted theoretical strong order 1/2 and weak order 1.

Remark 8 (Statistical error) The estimates (4.38) and (4.39) contain statistical error due to the finite number M of realizations. The problem and method parameters have been chosen such that this statistical error is negligible with respect to the time discretization error.

Finally, we look at the evolution of $F(t) := \mathbb{E}[X(t)]$ as a function of time. Being the expectation of the (time-dependent) random variable $X(t)$, $F(t)$ is a deterministic function of time. Figure 4.3 shows that $F(t)$ (unlike $X(t)$) is a smooth, differentiable function of time. This observation will be elaborated in the next section.

4.3.3 Population-Level Dynamics and Fokker-Planck Equation

Usually, one is not interested in the detailed stochastic behaviour of a single individual (cell, bacterium), but rather in the evolution of a large *population* of such cells. One then has given an initial *density* $\rho_0(x)$ of individuals as a function of space $x \in D \subset \mathbb{R}$, with D the domain. We interpret $\rho_0(x)$ as a probability density (this can always be done with a proper normalization), i.e., for an individual particle with position X_0 at time $t = 0$, we have

$$\Pr(x \leq X_0 < x + dx) = \rho_0(x)dx.$$

Note that lowercase x represents a possible position in the domain D , whereas uppercase X denotes the (random) position of an individual cell.

Each of the individuals behaves according to (4.27), generating a path $X(t)$. The question then becomes: defining the time-dependent density $\rho(x, t)$ as

$$\Pr(x \leq X(t) < x + dx) = \rho(x, t)dx,$$

can one obtain a corresponding evolution equation for $\rho(x, t)$? In this section, we will show that $\rho(x, t)$ satisfies a advection-diffusion partial differential equation (PDE), the Fokker-Planck equation. We perform the derivation only for a pure Brownian motion, after which we simply state the result for the general SDE (4.27). Our exposition closely follows [11].

In the pure diffusion case, we start from a (scaled) Brownian motion,

$$dX(t) = b dW(t), \quad (4.40)$$

with a constant scaling parameter $b > 0$.

Given the probability density $\rho(x, t)$ at time t , we can write the density $\rho(x, t + \Delta t)$ at time $t + \Delta t$ as

$$\rho(x, t + \Delta t) = \int_{-\infty}^{\infty} \rho(x + y, t) \cdot \psi(x, y, \Delta t) dy, \quad (4.41)$$

where the *transition probability kernel* $\psi(x, y, \Delta t)$ is the probability of ending up at position x at time $t + \Delta t$, given that one started at position $x + y$ at time t , i.e.,

$$\psi(x, y, \Delta t) = \Pr(X(t + \Delta t) \in [x, x + dx] | X(t) \in [x + y, x + y + dx]) \quad (4.42)$$

Eq. (4.41) states that the probability of finding a particle at position x at time $t + \Delta t$ is equal to the probability of finding the particle at position $x + y$ at time t , multiplied by the probability of moving from $x + y$ to x during the time step of size Δt , and this integrated over every possible value of y (i.e., integrated over every possible position $x + y$ at time t).

Remark 9 (Notation) It may seem odd to introduce the auxiliary variable y , instead of simply integrating over all possible original positions $z = x + y$. This is done because we intend to make a Taylor expansion of $\rho(x + y, t)$ around $\rho(x, t)$.

Let us now obtain an expression for $\psi(x, y, \Delta t)$. We know that $-y$ is the increment that was generated to move from $x + y$ to x . This increment is normally distributed with mean 0 and variance $b^2 \Delta t$, and therefore,

$$\psi(x, y, \Delta t) = \frac{1}{b\sqrt{2\pi\Delta t}} \exp\left(-\frac{y^2}{2b^2\Delta t}\right). \quad (4.43)$$

Now, we are ready to expand Eq. (4.41) by performing a Taylor expansion of $\rho(x + y, t)$ around $\rho(x, t)$,

$$\rho(x + y, t) = \rho(x, t) + y\partial_x\rho(x, t) + \frac{y^2}{2}\partial_{xx}\rho(x, t) + \text{h.o.t.} \quad (4.44)$$

in which h.o.t. stands for “higher order terms”. This leads to

$$\begin{aligned}
 \rho(x, t + \Delta t) &= \int_{-\infty}^{\infty} \rho(x, t) \psi(x, y, \Delta t) dy + \int_{-\infty}^{\infty} y \partial_x \rho(x, t) \psi(x, y, \Delta t) dy \\
 &\quad + \int_{-\infty}^{\infty} \frac{y^2}{2} \partial_{xx} \rho(x, t) \psi(x, y, \Delta t) dy + \text{h.o.t.} \\
 &= \rho(x, t) \underbrace{\int_{-\infty}^{\infty} \psi(x, y, \Delta t) dy}_{I_1} + \partial_x \rho(x, t) \underbrace{\int_{-\infty}^{\infty} y \psi(x, y, \Delta t) dy}_{I_2} \\
 &\quad + \frac{1}{2} \partial_{xx} \rho(x, t) \underbrace{\int_{-\infty}^{\infty} y^2 \psi(x, y, \Delta t) dy}_{I_3} + \text{h.o.t.} \tag{4.45}
 \end{aligned}$$

We will now separately look into each of the terms $I_{1,2,3}$. First, we have

$$I_1 = 1, \tag{4.46}$$

since $\psi(x, y, \Delta t)$ is independent of x and a probability density for y . Next, we have

$$I_2 = 0, \tag{4.47}$$

since y is an odd function and $\psi(x, y, \Delta t)$ is even. Finally, we obtain

$$I_3 = \text{Var} [y] = b^2 \Delta t, \tag{4.48}$$

in which $\text{Var} [y]$ is to be interpreted as the variance of y with respect to the probability density ψ . This leads to

$$\rho(x, t + \Delta t) = \rho(x, t) + \frac{b^2}{2} \Delta t \partial_{xx} \rho(x, t) + \text{h.o.t.}, \tag{4.49}$$

which is a time-discretized version of the diffusion equation. Taking the limit of $\Delta t \rightarrow 0$, we obtain the desired result:

$$\partial_t \rho(x, t) = \frac{b^2}{2} \partial_{xx} \rho(x, t). \tag{4.50}$$

Thus, the density evolves according to the diffusion equation. This is to be expected, as the particles have no preferred direction and will therefore spread evenly. This derivation explains why the Brownian motion is also called a diffusion process.

In the more general case of the SDE (4.27), the drift term will introduce a systematic bias in the motion of individual particles, resulting in an advective behaviour of the density $\rho(x, t)$. It can be shown that, in that case, the density $\rho(x, t)$ satisfies the

advection-diffusion equation

$$\partial_t \rho(x, t) + \partial_x (a(x)\rho(x, t)) = \frac{1}{2} \partial_{xx} \left((b(x))^2 \rho(x, t) \right). \quad (4.51)$$

We refer to [19, 34] for details.

4.4 More Realistic Microscopic Processes

Whereas stochastic differential equations are useful to describe a wide range of stochastic processes in biological applications, they remain in essence *phenomenological*, and thus descriptive. Even if these models have predictive power, it is often impossible to use such models for a detailed understanding of the *mechanism* that generate the dynamics. In this section, we therefore discuss more mechanistic velocity-jump processes and their relation to bacterial chemotaxis. The section follows the same structure as that of the previous sections: we first discuss the stochastic individual-based model (Sect. 4.4.1), after which we continue with an equivalent continuum description (Sect. 4.4.2). In Sect. 4.4.3, we relate the resulting stochastic processes with the SDEs of Sect. 4.3.2, and we conclude with some bibliographical remarks on generalizations in Sect. 4.4.4.

4.4.1 Velocity-Jump Processes for Bacterial Chemotaxis

Generally, the motion of flagellated bacteria consists of a sequence of *run* phases, during which a bacterium moves in a straight line at constant speed. The bacterium changes direction in a *tumble* phase, which is typically much shorter than the run phase and acts as a reorientation. Hence, the motion of an individual bacterium can be modeled as a velocity-jump process. To bias movement towards regions with high concentration of chemoattractant, the bacterium adjusts its turning rate to increase, resp. decrease, the chance of tumbling when moving in an unfavorable, resp. favorable, direction [2, 50]. The velocity-jump models described here are based on [15] and [44, 45].

We consider bacteria that are sensitive to the concentration of a chemoattractant $S(x) \geq 0$ for $x \in \mathbb{R}^d$, where x is the present position of the bacterium. While we do not consider time dependence of chemoattractant via production or consumption by the bacteria, a generalization to this situation is straightforward, at least for the definition of the models and the numerical method. Bacteria move with a constant speed v (run), and change direction at random instances in time (tumble), in an attempt to move towards regions with high chemoattractant concentrations.

The position of an individual bacterium is given by $X(t)$, its velocity is

$$\frac{dX_c(t)}{dt} = \epsilon V_c(t), \quad V_c(t) \in \mathbb{V} = \mathbb{S}^{d-1},$$

with \mathbb{S}^{d-1} the unit sphere in \mathbb{R}^d . Hence, $V_c(t)$ represents the direction and the scaling parameter $\epsilon > 0$ represents the size of the velocity. (The reason for the introduction of the subscript c will become clear in Sect. 4.6.) The velocity of each bacterium is switched at random jump times $(T_c^k)_{k \geq 1}$ that are generated via a Poisson process with a time dependent turning rate $\lambda_c^\epsilon(x, v)$ that depends on the bacterium's current position and velocity. The new velocity at time T_c^k is generated at random according to a centered probability distribution $\mathcal{M}(dv)$ with $\int v \mathcal{M}(dv) = 0$, typically

$$\mathcal{M}(dv) = \sigma_{\mathbb{S}^{d-1}}(dv),$$

where $\sigma_{\mathbb{S}^{d-1}}$ is the uniform distribution on the unit sphere.

The turning rate is assumed to satisfy

$$0 < \lambda_{\min} \leq \lambda_c^\epsilon(x, v) \leq \lambda_{\max}, \tag{4.52}$$

as well as, for small values of ϵ ,

$$\lambda_c^\epsilon(x, v) := \lambda^0 - \epsilon A_\epsilon^T(x)v + O(\epsilon^2). \tag{4.53}$$

Typically, $\lambda_c^\epsilon(x, v)$ is a function of $\nabla S(x)$, so that the model (4.53) may describe a large bacterium that is able to directly sense chemoattractant gradients. When the turning rate (4.53) is proportional to $\nabla S(x)v$, it can be interpreted as follows: the rate at which a bacterium will change its velocity direction depends on the alignment of the velocity with the gradient of the chemoattractant concentration $\nabla S(x)$, resulting in a transport towards areas with higher chemoattractant concentrations.

The resulting stochastic process can be written as

$$\begin{cases} \frac{dX_c(t)}{dt} = \epsilon V_c(t), \\ \int_{T_c^k}^{T_c^{k+1}} \lambda_c^\epsilon(X_c(t), V_c(t)) dt = \theta^{k+1}, \\ V_c(t) = \mathcal{V}^k \quad \text{for } t \in [T_c^k, T_c^{k+1}], \end{cases} \tag{4.54}$$

with initial condition $X(0), \mathcal{V}(0) \in \mathbb{R}^d$. In (4.54), $(\theta^k)_{k \geq 1}$ denote i.i.d. random variables with normalized exponential distribution, and $(\mathcal{V}^k)_{k \geq 1}$ denote i.i.d. random variables with distribution $\mathcal{M}(dv)$.

4.4.2 Population-Level Dynamics and Kinetic Equations

At the population level, there are two main differences with respect to the Fokker-Planck Eq. (4.51) for SDEs. First, the stochastic process is written in terms of positions x and velocities v , implying that the density of interest will be a density $p_c(x, v, t)$ in position-velocity phase space. Second, the stochastic process is discontinuous in the velocity component, which will result in a “collision operator” that models the discontinuous velocity changes probabilistically.

The resulting evolution equation for the density $p_c(x, v, t)$ turns out to be a Boltzmann equation with BGK-type collision operator,

$$\partial_t p_c + \epsilon v \cdot \nabla_x p_c = (R(\lambda_c^\epsilon p) - \lambda_c^\epsilon p_c), \quad (4.55)$$

where

$$R(p_c) := \int_{\mathbb{V}} p_c(\cdot, v, \cdot) \mathcal{M}(dv)$$

is the operator integrating velocities with respect to \mathcal{M} , and λ_c^ϵ is defined as in (4.53). We will not derive this equation here, but instead refer the interested reader to [18] for the derivation of master equations associated to Markov jump processes. Here, we suffice by pointing out that the advection term models the effect of the velocity on positions, and the righthand side models the effect of the random velocity changes.

4.4.3 Coarse-Graining and Approximate Macroscopic Descriptions

The explicit modeling of these individual velocity changes is necessary to have a biologically relevant mechanistic description of bacterial motion. In general, however, one is not really interested in the detailed phase-space distribution $p_c(x, v, t)$, but rather in the the position density $\rho_c(x, t) = R(p_c(x, v, t))$, and this for (at least) two reasons: (i) it is usually impossible to obtain experimental data on the velocity distribution of the bacteria; and (ii) the bacteria typically travel only a microscopic distance between velocity changes, such that the observed macroscopic motion is the averaged effect on long time scales of a large number of velocity changes. One can expect the position density $\rho_c(x, t)$ to satisfy a partial differential equation advection-diffusion type, such as (4.51). The velocity-jump process (4.54) and the advection-diffusion SDE (4.27) are therefore related.

To consider the behaviour of Eq. (4.55) on long time scales and for small bacterial velocities, we let ϵ tend to 0 and introduce the rescaled time $\bar{t} = t\epsilon^2$. (Then, when \bar{t} is $O(1)$, this corresponds to a physical time t that is $O(1/\epsilon^2)$.) In that case, the position density $\rho_c(x, t)$ satisfies the advection-diffusion PDE

$$\partial_{\bar{t}}\rho_c = \frac{1}{\lambda_0} \operatorname{div}_x (D\nabla_x \rho_c - DA_0(x)\rho_c), \quad (4.56)$$

in which the diffusion matrix is given by the covariance of the Maxwellian distribution,

$$D = \int_{\mathbb{S}^{d-1}} v \otimes v \mathcal{M}(dv) \in \mathbb{R}^{d \times d}. \quad (4.57)$$

We refer to [35] for justifications of this result based on Hilbert expansions. The result implies that the position density obtained via (4.56) and via (4.55) are the same in the limit when ϵ tends to 0.

Additionally, it is shown in [45] that the position $X^c(\bar{t})$ of an individual trajectory generated by the stochastic process (4.54) converges to a trajectory of the SDE

$$dX_c(\bar{t}) = \frac{DA_0(X_c(\bar{t}))}{\lambda^0} d\bar{t} + \left(\frac{2D}{\lambda_0}\right)^{1/2} dW(\bar{t}), \quad (4.58)$$

where $\bar{t} \mapsto W_{\bar{t}}$ is a standard Brownian motion, as ϵ tends to zero. (Note that this second result implies the convergence of the position densities to a solution of (4.56), but not vice versa.)

4.4.4 Further Comments and Remarks

The main modeling limitation of the models discussed so far is that they deal with *non-interacting* particles, i.e., every individual follows its own stochastic path, independently of all other individuals present. Many generalizations exist to introduce interactions between particles, either as two-particle collision operators, via long-range interactions or via interactions of individual particles with the position density. Giving an overview of all these generalizations would lead too far. We refer to the two excellent books [49] and [40] and references therein.

4.5 Monte Carlo Simulation and Variance Reduction

4.5.1 The Need for Monte Carlo Simulation

In most situations, we are interested in the evolution of a large population of individuals (cells, bacteria). To simulate this evolution, two courses of action are possible:

- a (stochastic) Monte Carlo simulation of an ensemble of M realizations of the stochastic process (such as (4.27) or (4.54)), from which information on the population density can be obtained using histograms or kernel density estimation [47, 48];

- A deterministic (grid-based) simulation of the corresponding PDE (such as (4.51) or (4.55)).

Both options have advantages and drawbacks. The clear drawback of a Monte Carlo simulation is the appearance of statistical error on the obtained population density, which is absent in a deterministic simulation—the variance on the obtained result being of the order of $O(1/\sqrt{M})$ [7]. The drawback of a grid-based simulation is that the computational cost of mesh refinement depends crucially on the number of dimensions of the PDE. Considering a system of particles in 3 spatial dimensions, the kinetic equation (4.55) is a 6-dimensional PDE. Doubling the number of mesh points in each spatial dimension already increases the total number of unknowns by a factor of 2^6 , even if one can still take the same time step. In contrast, the computational cost of refining a Monte Carlo simulation is independent of the dimension of the problem: one can simply augment the number M of simulated particles.

In more realistic applications, the computational complexity of the PDE-based description can be even higher, due to several reasons. When particles also have internal state, the dimension of the kinetic equation (4.55) increases even further (see Sect. 4.6). Moreover, if the particles are interacting, the collision operator becomes non-local, requiring the evaluation of an integral over velocity space at each spatial mesh point. The situation becomes even more difficult when particles experience long-range interactions.

4.5.2 Variance Reduction Techniques

Because of the problems associated with simulating high-dimensional PDEs, Monte Carlo simulation is a viable alternative, provided one can control the variance of the simulation. As a consequence, there exists a large literature on variance reduction techniques. The most popular techniques can roughly be categorized in two classes: importance sampling and the use of a control variate. These techniques are well-established in the computation of integrals with respect to a known probability distribution, see [7] and references therein.

In importance sampling, the key idea is to adaptively sample the density of interest using *weighted particles*, such that more particles are placed (with correspondingly lower weights) in regions in which the variance is expected to be higher. Goal is to obtain a variance that is evenly distributed over the computational domain. With a control variate, the key idea is to compute an approximation to the quantity of interest *deterministically* based on the solution of a related but simpler problem, for instance analytically or by numerically solving a PDE of lower dimension. One then uses the Monte Carlo only to sample the *correction* with respect to the deterministically computed quantity.

Several research groups are currently working along these lines to develop hybrid Monte Carlo/PDE methods. We refer to [13, 14] and related papers and to [1, 42] and related papers in the context of the Boltzmann equation, and to [4] for an example in

the context of radiation transport. We have developed a strategy based on a control variate [44] that will be detailed for bacterial chemotaxis in Sect. 4.6. A related method that we are currently developing for tumor growth will be discussed in Sect. 4.7.

4.6 Application 1: Bacterial Chemotaxis

As a first application, we return to bacterial chemotaxis, which was also used in Sect. 4.4.1. Since many species are unable to sense chemoattractant gradients reliably due to their small size, adjustment of their turning rate to bias motion in favorable directions is often done via an intracellular mechanism that allows the bacterium to retain information on the history of the chemoattractant concentrations along its path [6]. The resulting model, which will be called the “internal state” or “fine-scale” model in this text, can be formulated as a velocity-jump process, combined with an ordinary differential equation (ODE) that describes the evolution of an internal state that incorporates this memory effect [16]. The probability density distribution of the velocity-jump process evolves according to a kinetic equation, in which the internal variables appear as additional dimensions. A direct deterministic simulation of this equation is therefore prohibitively expensive, and one needs to resort to a stochastic particle method.

Unfortunately, a direct fine-scale simulation using stochastic particles presents a large statistical variance, even in the diffusive asymptotic regime when ϵ is small. In that regime, the bacterial density of the fine-scale model is known explicitly to satisfy a Keller-Segel advection-diffusion equation. Consequently, it is difficult to assess accurately how the solutions of the fine-scale model differ from their advection-diffusion limit in intermediate regimes.

In this section, we discuss a numerical method to simulate individual-based models for chemotaxis of bacteria with internal dynamics with reduced variance, introduced in [44]. The variance reduction is based on a coupling technique (control variate): the main idea is to simultaneously simulate, using the same random numbers, a simpler, “coarse” process where the internal dynamics is replaced by a direct “gradient sensing” mechanism (see [2, 38, 41] for references on such gradient sensing models). The probability density of the latter satisfies a kinetic equation without the additional dimensions of the internal state, and converges to a similar advection-diffusion limit as the model with internal state, see e.g. [10, 35, 37, 45]. The precise coarse model will be (4.54) with a suitable choice for $A_\epsilon(x)$ in (4.53) (see later), such that the coarse and fine-scale model have *exactly* the same advection-diffusion limit.

We first discuss the fine-scale model with internal dynamics in Sect. 4.6.1. The model is a simplification (for expository purposes) of the more general model in [44, 45]. We describe the variance reduction technique in detail in Sect. 4.6.3. Some numerical results are given in Sect. 4.6.4. A detailed analysis of the method can be found in [44].

4.6.1 Bacterial Chemotaxis with Internal State

We again consider bacteria that are sensitive to the concentration of a chemoattractant $S(x) \geq 0$ for $x \in \mathbb{R}^d$. As in Sect. 4.4.1, the bacteria follow a velocity-jump process in which $X(t)$ is the position of an individual bacterium, and the normalized velocity is given by

$$\frac{dX(t)}{dt} = \epsilon V(t), \quad V(t) \in \mathbb{V} = \mathbb{S}^{d-1},$$

with \mathbb{S}^{d-1} the unit sphere in \mathbb{R}^d . Hence, $V(t)$ represents the direction and the parameter ϵ represents the size of the velocity. The difference with respect to the process (4.54) with direct gradient sensing is in the definition of the turning rate. As in [15], the turning rate is made to depend upon an internal state $y \in \mathbb{Y} \subset \mathbb{R}$ of each individual bacterium, which models the memory of the bacterium and is subject to an evolution mechanism attracted by the chemoattractant concentration $S(x)$. (The model in [44, 45] is more general and can take into account multiple chemoattractants and higher-dimensional internal states.)

The internal state adapts to the local chemoattractant concentration through an ODE,

$$\frac{dY(t)}{dt} = F_\epsilon(Y(t), S(X(t))), \quad (4.59)$$

which is required to have a unique fixed point $y^* = S(x^*)$ for every fixed value $x^* \in \mathbb{R}^d$. We also introduce the deviations from equilibrium $Z(t) = S(X(t)) - Y(t)$.

The velocity of each bacterium is switched at random jump times $(T^k)_{k \geq 1}$ that are generated via a Poisson process with a time dependent rate given by $\lambda(Z(t))$, where $z \mapsto \lambda(z)$ is a smooth function satisfying

$$0 < \lambda_{\min} \leq \lambda(z) \leq \lambda_{\max}, \quad (4.60)$$

as well as (for small values of z),

$$\lambda(z) = \lambda^0 - bz + c_\lambda \mathcal{O}(|z|^\gamma), \quad (4.61)$$

with $b \in \mathbb{R}$, $\gamma \geq 2$. As before, the new velocity at time T^k is generated at random according to a centered probability distribution $\mathcal{M}(dv)$ with $\int v \mathcal{M}(dv) = 0$, typically

$$\mathcal{M}(dv) = \sigma_{\mathbb{S}^{d-1}}(dv),$$

where $\sigma_{\mathbb{S}^{d-1}}$ is the uniform distribution on the unit sphere.

The resulting fine-scale stochastic evolution of a bacterium is then described by the following differential velocity-jump equation,

$$\begin{cases} \frac{dX(t)}{dt} = \epsilon V(t), \\ \frac{dY(t)}{dt} = F_\epsilon(Y(t), S(X(t))), \\ \int_{T^k}^{T^{k+1}} \lambda(Z(t)) dt = \theta^{k+1}, \quad \text{with } Z(t) := S(X(t)) - Y(t), \\ V(t) = \mathcal{V}^k \quad \text{for } t \in [T^k, T^{k+1}), \end{cases} \quad (4.62)$$

with initial condition $X(0) \in \mathbb{R}^d$, $Y(0) \in \mathbb{R}$ and $T^0 = 0$. In (4.62), $(\theta^k)_{k \geq 1}$ denote i.i.d. random variables with normalized exponential distribution, and $(\mathcal{V}^k)_{k \geq 1}$ denote i.i.d. random variables with distribution $\mathcal{M}(dv)$.

In the numerical experiments, we will use a specific example, adapted from [15]. For the internal dynamics (4.59), we choose a linear equation

$$\frac{dy}{dt} = \frac{S(x) - y}{\tau} = \frac{z}{\tau}. \quad (4.63)$$

For the turning rate $z \mapsto \lambda(z)$, we choose the following nonlinear strictly decreasing smooth function

$$\lambda(z) = 2\lambda_0 \left(\frac{1}{2} - \frac{1}{\pi} \arctan \left(\frac{\pi}{2\lambda_0} z \right) \right). \quad (4.64)$$

The probability distribution density of the fine-scale process with internal state at time t with respect to the measure $dx \mathcal{M}(dv) dy$ is denoted as $p(x, v, y, t)$, suppressing the dependence on ϵ for notational convenience, and evolves according to the Kolmogorov forward evolution equation (or master equation). In the present context, the latter is the following kinetic equation

$$\partial_t p + \epsilon v \cdot \nabla_x p + \operatorname{div}_v (F_\epsilon(x, y) p) = \lambda(S(x) - y) (R(p) - p), \quad (4.65)$$

where

$$R(p) := \int_{\mathbb{V}} p(\cdot, v, \cdot) \mathcal{M}(dv)$$

is again the operator integrating velocities with respect to \mathcal{M} .

4.6.2 Relation Between Fine-Scale and Coarse Process

In [45], it is shown, using probabilistic arguments, that, in the limit $\epsilon \rightarrow 0$, both the equation for the coarse process (4.55) and the equation for the process with internal state (4.65) converge to an advection-diffusion limit on diffusive time scales. Convergence is to be understood *pathwise*, i.e., in the sense of individual trajectories.

For the coarse process, this result has already been stated, see Sect. 4.4.3.

In the same way, a standard probabilistic diffusion approximation argument can be used to derive the pathwise diffusive limit of the process with internal state (4.62), see [45]. For $\epsilon \rightarrow 0$, the process $\bar{t} \mapsto X^\epsilon(\bar{t})$, solution of (4.62), converges towards an advection-diffusion process, satisfying the stochastic differential equation (SDE) (4.58), where A_0 originates from

$$A_0(x) = b \lim_{\epsilon \rightarrow 0} \frac{\tau}{\lambda^0 \tau + 1} \nabla S(x), \quad (4.66)$$

in which b , τ , and λ^0 were introduced in (4.61)–(4.63) as parameters of the process with internal state. Again, the diffusion matrix D is given by the covariance of the Maxwellian distribution (4.57).

Introducing the bacterial density of the process with internal state as

$$\rho(x, t) = \int_{\mathbb{Y}} \int_{\mathbb{V}} p(x, v, y, t) \mathcal{M}(dv) dy, \quad (4.67)$$

this implies that the evolution of ρ converges to (4.56) on diffusive time scales in the limit of $\epsilon \rightarrow 0$.

4.6.3 Asymptotic Variance Reduction

As discussed in Sect. 4.5, obtaining the position density of bacteria by solving the kinetic equation (4.65) over diffusive time scales can be cumbersome, due to the additional dimensions associated with the internal state. The alternative is to use stochastic particles. However, a particle-based simulation of Eq. (4.65) is subject to a large statistical variance of the order $\mathcal{O}(M^{-1/2})$, where M is the number of simulated particles. Additionally, the asymptotic analysis shows that the position bacterial density approaches an advection-diffusion limit (4.56) when $\epsilon \rightarrow 0$; moreover, this advection-diffusion limit is shared with a simpler, coarse model without internal dynamics. Consequently, to accurately assess the deviations of the process with internal state (4.62) as compared to its advection-diffusion limit (for small but non-vanishing (intermediate) values of ϵ), the required number of particles needs to increase substantially with decreasing ϵ , which may become prohibitive from a computational point of view.

The idea that will be discussed here is to construct a hybrid method, based on the principle of control variates, that couples the process with internal dynamics with the coarse process, which is simulated simultaneously using a grid-based method. The stochastic particles are then only used to perform a Monte Carlo simulation of the deviations of the model with internal state with respect to the coarse model.

To explain the method, let us first assume that we are able to compute the exact solution of the kinetic equation for the coarse process, (4.55), with infinite precision in space and time. From now on, we will refer to (4.55) as the *control process* as it is used as a control variate. (This explains the addition of the subscript c .)

The algorithm of asymptotic variance reduction is based on a coupling between an ensemble of realizations evolving according to the process with internal state (4.62), denoted as

$$\{X_m(t), V_m(t), Y_m(t)\}_{m=1}^M,$$

and an ensemble of realization of the control process (4.54), denoted as

$$\{X_{m,c}(t), V_{m,c}(t)\}_{m=1}^M.$$

We denote the empirical measure of the particles with internal state in position-velocity space as

$$\mu_{\bar{t}}^M(x, v) = \frac{1}{M} \sum_{m=1}^M \delta_{X_m(\bar{t}/\epsilon^2), V_m(\bar{t}/\epsilon^2)},$$

and, correspondingly, the empirical measure of the control particles as

$$\mu_{\bar{t}}^{M,c} = \frac{1}{M} \sum_{m=1}^M \delta_{X_{m,c}(\bar{t}/\epsilon^2), V_{m,c}(\bar{t}/\epsilon^2)},$$

with $\delta_{x,v}$ the Dirac delta centered at (x, v) . (These empirical measures can be interpreted as a discrete particle approximation to the phase-space density of the bacteria.)

A coupling between the two ensembles is obtained by ensuring that both simulations use *the same random numbers* $(\theta^k)_{k \geq 1}$ and $(\mathcal{V}^k)_{k \geq 0}$, which results in a strong correlation between $(X_m(t), V_m(t))$ and $(X_{m,c}(t), V_{m,c}(t))$ for each realization. Simultaneously, the kinetic equation for the control process (4.55) is also solved using a deterministic method (which, for now, is assumed to be exact). We formally denote the corresponding semi-group evolution (a formal description of the exact solution) as

$$e^{\bar{t}L_c}, \quad \text{with } L_c(p_c) = -\epsilon v \cdot \nabla_x p_c + (R(\lambda_c^\epsilon p_c) - \lambda_c^\epsilon p_c).$$

Besides the two particle measures $\mu_{\bar{t}}^M$ and $\mu_{\bar{t}}^{M,c}$, we denote by $\bar{\mu}_{\bar{t}}^M$ the variance reduced measure, which will be defined by the algorithm below. Since, with increasing diffusive time, the variance of the algorithm increases due to a loss of coupling between the particles with internal state and the control particles, the variance reduced algorithm will also make use of a reinitialization time step $\bar{\delta}t_{ri}$, which is defined on the diffusive time scale. The corresponding time instances are denoted as $\bar{t}^\ell = \ell \bar{\delta}t_{ri}$ on the diffusive time scale, or equivalently, on the original time scale as $t^\ell = \ell \bar{\delta}t_{ri} / \epsilon^2$.

Starting from an initial probability measure μ_0 at time $t = 0$, we sample μ_0 to obtain the ensemble $\{X_m(t), V_m(t), Y_m(t)\}_{m=1}^M$, corresponding to μ_0^M , and then set $\mu_0^{M,c} := \mu_0^M$, i.e., $X_{m,c}(0) = X_m(0)$ and $V_{m,c}(0) = V_m(0)$ for all $m = 1, \dots, M$. Furthermore, we set the variance reduced estimator as $\bar{\mu}_0^M := \mu_0 = \mathbb{E}(\mu_0^M)$. We then use the following algorithm to advance from \bar{t}^ℓ to $\bar{t}^{\ell+1}$, (see also Fig. 4.4):

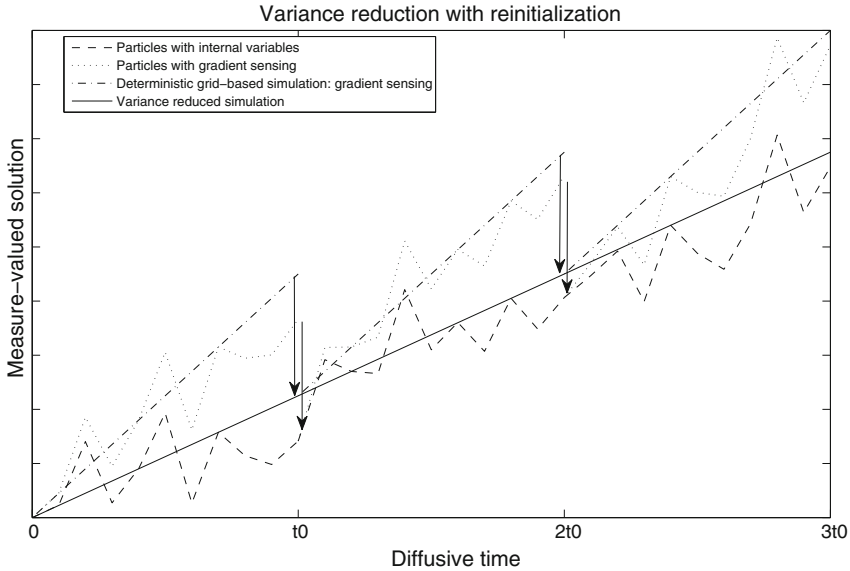


Fig. 4.4 A schematic description of Algorithm 3. The *dashed line* represent the evolution of M bacteria with internal state. The *dotted line* represent the coupled evolution of M bacteria with gradient sensing, subject to regular reinitializations. The *dashed-dotted line* is computed according to a deterministic method simulating the density of the model with gradient sensing, and subject to regular reinitializations. The *solid line* is the variance reduced simulation of the internal state dynamics, and is computed by adding the difference between the particle computation with internal state, and the particle simulation with gradient sensing to the deterministic gradient sensing simulation. At each reinitialization step, the two simulations (deterministic and particles) of the gradient sensing dynamics are reinitialized to the values of their internal state simulation counterpart (as represented by the *arrows*)

Algorithm 3 At time t^ℓ , we have that the particle measure $\mu_{\tilde{t}^\ell}^{M,c} = \mu_{\tilde{t}^\ell}^M$, and the variance reduced measure is given by $\bar{\mu}_{\tilde{t}^\ell}^M$. To advance from time \tilde{t}^ℓ to $\tilde{t}^{\ell+1}$, we perform the following steps :

- Evolve the particles $\{X_m(t), V_m(t), Y_m(t)\}_{m=1}^M$ from t^ℓ to $t^{\ell+1}$, according to (4.62);
- Evolve the particles $\{X_{m,c}(t), V_{m,c}(t)\}_{m=1}^M$ according to (4.54), using the same random numbers as for the process with internal state;
- Compute the variance reduced evolution

$$\bar{\mu}_{\tilde{t}^{\ell+1}}^M = \bar{\mu}_{\tilde{t}^\ell}^M e^{\delta \tilde{t}_{ri} / \epsilon^2 L_c} + \left(\mu_{\tilde{t}^{\ell+1}}^M - \mu_{\tilde{t}^{\ell+1}}^{M,c} \right). \tag{4.68}$$

(Note that this implies that we start the deterministic simulation for this time step from $\bar{\mu}_{\tilde{t}^\ell}^M$.)

- Reinitialize the control particles by setting

$$X_{m,c}(t^{\ell+1}) = X_m(t^{\ell+1}), \quad V_{m,c}(t^{\ell+1}) = V_m(t^{\ell+1}), \quad m = 1, \dots, M,$$

i.e., we set the state of the control particles to be identical to the state of the particles with internal state.

In (4.68), we use the symbol $\bar{t}^{\ell+1}$ to emphasize that the involved particle positions and velocities are those obtained *before* the reinitialization. An easy computation shows that the algorithm is unbiased in the sense that for any $\ell \geq 0$,

$$\mathbb{E} \left[\bar{\mu}_{\bar{t}^\ell}^M \right] = \mathbb{E} \left[\mu_{\bar{t}^\ell}^M \right],$$

since the particles with internal dynamics are unaffected by the reinitialization, and, additionally,

$$\mathbb{E} \left[\mu_{\bar{t}^{\ell+1}}^{M,c} \right] = \mathbb{E} \left[\bar{\mu}_{\bar{t}^\ell}^M e^{\bar{\delta} t_{ri} / \epsilon^2 L_c} \right].$$

Moreover, the variance is controlled by the coupling between the two processes. Indeed, using the independence of the random numbers between two steps of Algorithm 3, and introducing φ as a position and velocity dependent test function, we get (see [44]),

$$\text{stdev}(\bar{\mu}_{\bar{t}^\ell}^M(\varphi)) \leq C \frac{\epsilon}{M}, \quad (4.69)$$

where in the last line, C is independent of ℓ , ϵ , and M .

Remark 10 (Sharpness of the variance estimate) In some generic situations, we can argue that the statistical error in Algorithm 3 coming from the coupling is “sharp” with respect to the order in ϵ . This means that the difference between the probability distribution of the model with internal state and the probability distribution of the model with gradient sensing is of the same order. This would imply that, with the asymptotic variance reduction technique, one is able to reliably assess the true deviation of the process with internal variables from the control process using a number of particles M that is independent of ϵ .

Remark 11 (Effect of time discretization) The analysis in [44] reveals that the variance reduction is *asymptotic*, in the sense that the variance vanishes in the diffusion limit. To ensure this asymptotic variance reduction during actual simulations, one needs to ensure that the time discretization preserves the diffusion limits of the time-continuous process. An appropriate time discretization is highly non-trivial, and is discussed in [44].

4.6.4 Numerical Results

To illustrate the algorithm, we consider a simulation of the density of an ensemble of particles, with and without variance reduction. We restrict ourselves to one space

dimension, with domain $x \in [0, 20]$ and periodic boundary conditions. In this case, the kinetic equation corresponding to the control process reduces to the system

$$\begin{cases} \partial_t p_c^+ + \epsilon \partial_x p_c^+ = -\frac{\lambda_c(x, +1)}{2} p_c^+ + \frac{\lambda_c(x, -1)}{2} p_c^- \\ \partial_t p_c^- - \epsilon \partial_x p_c^- = \frac{\lambda_c(x, +1)}{2} p_c^+ - \frac{\lambda_c(x, -1)}{2} p_c^- \end{cases}. \quad (4.70)$$

of two PDEs, which is straightforward to simulate using finite differences.

We fix the chemoattractant concentration field as

$$S(x) = \alpha \left(\exp \left(-\beta (x - \xi)^2 \right) + \exp \left(-\beta (x - \eta)^2 \right) \right), \quad (4.71)$$

with parameters $\alpha = 2$, $\beta = 1$, $\xi = 7.5$ and $\eta = 12.5$. For the internal dynamics, the model ((4.63–4.64), (4.53–4.66)) is used. The parameters are $\epsilon = 0.5$, $\lambda_0 = 1$, $\tau = 1$, $\delta t = 0.1$.

All simulations are performed with $M = 5000$ particles. The initial positions are uniformly distributed in the interval $x \in [13, 15]$; the initial velocities are chosen uniformly, i.e., each particle has an equal probability of having an initial velocity of $\pm\epsilon$. The initial condition for the internal variable is chosen to be in local equilibrium, i.e., $Y_m(0) = S(X_m(0))$. The initial positions and velocities of the control particles are chosen to be identical.

We discretize the continuum description (4.70) on a mesh with $\Delta x = 0.1$ using a third-order upwind-biased scheme, and perform time integration using the standard fourth order Runge–Kutta method with time step $\delta t_{pde} = 10^{-1}$. The initial position density is given as

$$\rho(x, 0) = \begin{cases} 0.25, & x \in [13, 15], \\ 0, & \text{otherwise.} \end{cases} \quad (4.72)$$

Simulation without variance reduction First, we simulate both stochastic processes up to time $\bar{t} = 50$ ($t = 50/\epsilon^2$) and estimate the density of each of these processes $\hat{\rho}_c^M(x, \bar{t})$, resp. $\hat{\rho}_c^M(x, \bar{t})$, without variance reduction. The density is obtained via binning in a histogram, in which the grid points of the deterministic simulation are the centers of the bins. Figure 4.5 (left) shows the results for a single realization. We see that, given the fluctuations on the obtained density, it is impossible to conclude on differences between the two models. This observation is confirmed by computing the average density of both processes over 100 realizations. The mean densities are shown in Fig. 4.5 (right), which also reveals that the mean density of the control process is within the 95% confidence interval of the process with internal state. Both figures also show the density that is computed using the continuum description, which coincides with the mean of the density of the control particles.

Simulation with variance reduction Next, we compare the variance reduced estimation (4.68) with the density of the control PDE. We reinitialize the control par-




Fig. 4.5 Bacterial density as a function of space at $t = 50/\epsilon^2$ without variance reduction. *Left* one realization. *Right* mean over 100 realizations and 95 % confidence interval. The *solid line* is the estimated density from a particle simulation using the process with internal state; the *dashed line* is estimated from a particle simulation using the control process. Both used $M = 5000$ particles. The *dotted line* is the density obtained from the deterministic PDE (4.70)




Fig. 4.6 Bacterial density as a function of space at $t = 50/\epsilon^2$ with variance reduction and reinitialization. *Left* variance reduced density estimation of one realization with $M = 5000$ particles (*solid*) and density obtained from a deterministic solution for the control process (4.70) (*dashed*). *Right* mean over 100 realization and 95 % confidence interval (*solid*) and density obtained from a deterministic solution for the control process (4.70) (*dashed*)

ticles after each coarse-scale step, i.e., each k steps of the particle scheme, where $k\delta t = \delta t_{pde}$, (here $k = 1$). The results are shown in Fig. 4.6. We see that, using this reinitialization, the difference between the behaviour of the two processes is visually clear from one realization (left figure). Also, the resulting variance is such that the density of the control PDE is no longer within the 95 % confidence interval of the variance reduced density estimation (right figure). We see that there is a significant difference between both models: the density corresponding to the control process is more peaked, indicating that bacteria that follow the control process are more sensitive to sudden changes in chemoattractant gradient. This difference can be interpreted from the fact that the bacteria with internal state do not adjust themselves instantaneously to their environment, but instead with a time constant τ .

4.7 Application 2: Tumor Growth

As a second application, we turn to tumor growth, which is a complex biological phenomenon consisting of processes on different scales. On the cellular level, one has to track the random motion of the cells, as well as cell division and cell death (apoptosis). As before, we deal with random motion by means of a velocity-jump process, in which we additionally ensure that the concentration of cells in a certain volume remains restricted. To achieve this, the spatial evolution of the individuals will be coupled through the local cell density. Cell division and apoptosis are modeled by means of two extra (intracellular) variables (a cell cycle variable ϕ and an apoptosis variable z), which in turn depend on intracellular and environmental concentrations of a number of chemical compounds (described via reaction diffusion equations). The resulting *fine-scale* model therefore consists of a velocity-jump process, supplemented with a set of ODEs describing the (sub)-cellular state of the individual cells and a set of reaction-diffusion PDEs describing the environment.

As for bacterial chemotaxis, it is possible to equivalently write the fine-scale model as a kinetic equation for each cell type p that then models the phase space density $p_p(x, v, \phi, z, t)$. However, a direct simulation of this fine-scale kinetic model is again not feasible because of the high-dimensional character of the resulting system, while a stochastic particle discretization is significantly influenced by Monte Carlo noise. Therefore, we propose a tailored variance reduction technique. The key point is to simulate the kinetic description of a simpler *control model* that only contains the motion of the cells and to couple this deterministic simulation with a stochastic agent-based simulation to obtain information on cell divisions and apoptosis.

We first give a detailed overview of the different layers of the model in Sect. 4.7.1. This model is similar to the model (4.62) used to describe bacterial chemotaxis, and reproduces the features of the cellular automaton model proposed by Owen et al. [39]. We describe the variance reduction algorithm in Sect. 4.7.2. Finally, we illustrate the technique with some numerical experiments in Sect. 4.7.3.

4.7.1 Model

The model consists of two main components: an agent-based model, describing the individual cellular motion and internal processes attached to each cell (cell cycle and apoptosis) and the environment, modeled by a set of reaction diffusion equations. We start by giving an overview of the model structure and notations that will be used throughout the section (Sect. 4.7.1.1), after which we describe the agent-based model (Sect. 4.7.1.2) and the evolution laws for the environment (Sect. 4.7.1.3).

4.7.1.1 Overview and Notation

We consider three types of cells, indexed by $1 \leq p \leq P = 3$: normal cells ($p = 1$), cancer cells ($p = 2$), and endothelial cells (that build up blood vessels, $p = 3$). For

each of these cell types, we consider an ensemble of $M_p(t)$ cells, characterized by their positions $X_{m,p}(t)$, velocities $V_{m,p}(t)$, cell cycle phase $\Phi_{m,p}(t)$, and remaining internal state variables $Z_{m,p}(t)$, for $1 \leq m \leq M_p(t)$ and $1 \leq p \leq P$. (The time dependence of the number of particles is due to the fact that cells divide and die.) In the numerical experiments in this chapter, blood vessel growth is not taken into consideration. We thus only consider normal cells and cancer cells.

To keep a consistent notation throughout the section, we introduce the following convention. If, at a moment $t = t^*$, the cell with index m^* in population p divides, we set

$$M_p(t^*) = M_p(t^*_-) + 1, \quad (4.73)$$

in which the symbol t^*_- is used to emphasize that the involved number of cells is meant to be taken *just before* the division. Simultaneously, we introduce a new cell as specified below (see Eq. 4.79). When a cell undergoes apoptosis, it is removed from the simulation. To avoid cumbersome renumbering of the cells in the text, we associate a weight $w_{m,p}(t)$ to each of the cells. If the cell is alive, the corresponding weight is one; upon apoptosis, it becomes zero. The active number of cells is therefore

$$\bar{M}_p(t) = \sum_{m=1}^{M_p(t)} w_{m,p}(t). \quad (4.74)$$

Given the positions of these cells, the empirical cell number density is then obtained as

$$\rho_p(x, t) = \sum_{m=1}^{M_p(t)} w_{m,p}(t) \delta_{X_{m,p}(t)}. \quad (4.75)$$

The agent-based cellular model is coupled with the environment, consisting of oxygen and vascular endothelial growth factor (VEGF). We denote by $C(x, t)$ the concentration of oxygen and by $G(x, t)$ the concentration of VEGF; both fields evolve according to PDEs in which the cell number density (4.75) appears. The different behaviour for different cell types originates both from different intracellular mechanisms and from the cell-type dependency of the coefficients in the agent-based model, see below. A table containing the parameter values for all cell types is given in Table 4.1.

4.7.1.2 Agent-Based Model

Motion Cellular motion is composed of two components: random motion, which will be modeled by means of a velocity jump process, and deterministic chemotactic motion towards high VEGF concentrations. The deterministic term also contains a volume factor, restricting motion towards regions where the number density $\rho_p(X_{m,p}(t), t)$ is higher than a threshold value $\rho_{\max,p}$

$$\frac{dX_p(t)}{dt} = \epsilon V_p(t) + \epsilon^2 \chi_p \nabla G(X_p(t), t) \left(1 - \frac{\rho_p(X(t), t)}{\rho_{\max, p}} \right), \tag{4.76}$$

$$V_p(t) = \mathcal{V}_p^k \text{ for } t \in [T_p^k, T_p^{k+1}], \quad \mathcal{V}_p^k \in \mathbb{V} = V_p^* \mathbb{S}^{d-1}, \tag{4.77}$$

in which $0 < \epsilon \ll 1$ and χ_p is the cell-type dependent chemotactic sensitivity. In this chapter, we choose the same velocity V_p^* for all cell types. Similar to the bacterial chemotaxis case, the velocities \mathcal{V}_p^k are sampled from the uniform distribution

$$\mathcal{M}_p(dv) = \sigma_{V_p^* \mathbb{S}^{d-1}}(dv),$$

on the sphere $V_p^* \mathbb{S}^{d-1}$, and the jump times T_p^k are generated from a Poisson process with constant rate λ_p ,

$$\int_{T_p^k}^{T_p^{k+1}} \lambda_p dt = \lambda_p (T_p^{k+1} - T_p^k) = \theta_p^{k+1},$$

where θ_p^{k+1} are i.i.d. random numbers, sampled from a normalized exponential distribution. During the numerical experiments, we choose $\lambda_p = 1$, independent of the cell type. In the bacterial chemotaxis case, bacteria possessed internal dynamics

Table 4.1 Parameter values related to the populations

Parameter	Normal population	Cancer population	Endothelial cells
χ	0.0	0.0	0.33333
ρ_{\max}	0.1	0.2	0.1
ϵ	3.536×10^{-4}	3.536×10^{-4}	–
C_ϕ	399.96	186.64	–
T_{\min}	1.8×10^5	9.6×10^4	–
z_{high}	0.8	–	–
z_{low}	0.08	–	–
n_{thr}	0.75	–	–
c_1	3.3333×10^{-5}	–	–
c_2	1.6667×10^{-4}	–	–
c_3	3.3333×10^{-5}	–	–
c_4	3.3333×10^{-5}	–	–
c_5	1.6667×10^{-4}	–	–
J_5	0.04	–	–
C_{VEGF}	0.01	–	–
C_{p53}	0.01	–	–
A	–	1	–
B	–	1	–

to bias their chemotactic motion. In the tumor model, internal dynamics governs cell division and cell death (apoptosis). (Hence, the time-dependent number of cells $M_p(t)$.) Let us now look at these intracellular mechanisms (Table 4.1)

Cell Cycle Dynamics Cells evolve according to a cell cycle, and division occurs as soon as this cycle is completed. In our model, this is represented by a cell phase cycle variable $\Phi_p(t)$ that is zero at cell birth. The cell cycle is completed when $\Phi_p(t) = 1$. The cell cycle speed depends on the local oxygen concentration $C(X_p(t), t)$ that is observed by the cell as it is simultaneously moving through space and evolving through the cycle. The higher the oxygen concentration, the faster the cycle proceeds, while the cell cycle is put on hold once the oxygen concentration is approaching zero. This behaviour is modeled by means of the following ODE,

$$\frac{d\Phi_p(t)}{dt} = \frac{C(X_p(t), t)}{\tau_{\min,p}(C_{\phi,p} + C(X_p(t), t))}, \quad (4.78)$$

in which we introduce the cell-type dependent parameters $C_{\phi,p}$ and $\tau_{\min,p}$, the minimal time needed for a cell to complete one cell cycle. From Table 4.1, we see that cancer cells are able to proceed twice as fast as normal cells during the cell cycle in a given environment.

If, for the cell with index m^* in population p , at time $t = t^*$, we obtain $\Phi_p(t^*) \geq 1$, we introduce a new cell in the simulation. We adjust $M_p(t)$ according to (4.73) and set $\Phi_{m^*,p}(t) = 0$. The new cell inherits the complete state from the cell that divides:

$$\begin{aligned} X_{M_p(t),p}(t^*) &= X_{m^*,p}(t^*), & V_{M_p(t),p}(t^*) &= V_{m^*,p}(t^*), \\ \Phi_{M_p(t),p}(t^*) &= \Phi_{m^*,p}(t^*), & Z_{M_p(t),p}(t^*) &= Z_{m^*,p}(t^*). \end{aligned} \quad (4.79)$$

More details on this cell cycle model can be found in [39] and its supplementary material.

Remaining Internal State Variables To account for apoptosis, we introduce a second sub-cellular model, consistent with [39],

$$\frac{dZ_p(t)}{dt} = F_p(X_p(t), Z_p(t)),$$

for the internal state $Z_p(t) \in \mathbb{R}^q$, in which q may depend on the cell type. This internal dynamics follows a different mechanism, depending on the cell type. For the normal tissue ($p = 1$), the variable $Z_1(t)$ contains two components, namely the p53 concentration $Z_{1,1}(t)$ and the intracellular VEGF concentration $Z_{1,2}(t)$. The former can be seen as an estimator for the number of mutations that a cell has undergone during its lifetime. The latter models the process that allows cells to store VEGF during hypoxic conditions and release it once this intracellular concentration has reached a certain threshold level.

$$\frac{dZ_{1,1}(t)}{dt} = c_1 - c_2 \frac{C(X_1(t), t)}{C_{p53} + C(X_1(t), t)} Z_{1,1}(t), \quad (4.80)$$

$$\frac{dZ_{1,2}(t)}{dt} = c_3 - c_4 \frac{Z_{1,1}(t)Z_{1,2}(t)}{J_5 + Z_{1,2}(t)} + c_5 \frac{C(X_1(t), t)}{C_{VEGF} + C(X_1, t)} Z_{1,2}(t), \quad (4.81)$$

where c_i , $1 \leq i \leq 5$, C_{p53} , J_5 , and C_{VEGF} are parameters of which the value can be found in Table 4.1. As soon as oxygen is available, the second term in (4.80) ensures exponential decay of $Z_{1,1}(t)$; the first term in (4.80) models linear growth of $Z_{1,1}(t)$, an effect that is only dominant in the absence of oxygen. A cell undergoes apoptosis if $Z_{1,1}(t)$ reaches a threshold value $\gamma_{\text{apt}}(\rho_1(X_1(t), t))$ that depends on the local density of normal cells. The threshold value is lower in case of a harsh environment (with low cell number density), i.e.,

$$\gamma_{\text{apt}}(\rho) = \begin{cases} z_{\text{high}} & \text{if } \rho \leq \rho_{\text{thr}} \\ z_{\text{low}} & \text{else} \end{cases}.$$

See Table 4.1 for parameter values.

The internal dynamics of tumor cells does not depend on the p53 concentration, since this mechanism to regulate the normal cell cycle does not function properly anymore in a tumor. Cancer cells are able to go into a quiescent state when expressed to hypoxic circumstances, meaning that they don't consume any nutrients. However, the maximal duration of this quiescent state is limited, which implies that cancer cells will also undergo apoptosis when the hypoxia holds too long. On the other hand, cancer cells have the ability to recover quickly once oxygen becomes available again. This mechanism can be modeled by the following equation:

$$\frac{dZ_2(t)}{dt} = \underbrace{A \text{H}(C_{\text{threshold}} - C(X_2(t), t))}_{\text{Linear increase during hypoxia}} - \underbrace{B Z_2(t) \text{H}(C(X_2(t), t) - C_{\text{threshold}})}_{\text{Exponential decay if } C(X_2(t), t) > C_{\text{threshold}}}, \quad (4.82)$$

where A , B are constants (see Table 4.1) and H is the Heaviside function. The first term models the reaction to a hypoxic state, i.e., when the local oxygen concentration $C(X_2(t), t)$ drops below the threshold level $C_{\text{threshold}}$. During this hypoxic period, the internal variable $Z_2(t)$ increases linearly as a function of time. On the other hand, the second term describes the recovery of the cancer cells if the environment is not hypoxic anymore, which is captured by the exponential decay term of $Z_2(t)$. Cancer cells die as soon as $Z_2(t) \geq 1$, corresponding to $\gamma_{\text{apt}} = 1$.

Complete Agent-Based Model Combining all components described above, we end up with the following set of differential equations governing the behaviour of an individual cell:

$$\left\{ \begin{array}{l} \frac{dX_p(t)}{dt} = \epsilon V_p(t) + \epsilon^2 \chi_p \nabla G(X_p(t), t) \left(1 - \frac{\rho_p(X_p(t), t)}{\rho_{\max}} \right), \\ (T^{k+1} - T^k) \lambda_p = \theta_p^{k+1}, \\ V_p(t) = \mathcal{V}_p^k \quad \text{for } t \in [T_p^k, T_p^{k+1}), \\ \frac{d\Phi_p(t)}{dt} = \frac{C(X_p(t), t)}{\tau_{\min, p}(C_{\phi, p} + C(X_p(t), t))}, \\ \frac{dZ_p(t)}{dt} = F_p(Z_p(t), t), \end{array} \right. \quad (4.83)$$

combined with the division rule (4.79) and the rule that adjusts the weights $w_{m,p}(t)$ upon apoptosis.

Coarse (population-Level) Description As in Sect. 4.4.2, we again have a kinetic equation of the phase space density $p_p(x, v, z, \phi, t)$, which becomes quite complicated due to the cell cycle and apoptosis that govern cell division and cell death. Ideally, a coarse-grained model would be written in terms of the cell number density $\rho_p(x, t)$. We expect to obtain a reaction-advection-diffusion equation. However, due to the modeling detail for the cell cycle and apoptosis, it is unrealistic to expect one can write the reaction term as a closed-form function of $\rho_p(x, t)$. When ignoring the intracellular dynamics (and therefore considering a system with constant number of cells) and using a diffusive scaling $\bar{t} = t\epsilon^2$ [36], one can obtain an advection-diffusion equation where no reactions (cell divisions, cell deaths) are taken into account [30]:

$$\partial_{\bar{t}} \rho_p(x, \bar{t}) = D_p \nabla^2 \rho_p(x, \bar{t}) - \chi_p \nabla \cdot \left[\rho_p(x, \bar{t}) \left(1 - \frac{\rho_p(x, \bar{t})}{\rho_{\max}} \right) \nabla G(x, t) \right], \quad (4.84)$$

in which $D_p = \int_{\mathbb{V}} v \otimes v \mathcal{M}(dv)$.

In the variance reduction algorithm, we will also consider the *kinetic number density*,

$$N_p(x, v, t) = \int \int p_p(x, v, z, \phi, t) dz d\phi = \sum_{m=1}^{M_p(t)} w_{m,p}(t) \delta_{X_{m,p}(t), V_{m,p}(t)}, \quad (4.85)$$

which counts the number of particles with a position x and velocity v at time t , regardless of their internal state.

4.7.1.3 Environment

The cellular environment consists of two diffusible components regulating the behavior of the cells in various ways: oxygen $C(x, t)$ and VEGF concentration $G(x, t)$. Oxygen is evidently important for the cells to proceed through the cell cycle and survive, see Eqs. (4.78) and (4.80). The local oxygen concentration is determined from the following advection-diffusion equation:

Table 4.2 Parameter values reaction diffusion equations

Parameter	Oxygen	VEGF
D	2.4167×10^{-9}	1.6667×10^{-11}
ψ	6	1.6667×10^{-9}
δ	0	$1.667e - 4$
k_{normal}	-0.2167	0.01
k_{cancer}	-0.2167	0.01

$$\partial_t C(x, t) = \underbrace{D_C \nabla^2 C(x, t)}_{\text{diffusion}} + \underbrace{\psi_C \rho_v(x, t) (C_{\text{blood}}(x) - C(x, t))}_{\text{exchange with blood}} - \underbrace{C(x, t) \sum_{p=1}^P k_{C,p} \rho_p(x, t)}_{\text{consumption by cells}}, \quad (4.86)$$

where D_C is the diffusion coefficient, ψ_C denotes the permeability of the oxygen through the vessels, $\rho_v(x, t)$ describes the surface area occupied by blood vessels at position x , and $C_{\text{blood}}(x)$ defines the oxygen concentration in a blood vessel located at position x . In general the blood vessel concentration $\rho_v(x, t)$ is computed from the agent-based evolution of endothelial cells. In this paper, we present numerical tests for a case where $\rho_v(x, t) \equiv \rho_v(x)$ is fixed during the simulation, and no vessel growth is incorporated. We choose $\rho_v(x) = 1$ in grid cells where blood vessels are present. During the experiments we have taken $C_{\text{blood}}(x, t) = 400$ at vessel locations and zero elsewhere. The last term in Eq. (4.86) reflects the fact that all cell types consume oxygen with a cell specific rate $k_{C,p}$. Recall that the cell number density $\rho_p(x, t)$ is defined via Eq. (4.75).

A similar approach is used to describe the local concentration of VEGF, which is responsible for the growth of new blood vessels. This is especially important for larger tumors, since endothelial cells will grow towards regions with higher VEGF concentrations. Initially, the tumor can benefit from the existing vasculature, but when the tumor occupies a larger volume, the oxygen supply does not suffice and the cells are obliged to use their ability to ask for new vessels by secreting VEGF. Endothelial cells – the building blocks of blood vessels– can then react and move chemotactically towards the hypoxic regions. The corresponding reaction diffusion equation for VEGF reads:

$$\partial_t G(x, t) = \underbrace{D_G \nabla^2 G(x, t)}_{\text{diffusion}} - \underbrace{\psi_G \rho_v(x, t) G(x, t)}_{\text{exchange with blood}} + \underbrace{\sum_{p=1}^P k_{G,p} \rho_p(x, t)}_{\text{production}} - \underbrace{\delta_G G(x, t)}_{\text{decay}}. \quad (4.87)$$

In contrast to oxygen, the VEGF concentration is assumed to be zero in the blood and the growth factors contained in the tissue is degrading with rate δ_G when time evolves (Table 4.2).

Because the diffusible components equilibrate much faster than the individual cells, we adopt a steady state approximation, which implies that $\partial_t G(x, t)$ and $\partial_t C(x, t)$ are set to zero. Then, a time step of the agent-based model is performed first, after which the steady state equations for $C(x, t)$ and $G(x, t)$ are solved.

4.7.2 Variance Reduction

In this section, we propose an variance reduction algorithm similar to the technique used for bacterial chemotaxis. The main differences are due to the fact that (i) the tumor model is not conservative; and (ii) the internal dynamics only relates to cell division and apoptosis and not to advection-diffusion behaviour.

Again, the algorithm conceptually relies on the combination of three simulations: a stochastic simulation with the full fine-scale model, as well as with a coarse, approximation, combined with a deterministic, grid-based simulation of the coarse model. The full fine-scale model uses $M_p(t)$ particles with state variables

$$\{X_{m,p}(t), V_{m,p}(t), \Phi_{m,p}(t), Z_{m,p}(t)\}_{m=1}^{M_p(t)}. \quad (4.88)$$

As the coarse agent-based model, we conceptually consider an agent-based model in which the internal state has been suppressed and only position and velocity remain:

$$\{X_{m,p}^c(t), V_{m,p}^c(t)\}_{m=1}^{M_p(t)}.$$

Since no internal dynamics is present, cells cannot divide or die. (In practice, we will use the results obtained from the full fine-scale model, in which we neglect apoptosis and cell division, see later.) The only dynamics is motion, which can be modeled with a kinetic equation for the phase space density $N_p^c(x, v, t)$,

$$\partial_t N_p^c + \epsilon v \cdot \nabla_x N_p^c = \lambda \left(R(N_p^c) - \lambda N_p^c \right), \quad (4.89)$$

see also (4.55). We again call this coarse approximation the *control process*. (Recall that the kinetic number density for the original process is denoted as $N_p(x, v, t)$, see (4.85).) We also introduce the formal semigroup notation

$$e^{tL_p^c}, \quad \text{with } L_p^c(N_p^c) = -(\epsilon v + \epsilon^2 \chi \nabla G(x, t)) \cdot \nabla_x N_p^c + \left(\lambda(R(N_p^c) - N_p^c) \right) \quad (4.90)$$

that represents the exact solution of the kinetic equation (4.89). In practice, this exact solution will be approximated by a deterministic simulation on a grid.

It should be clear that the advection-diffusion behaviour in both agent-based models is identical. Thus, the only difference between the two models occurs when cells divide or die. Assuming no reactions take place, the three processes thus have the

same expectation. This observation leads to the following variance reduction algorithm. As an initial condition, we start from $M_p(0)$ particles sampled from the kinetic probability densities $\mu_{p,0}^{M_p(0)}(x, v)$, resulting in the number density $N_p(x, v, 0)$. For each particle, we choose a given internal state, for instance $\Phi_{m,p}(0) = Z_{m,p}(0) = 0$, $1 \leq m \leq M_p(0)$, $1 \leq p \leq P$. (These internal states could also be sampled from an appropriate probability distribution.) Additionally, we introduce the variance reduced measure $\bar{N}(x, v, t)$, which we initialize as $\bar{N}_p(x, v, 0) = N_p(x, v, 0)$. We denote the time step δt and the discrete time instances $t^\ell = \ell \delta t$, $\ell = 0, 1, \dots$

Algorithm 4 (*Variance reduction for tumor growth*) We advance the variance reduced kinetic number density $\bar{N}(x, v, t)$ from time t^ℓ to $t^{\ell+1}$ as follows:

- Evolve the particle states (4.88) from t^ℓ to $t^{\ell+1}$ using the agent-based model (4.83).
- Compute the kinetic number density for the stochastic fine-scale model using (4.85), as well as the kinetic number density for the coarse process as

$$N_p^c(x, v, t^{\ell+1}) = \sum_{m=1}^{M_p(t^\ell)} w_{m,p}(t^\ell) \delta_{X_{m,p}(t^{\ell+1}), V_{m,p}(t^{\ell+1})}, \quad (4.91)$$

i.e., we compute the kinetic number density for the control process based on particle positions and velocities at time $t^{\ell+1}$, taking into account only the particles that were present in the simulation at time t^ℓ .

- Evolve the control kinetic number density $N_p^c(x, v, t)$ using a grid-based method based on (4.90) and add the reactions (the difference in kinetic number density due to cell division and apoptosis)

$$\bar{N}_p(x, v, t^{\ell+1}) := \bar{N}_p(x, v, t^\ell) e^{\delta t / \epsilon^2 L^c} + N_p(x, v, t^{\ell+1}) - N_p^c(x, v, t^{\ell+1}), \quad (4.92)$$

Again, in the absence of discretization errors in the grid-based method, we have

$$\mathbb{E} \left[\bar{N}_p(x, v, t^{\ell+1}) \right] = \mathbb{E} \left[N_p(x, v, t^{\ell+1}) \right],$$

since $\bar{N}_p(x, v, t^\ell) = N_p^c(x, v, t^\ell)$ and $\mathbb{E} \left[N_p^c(x, v, t^\ell) e^{\delta t / \epsilon^2 L^c} \right] = \mathbb{E} \left[N_p^c(x, v, t^{\ell+1}) \right]$. Moreover, we expect the variance of \bar{N}_p to be significantly lower than that of N_p , since all randomness due to random motion has been removed and only the location of the reactions remains random.

4.7.3 Numerical Experiments

In this section, we illustrate both the model and the variance reduction algorithm by means of some numerical experiments. First, we initialize a normal tissue consisting




Fig. 4.7 Evolution of normal cell population (*black*) with small tumor (*gray*). *Left* initial state (1600 normal cells and 20 cancer cells). *Right* result after 490 steps

of 1600 cells, with initial positions sampled from the uniform distribution on the domain $[0.3, 0.6] \times [0.3, 0.6]$. The initial velocities are sampled from the uniform distribution on the sphere, and the internal state variables $\Phi(0) = 0.0$, $Z(0) = 0$ for every particle. We initialize a small tumor, containing 20 cells, with initial positions samples from a uniform distribution in $[0.445, 0.455] \times [0.445, 0.455]$, also with initial velocities sampled from the uniform distribution on the sphere with radius $V_p^* = 3.5 \times 10^{-6}$, and the internal state variables $\Phi(0) = Z(0) = 0$ for every particle. In the following numerical experiments, we adopt a static vasculature, consisting of two straight vertical vessels at $x = 0.4$ and at $x = 0.8$ to be specific. The first agent-based experiment was performed in a non-scaled way with discretization parameters $\delta t = 1.8 \times 10^3$ s and $\Delta x = 4 \times 10^{-5}$ m for the agent-based model. We use reflective boundary conditions. In Fig. 4.7, we show the evolution of both cell populations. On the left, one can observe the initial configuration, while on the right hand side, we see the resulting configuration after 490 steps. Apart from the fact that cells performed a random walk, a significant amount of the normal tissue died because of the presence of the tumor.

In a second experiment, we illustrate the performance of the variance reduction algorithm as it was explained in Sect. 4.7.2 in a similar setting as the previous experiment but on a diffusive scale, i.e. the normal tissue containing 2500 cells is uniformly distributed on the square $[0.2, 0.8] \times [0.2, 0.8]$ and a small tumor originally consisting of 20 cells is also uniformly distributed within the area $[0.39, 0.41] \times [0.39, 0.41]$. Furthermore, the scaled cellular velocity was chosen $\bar{V}^* = \sqrt{2}/2$ and we modified the minimal cell cycle durations ($T_{\min, \text{cancer}} = 9.6 \times 10^3$ s, $T_{\min, \text{normal}} = 1.8 \times 10^4$ s) to demonstrate the performance of the algorithm in an extreme (not biologically realistic) setting where cells are able to divide very quickly. As in the bacterial chemotaxis application, the coarse equation is modeled on a diffusive timescale, where we choose $\delta \bar{t} = 0.8\epsilon^2 = 1 \times 10^{-7}$, $\Delta \bar{x} = 2 \times 10^{-2}$ as mesh parameters to simulate the deterministic kinetic equation (4.89) needed to apply the variance reduction algorithm. To simulate this, a second order central finite volume scheme was adopted to discretize




Fig. 4.8 Illustration variance reduction algorithm for tumor growth: Mean cell distribution (*first column*) and variance (*2nd column*) for normal and cancer population after 5000 timesteps. The *first two rows* show the results without applying the variance reduction algorithm, which can be compared with the results displayed in the *last two rows* where the variance reduction algorithm was applied

the spatial derivative in the kinetic equation (4.89) and a first order forward Euler scheme for the time derivative.

The mean normal and cancer cell distribution and there variance with and without applying the variance reduction algorithm are displayed. Furthermore the mean oxygen distribution and the corresponding mean reaction field are also shown next to the variance. First, we observe that variance on the normal cell distributed has been reduced significantly for both cell populations by applying Algorithm 4. Additionally, one can observe a clear correlation with the oxygen field. The latter is a logical consequence from the fact that the progress of the cell cycle is closely related to the local oxygen concentration along the track, meaning that a cell is more likely to

divide in oxygen rich environments. Moreover, the variance is almost eliminated in the other regions (Fig. 4.8).

In this section, the parameters used in the numerical experiments are listed. They are based on the parameters used in [39]. One can find them in the supplementary material corresponding to [39].

4.8 Conclusions

We gave a broad overview of the use of stochastic modeling and simulation techniques in computational biology, focussing on some common individual-based modeling and simulation methods. We paid particular attention to the equivalence between the stochastic process that governs the evolution of individual agents and the deterministic behaviour of the involved probability distributions, and we discussed numerical methods that exploit this relation for variance reduction purposes. Using examples involving intracellular chemical reactions, bacterial chemotaxis and tumor growth, we showed the effects of stochasticity at different scales and different levels of description.

A main focus of the chapter was the design of dedicated simulation algorithms. Two main computational bottlenecks arise. The first is related to the time-scale separation between the fast processes (that determine the maximal time step that is allowed) and slow processes (that determine the time scale over which the simulation needs to be performed). The second is related to noise that appears in the simulation and that requires dedicated variance reduction techniques. These problems are not completely solved, and stochastic simulation therefore remains an active research topic.

Acknowledgments This work was partially supported by the Research Council of the K.U. Leuven through grant OT/13/66 and by the Interuniversity Attraction Poles Programme of the Belgian Science Policy Office through grant IUAP/V/22 and by the agency for Innovation by Science and Technology in Flanders (IWT). We thank Bert Mortier, Delphine Vandamme and the editors for their careful proofreading and comments that helped us improve the text.

References

1. Al-Mohssen, H.A., Hadjiconstantinou, N.G.: Low-variance direct Monte Carlo simulations using importance weights. *ESAIM. Math. Model. Numer. Anal.* **44**(05), 1069–1083 (2010)
2. Alt, W.: Biased random walk models for chemotaxis and related diffusion approximations. *J. Math. Biol.* **9**(2), 147–177 (1980)
3. Auger, A., Chatelain, Ph. Koumoutsakos, P.: R-leaping: accelerating the stochastic simulation algorithm by reaction leaps. *J. Chem. Phys.* **125**(8), 084103 (2006)
4. Bal, G., Davis, A.B., Langmore, I.: A hybrid (Monte Carlo/deterministic) approach for multi-dimensional radiation transport. *J. Comput. Phys.* **230**(20), 7723–7735 (2011)
5. Berg, H.C.: *Random Walks in Biology*. Princeton University Press, Princeton (1993)

6. Bren, A., Eisenbach, M.: How signals are heard during bacterial chemotaxis: protein-protein interactions in sensory signal propagation. *J. Bacteriol.* **182**(24), 6865–6873 (2000)
7. Caffisch, R.E.: Monte Carlo and quasi-Monte Carlo methods. *Acta Numer.* **7**, 1–49 (1998)
8. Cao, Y., Gillespie, D.T., Petzold, L.R.: The slow-scale stochastic simulation algorithm. *J. Chem. Phys.* **122**(1), 014116 (2004)
9. Cao, Y., Gillespie, D.T., Petzold, L.R.: Efficient step size selection for the tau-leaping simulation method. *J. Chem. Phys.* **124**(4), 044–109 (2006)
10. Chalub, F.A.C.C., Markowich, P.A., Perthame, B., Schmeiser, C.: Kinetic models for chemotaxis and their drift-diffusion limits. *Monatsh. für Math.* **142**(1–2), 123–141 (2004)
11. Chorin, A.J., Hald, O.H.: *Stochastic Tools in Mathematics and Science*. Springer, Texts in applied mathematics (2013)
12. Chung, C.A., Lin, T.-H., Chen, S.-D., Huang, H.-I.: Hybrid cellular automaton modeling of nutrient modulated cell growth in tissue engineering constructs. *J. Theor. Biol.* **262**(2), 78–267 (2010)
13. Degond, P., Dimarco, G., Pareschi, L.: The moment-guided Monte Carlo method. *Int. J. Numer. Meth. Fluids* **67**(2), 1–24 (2009)
14. Dimarco, G., Pareschi, L.: Hybrid multiscale methods II. kinetic equations. *Multiscale Model. Simul.* **6**(4), 1169–1197 (2008)
15. Erban, R., Othmer, H.G.: From individual to collective behavior in bacterial chemotaxis. *SIAM J. Appl. Math.* **65**(2), 361–391 (2004)
16. Erban, R., Othmer, H.G.: From signal transduction to spatial pattern formation in *E. coli*: a paradigm for multiscale modeling in biology. *Multiscale Model. Simul.* **3**(2), 362–394 (2005)
17. Erban, R., Chapman, J., Maini, Ph., A practical guide to stochastic simulations of reaction-diffusion processes. Technical report, 2007
18. Ethier, S.N., Kurtz, T.G.: *Markov processes: characterization and convergence*. Wiley, New York (2009)
19. Gardiner, C.: *Stochastic methods: a handbook for the natural and social sciences*, vol. 13. Springer Series in Synergetics. Springer, Berlin (2009)
20. Gardner, T.S., Cantor, C.R., Collins, J.J.: Construction of a genetic toggle switch in *Escherichia coli*. *Nature* **403**(6767), 339–342 (2000)
21. Gerlee, P., Anderson, A.R.A.: An evolutionary hybrid cellular automaton model of solid tumour growth. *J. Theor. Biol.* **246**(4), 583–603 (2007)
22. Ghanem, R.G., Spanos, P.D.: *Stochastic Finite Elements: A Spectral Approach*. Springer-Verlag New York Inc, New York (1991)
23. Gillespie, D.T.: Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.* **81**(25), 2340–2361 (1977)
24. Gillespie, D.T.: The chemical langevin equation. *J. Chem. Phys.* **113**(1), 297–306 (2000)
25. Gillespie, D.T.: Approximate accelerated stochastic simulation of chemically reacting systems. *J. Chem. Phys.* **115**(4), 1716 (2001)
26. Gillespie, D.T.: Stochastic simulation of chemical kinetics. *Annu. Rev. Phys. Chem.* **58**, 35–55 (2007)
27. Hairer, E., Wanner, G.: *Solving Ordinary Differential Equations II*. Springer Series in Computational Mathematics, Springer, Berlin (1996)
28. Hasty, J., McMillen, D., Collins, J.J.: Engineered gene circuits. *Nature* **420**(6912), 224–30 (2002)
29. Higham, D.J.: An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Rev.* **43**(3), 525–546 (2001)
30. Hillen, T., Painter, K.J.: A user’s guide to PDE models for chemotaxis. *J. Math. Biol.* **58**(1–2), 183–217 (2009)
31. Kiskowski, M.A., Alber, M.S., Thomas, G.L., Glazier, J.A., Bronstein, N.B., Pu, J., Newman, S.A.: Interplay between activator-inhibitor coupling and cell-matrix adhesion in a cellular automaton model for chondrogenic patterning. *Dev. Biol.* **271**(2), 87–372 (2004)
32. Kloeden, P.E., Platen, E.: *Numerical Solution of Stochastic Differential Equations*. Springer, Stochastic modeling and applied probability (1992)

33. Mallet, D.G., De Pillis, L.G.: A cellular automata model of tumor-immune system interactions. *J. Theor. Biol.* **239**(3), 50–334 (2006)
34. Øksendal, B.: *Stochastic Differential Equations*. Springer (2003)
35. Othmer, H.G., Hillen, T.: The diffusion limit of transport equations derived from velocity-jump processes. *SIAM J. Appl. Math.* **61**(3), 751 (2000)
36. Othmer, H., Hillen, T.: The diffusion limit of transport equations derived from velocity-jump processes. *SIAM J. Appl. Math.* **61**(3), 751–775 (2000)
37. Othmer, H.G., Hillen, T.: The diffusion limit of transport equations II: chemotaxis equations. *SIAM J. Appl. Math.* **62**(4), 1222 (2002)
38. Othmer, H.G., Dunbar, S.R., Alt, W.: Models of dispersal in biological systems. *J. Math. Biol.* **26**(3), 263–298 (1988)
39. Owen, M.R., Stamper, I.J., Muthana, M., Richardson, G.W., Dobson, J., Lewis, C.E., Byrne, H.M.: Mathematical modeling predicts synergistic antitumor effects of combining a macrophage-based, hypoxia-targeted gene therapy with chemotherapy. *Cancer Res.* **71**(8), 37–2826 (2011)
40. Pareschi, L., Toscani, G.: *Interacting Multiagent Systems: Kinetic Equations and Monte Carlo Methods*. Oxford University Press (2013)
41. Patlak, C.S.: Random walk with persistence and external bias. *Bull. Math. Biophys.* **15**(3), 311–338 (1953)
42. Radtke, G., Hadjiconstantinou, N.: Variance-reduced particle simulation of the Boltzmann transport equation in the relaxation-time approximation. *Phys. Rev. E* **79**(5), 1–9 (2009)
43. Roose, T., Chapman, S.J., Maini, Ph K: Mathematical models of avascular tumor growth. *SIAM Rev.* **49**(2), 179–208 (2007)
44. Rousset, M., Samaey, G.: Simulating individual-based models of bacterial chemotaxis with asymptotic variance reduction. *Math. Models Methods Appl. Sci.* **23**(12):2155–2191 (2013)
45. Rousset, M., Samaey, G.: Individual-based models for bacterial chemotaxis in the diffusion asymptotics. *Math. Models Methods Appl. Sci.* **23**(11), 2005–2037 (2013)
46. Schlögl, F.: Chemical reaction models for non-equilibrium phase transitions. *Zeitschrift für Physik* **253**(2), 147–161 (1972)
47. Scott, D.W.: *Multivariate density estimation: theory, practice, and visualization*. Wiley Series in Probability and Statistics. Wiley, New York (2009)
48. Silverman, B.W.: *Density Estimation for Statistics and Data Analysis*. Chapman and Hall/CRC, Boca Raton (1986)
49. Spohn, H.: *Large scale dynamics of interacting particles*, Springer, vol. 825 (1991)
50. Stock, A.M.: A nonlinear stimulus-response relation in bacterial chemotaxis. *Proc. Natl. Acad. Sci.* **96**(20), 10945–10947 (1999)
51. Stuart, A.M.: Inverse problems: a Bayesian perspective. *Acta Numer.* **19**, 451–559 (2010)
52. Weinan, E., Vanden-Eijnden, E.: Transition-path theory and path-finding algorithms for the study of rare events. *Annu. Rev. Phys. Chem.* **61**:391–420 (2010)

Part III
Model Selection and Parameter Fitting

Chapter 5

The Experimental Side of Parameter Estimation

Monica Schliemann-Bullinger, Dirk Fey, Thierry Bastogne,
Rolf Findeisen, Peter Scheurich and Eric Bullinger

Abstract The development of mathematical models is an iterative process. One of its key steps is the acquisition of experimental data. For developing quantitative and dynamical models, this data needs to contain time-courses as well as several perturbation experiments. This article gives an overview of and an introduction into various experimental techniques frequently used during the process of developing dynamical models of biological signalling networks. Among others, the article discusses western blotting, enzyme assays, flow cytometry, protein mass spectrometry, DNA microarrays and different fluorescent microscopy techniques and stresses the significance of single-cell versus population measurements.

M. Schliemann-Bullinger · R. Findeisen · E. Bullinger (✉)
Laboratory for Systems Theory and Automatic Control, University of Magdeburg,
Magdeburg, Germany
e-mail: eric.bullinger@ovgu.de

R. Findeisen
e-mail: rolf.findeisen@ovgu.de

M. Schliemann-Bullinger
e-mail: monica.schliemann-bullinger@ovgu.de

D. Fey
Systems Biology Ireland, University College Dublin, Belfield, Dublin 4, Ireland
e-mail: dirk.fey@ucd.ie

T. Bastogne
Centre de Recherche en Automatique de Nancy (CRAN), Université de Lorraine, CNRS UMR
7039, INRIA BIGS, Vandœuvre-lès-Nancy, France
e-mail: thierry.bastogne@univ-lorraine.fr

P. Scheurich
Institute of Cell Biology and Immunology, University of Stuttgart, Stuttgart, Germany
e-mail: Peter.Scheurich@izi.uni-stuttgart.de

Keywords Cell counting · Computer tomography · DNA microarrays · Electrophoretic mobility shift assays · Enzyme assay · Flow cytometry · Förster resonance energy transfer · Gel electrophoresis · Heterogeneous populations · Live cell imaging · Metabolic flux analysis · Microinjection · Microscopy · Photobleaching · Practical identifiability · Protein mass spectrometry · Quantification · Synthetic biology · Systems biology · Western blotting

5.1 Overview

Estimating parameters of a dynamical system not only requires suitable mathematical methods. Also needed is experimental data. This is in particular the case for parametrising quantitative, dynamical models. While some parameters can be observed in a steady-state, for example the steady-state protein concentrations, other parameters are only visible during transient perturbations. In the model of TNF-induced pro- and anti-apoptotic signalling [45] illustrated in (Fig. 5.1) whose states are concentrations of proteins and mRNA's, roughly half of the flows are zero in the absence of TNF, as of the 106 reactions, only 52 are non-zero: 30 turnover reactions and 22 reactions linked to the shuttling between cytoplasm and nucleus of the transcription factor nuclear factor kappa B (NF- κ B).

Obtaining quantitative, dynamical mathematical models and calibrating their parameters thus requires high quality, high resolution data using both time-courses and multiple perturbations/stimulation. This is particularly challenging in heterogeneous cell populations, see also Chap. 4 by Lejon and Samaey.

Throughout this book chapter, two papers will serve as core examples illustrating the experimental aspects of parameter estimation of dynamical models, as both combine experiments with modelling. The first one is the already mentioned TNF-induced pro- and anti-apoptotic signalling [45]. The other is [40], which studies how the pseudo-phosphatase STYX regulates ERK signalling and affects cell migration and differentiation in HeLa and PC12 cells, respectively. All three biological questions are experimentally studied in mammalian cells. The TNF signalling is studied in KYM-1 cells, a human rhabdomyosarcoma derived cell line. The well defined human carcinoma cell line HeLa is used for elucidating how spatio-temporally ERK1 and ERK2 signalling regulates cell migration, while PC12 cells, a catecholaminergic cell line from the rat, serves as biological model for elucidating cell differentiation. Additionally, several other papers will illustrate specific experimental techniques.

In the following, we will give an introductory overview into several experimental techniques to illustrate the challenges of parameter estimation of biological signalling models, see also Chap. 11 by Mannekee et al. In particular, Sect. 5.2 presents assays for quantifying proteins or enzyme activity and cell counting, Sect. 5.4 discusses various microscopy techniques yielding quantitative and time series data. High-throughput techniques are the topic of Sect. 5.3 (mass spectroscopy), and Sect. 5.5 (genomics) as well as Chap. 12 by Eriksson et al. The chapter concludes with a discussion section.

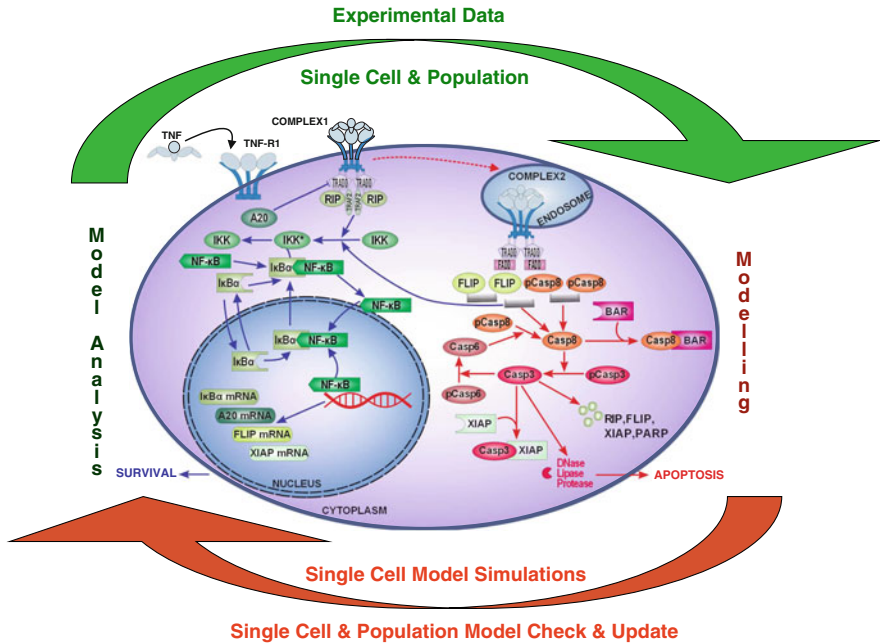


Fig. 5.1 Sketch of TNF-induced pro- and anti-apoptotic signalling [45] consisting of 43 proteins and protein complexes and 4 mRNA. There are 70 irreversible and 18 reversible reactions, totalling 106 kinetic parameters to be estimated. The process of modelling: Acquisition of experimental data, model development, model simulation and model analysis is depicted around the sketch of the signalling pathway. Figure reprinted from [44]

5.2 Assays

Assays are experimental techniques measuring quantitatively the amount or the functional activity of an analyte. For modelling cell signalling, two assays are common: enzyme assays and cell counting.

5.2.1 Gel Electrophoresis

Polyacrylamide gel based methods are a classical and popular way of resolving proteins and other macromolecules such as mRNA. Neglecting the technicalities, cellular extracts are placed onto a gel where the individual types of proteins' travelling speed depends on their mass or their charge. The identification of specific proteins is then achieved via antibodies.




Fig. 5.2 Western blot experiment. Immunodetection of protein bands was performed on an Odyssey infrared imaging system 2.1.12. Markers of known molecular weight are in the outer lanes with antibodies in red, while the proteins of interest are tagged by a green antibody. **a** Rabbit polyclonal antibody directed against I κ B α (41 kDA) and as secondary IRDye-conjugated antibody IRDye 800CW goat anti-rabbit IgG. **b** The protein actin serves as load control. It is detected on the same membrane as **a**, which was stripped of its antibodies before polyclonal goat antibody for actin (43 kDA) and, as secondary IRDye conjugated antibody IRDye 800CW donkey anti-goat. **a** and **b** Combining the binding specificities of the antibody and the molecular weight of the protein of interest allows for localising its band. **a** I κ B α is the *light grey lane* earmarked by the *black and white arrows*, see also Fig. 5.4. **b** Actin is the *dark grey lane*. Reprinted from [44]

5.2.2 Western Blotting

After gel electrophoresis for separating proteins, most often according to their molecular weight, the resolving proteins are transferred by semi-dry electro blotting onto a nitrocellulose membrane using a prestained protein marker with known broad range of the molecular weight. Figure 5.2 shows exemplarily such an experiment where prestained protein markers with the broad range (6–175 kDA) were used. To detect a particular protein, the western blot membrane is incubated with a specific antibody. This antibody can also consist of two antibodies when first a primary antibody for the protein of interest is incubated with the membrane before a secondary antibody recognising this tag is added. The (secondary) antibody can be detected, via their fluorophore or some other source of light or radiation.

Each column in Fig. 5.2 corresponds to a different probe of equal amounts of protein, while the rows corresponds to proteins of different molecular weight as heavier proteins do not travel as far. The lane of the protein of interest can be found by comparison with marker proteins of similar molecular weight. Membranes can be stripped from antibodies and then reprobed with other antibodies of interest, thus allowing several quantifications on a single membrane, thus under identical experimental conditions.

There is a large variety of (commercially) available antibodies. These can be so specific, as to discriminate between phosphorylated and unphosphorylated protein, see the example in Fig. 5.3.

As several samples can be placed side-by-side onto a gel, thus a single gel experiment can quantify molecules from cells in different experimental conditions or harvested at different time points. For example, Fig. 5.4 shows a western blot experiment, also called immunoblotting, of I κ B α , P-I κ B α and NF- κ B time samples after




Fig. 5.3 Time series experiment comparing the response of ERK in starved cells to serum under two conditions: knockdown of STYX via siRNA (*right*) and the control (*left*) [40]. **a** Western blot, where ERK and its activated form, the double phosphorylated ppERK as well as tubulin as control are shown. In the gel, the three bands could be localised using three specific antibodies. **b** Densitometric quantification of three experiments, in which band and background intensity were obtained and subtracted. See [40] for the details. Figure reprinted from [40]




Fig. 5.4 Illustration of time series measurements using western blot. The proteins IκBα, P-IκBα and NF-κB as well as actin and lamin A serving as loading control have been sampled every few minutes following a stimulation with TNF. **a** Western blot data, **b** Quantification based on the blot intensities using the Odyssey software, LI-COR Biosciences (*squares*, connected by a *dashed line*) and the simulation of the mathematical model [45] (*solid line*). **a** and **b** based on [44]

stimulation with TNF (Fig. 5.2). In the first two cases, the protein actin is used as loading control as it is not affected by the TNF stimulation. For the western blot detecting NF-κB in the nucleus, lamin A is the control protein of choice.

5.2.2.1 Quantification

Even though western blotting is commonly used, its quantification is challenging and often not properly performed [19]. The quantitative results of western blotting dependent on antibody (binding) properties, exposure time, gel quality among others, and requires the mapping of each band to some number, the so-called densitometry. This can be done by scanning the gel directly using a CCD camera or via the intermediate step of an X-ray film. Finally, the scanned bands are quantified using image analysis, see e.g. [29], where the integrated intensity of each band corresponds to the signal value. This has for example been applied to obtain the numerical values shown in Fig. 5.4b from the experimental data Fig. 5.4a. This quantification has multiple pitfalls. In general, CCD cameras are preferable as they have a larger dynamic range than films [55]. However, the gel itself can already be saturated due to a too large amount of proteins [55]. Additionally, the values need to be normalised using a common reference, although that introduces noise and opinions differ on how to do it [17], and the result depends on properties of the antibody that may not be fully known [2].

Due to the high variability induced by the experimental steps, it is very difficult to obtain really quantitative results, in particular between individual runs. Usually only signals on the same blot can be compared. However, it is possible to run several gels (at the same time), then cut the band of interest and transfer them in one go onto the same membrane. This technique is called multi-strip western blotting [3].

5.2.3 Electrophoretic Mobility Shift Assays

Electrophoretic mobility shift assays quantify how much of a certain (mixture of) protein binds to a chosen DNA or RNA. As the protein-DNA (or protein-RNA) complex is heavier, it travels at lower speed on the gel, thus leading to a shift of the band compared to a control experiment. Due to their high binding specificity, radioactive markers are usually used for the labelling of specific oligonucleotides with one protein binding site. Due to environmental as well as health and safety concerns, these experiments are usually performed in relatively low number. Figure 5.5 shows blots proportional to the amount of NF- κ B in the nucleus, under different stimulations and at various time points.

5.2.4 Enzyme Assay

Enzyme assays measure the enzymatic activity by quantifying the consumption of an enzyme's substrate. This consumption can be measured continuously over time, or only at temporal samples. If the enzyme amount is constant small compared to the substrate amount, the substrate will be consumed linearly over time. Figure 5.6 shows




Fig. 5.5 Electrophoretic mobility shift assays of nuclear NF-κB after 1 and 10 ng ml⁻¹ TNF for 30 min pulse or continuous stimulation, at several time points (*top row*), with Oct-1 as loading control. Figure reprinted from [45]

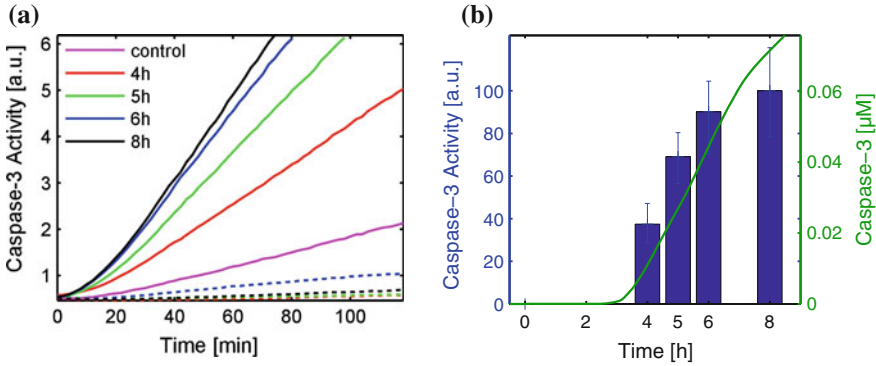


Fig. 5.6 **a** Caspase-3 activity assay for continuous stimulation of 1 ng ml⁻¹ TNF (*solid lines*) measured 4–8 h post stimulus. Control experiments without TNF (*magenta*) and with the caspase inhibitor Z-VAD-fmk (*dashed lines*). **b** Quantification of the slopes of the caspase activity assay in A, where the bars show the spread of second and third quartile. In *green* the corresponding simulation of the model [45]. Figure reprinted from [45]

continuous measurements of the activity of caspase-3 for samples taken at different time points post stimulation. From these measurements, the actual enzymatic activity needs to be inferred and corresponds to the slope of the measured curves.

As can be seen from the experimentally obtained curves in Fig. 5.6a, several measurement points are needed to be certain to ignore transient effects as visible here in the first 15 min, or saturations not visible here due to exhaustion of the substrate or to more complex effects. Figure 5.6b shows the statistics of the slopes at each time point, where the average background activity, measured via the control, has been subtracted.

5.2.5 Cell Counting

The number of cells is a useful marker when studying cell proliferation or viability. The distinction between living and dead cells can be made via dyes that are absorbed

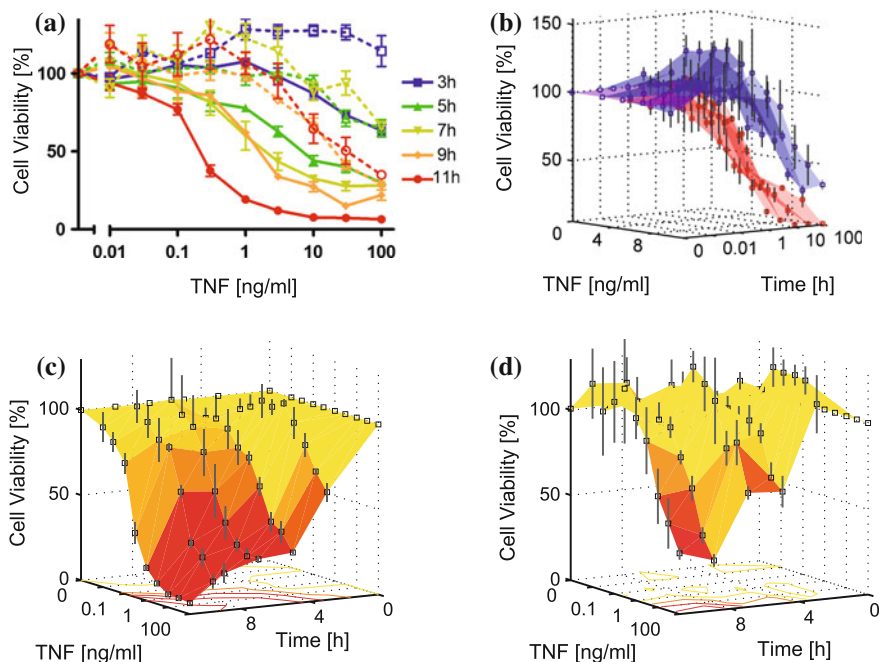



Fig. 5.7 Experimentally obtained dynamic cell viability dose-responses. **a–d** Cell viability over time for 10 different TNF concentrations ($0\text{--}100\text{ ng ml}^{-1}$ TNF) and $0\text{--}11$ h post stimulus, including error bars from three independent experiments. **a** Continuous stimulations (*solid lines*), 30 min pulse stimulations (*dashed lines*). **b–d** Three-dimensional view of the data of **a** with surfaces connecting the measured data points. **b** Continuous (*red*) and 30 min pulse (*blue*) stimulations. **c** Continuous stimulation only, **d** 30 min pulse stimulation only. **c–d** The isoclines (lines connecting points of same cell viability) are shown projected onto the 0% planes. Figures **a**, **c** and **d** reprinted from [45], **b** from [44]

only by dead cells, by MTT assay where living cells convert a substrate into a dye or other similar techniques. The amount of dye can then be quantified as a measure for the number of cells. Figure 5.7 shows the quantification of cytotoxicity assays for different doses of TNF, sampled at various time points.

For adherent cells, the distinction can also be made by measuring their electric impedance, which is different for living and dead cells. All the impedance-based cell analysers use the fact that cells have an electrical insulating membrane. Cells are placed in plate wells equipped with electrodes on the bottom. Since cells are insulators, the more they fix on electrodes the more they decrease the current exchange area, thus increasing the well impedance. The measured impedance signal is affected by the number of cells, the extent to which cells adhere to plate and the cell morphology. The raw impedance is finally converted to a dimensionless cell index variable to correct for well variations. Figure 5.8 shows cell index profiles of MDA-MB-231 cells exposed to various concentrations of nanoparticles.

Fig. 5.8 Cell profiles of MDA-MB-231 cells exposed to various concentrations of nanoparticles. *Blue curves* show the cell proliferation in the absence of nanomaterials while the *purple and red plots* (high concentration of nanoparticles) exhibit their time-resolved cytotoxic effects. Figure reprinted from [7]



For proliferating cells, care needs to be taken if the toxic effect occurs on a similar time scale to the time scale of the cell cycle. Then, the measurement is a combination of cell proliferation and death.

5.2.6 Flow Cytometry

In flow cytometry, cells are passed one by one at high speed in front of an optical sensor, at rates of up to several ten thousand cells per second. This technique can be used to measure cell size, DNA content (useful for determining the cell cycle




Fig. 5.9 Flow cytometry measurement of TNF receptor 1 in KYM-1 cells with approximation (*dashed line*). The *horizontal axis* is a relative intensity of the fluorescent signal in arbitrary unit. The *vertical axis* is the number of events per bin of fluorescence. This histogram has been obtained measuring 10,000 cells. Figure reprinted from [45, supplemental information]

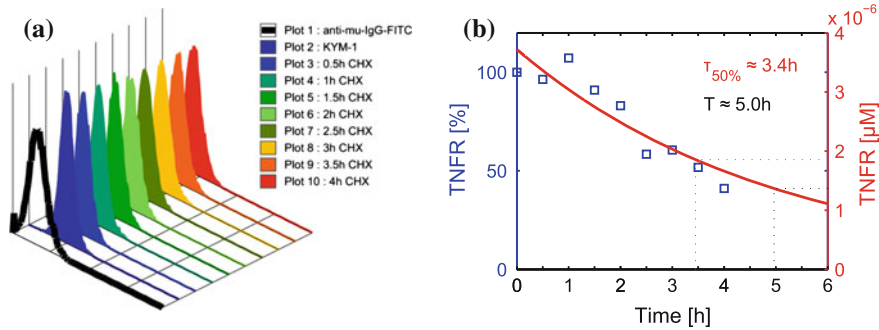


Fig. 5.10 Flow cytometry measurements counting TNF receptor 1 on the cell membrane, at different time points after blocking protein production with CHX. **a** Histograms, one colour per time point, each one similar to the histogram in (Fig. 5.9). **b** Quantification of the experimental data in **a**, normalised by the initial condition (blue squares) and approximated exponential decay (red). Figure reprinted from [45, supplemental information]

phase), and the amount of proteins labelled with a fluorescent marker. For example, (Fig. 5.9) shows the measurement of TNF receptor 1 on KYM-1 cells. Thus, it can be seen as a high-speed imaging without spatial resolution.

Certain flow cytometers are combined with a microscope thus measuring not only the overall fluorescence intensity, but also its spatial distribution in each cell. Another variation is fluorescence-activated cell sorting (FACS) where the cells are sorted according to the measurement of the flow cytometer, for example into cells with a fluorescent marker and those without.

Flow cytometry can be used to obtain time course data. As an example, the number of TNF receptor 1 was measured at several time points post stimulation with cycloheximide (CHX), a protein biosynthesis inhibitor, see Fig. 5.10. Figure 5.10a shows the distribution of the receptor number across the population of cells, one colour corresponds to a specific time sample. The temporal evolution of the mean values is plotted as squares in Fig. 5.10b. This allowed for estimating the degradation rate of this protein.

5.3 Protein Mass Spectrometry and Metabolic Flux Analysis

In protein mass spectrometry, proteins are first separated according to their mass using electrophoresis or chromatography. The proteins in the mass range of interest are then digested into peptides, whose abundance is then measured by mass analysis. Using bioinformatics approaches, the protein abundance is then calculated.

Mass spectrometry equipment is in itself very expensive, but individual experiments can also be costly, in particular for the labelling using isotopes, e.g. using of a heavy carbon (^{13}C) on one probe and a lighter one (^{12}C) on the other. This labelling




Fig. 5.11 Mass spectrometry measurements of the amino acid alanine in young rice plants grown under light cycles of 13 h light and 11 h dark, illustrated by the *black–white* bar on top. The concentration of alanine was measured in two experiments. Depicted are here their average (*diamonds*) and the respective standard error of the mean (*error bars*). Figure reprinted from [43]

allows for comparing in two samples the abundance of particular proteins, as the same protein will have a different weight in each sample. Label-free experiments are cheaper, but the quantification is of lower quality [47]. Thus, label-free mass spectrometry experiments would permit taking a higher number of temporal samples, at the price of a reduced quality of the quantification.


As example, the circadian variations of the central metabolism of rice plants was studied in [43], where 69 metabolites were measured with mass spectrometry every hour during a 24 h cycle. Each experiment was repeated once, just the strict minimum for calculating standard deviations. The time course of alanine is shown in Fig. 5.11.

Using substrates with standard carbon ^{12}C replaced at least in part by ^{13}C results in slightly heavier metabolites. This difference can be seen in mass spectrometry experiments. In combination with a model of the stoichiometry, this mass spectrometry data can be used to estimate the fluxes of metabolites. Such experiments, called metabolic flux analysis, can not only be done at steady-state, but also in temporal sample after a stimulation. This is called instationary ^{13}C -based metabolic flux analysis [63], and sampling times of 4 s are possible [26]. These flux measurement combined with concentration measurements are very useful for determining kinetic parameters. Often, Metabolic Flux Analysis is visualised by varying the width of the arrows in diagrams of the metabolism, with larger arrows corresponding to larger flows, see e.g. Fig. 5.12.

5.4 Microscopy

Microscopy-based experiments offer the possibility of measuring the presence of molecules in single cell. This is achieved via the addition of a fluorescent dye that reacts with the molecule of interest allows for an indirect measure of that molecules concentration. This is for example employed for intracellular calcium (Ca^{2+}), or for

Fig. 5.12 Metabolic fluxes as estimated from ^{13}C mass spectrometry and a stoichiometric model of the extended central metabolism [22]. The upper value represents the net flux (in direction of the arrow), the lower value (on blue background) the backward flux. Figure reprinted from [22]




caspace-3 as shown in Fig. 5.17. The dye can also be directly incorporated into the protein of interest using genetic engineering techniques, as for example to observe the translocation of NF- κ B from the cytoplasm to the nucleus (Fig. 5.14). For optical microscopy, the spatial resolution is approx. 200–300 nm, which can be improved using so-called super-resolution techniques microscopy [5].

For detecting the localisation of a protein of interest, two or more fluorophores are simultaneously employed. Common are the staining of a compartment by using markers specific for e.g. the nucleus, mitochondria or the cell membrane with one of the fluorophores. Then the microscopy measurements aim at quantifying how much of the protein of interest is in the selected compartment and how much outside. As example, a red staining of the nucleus is combined with a green one for the protein STYX, allowing to compare its localisation in the unstimulated and stimulated case, see Fig. 5.13.

Such multiple staining can also be used for arbitrary proteins, thus measuring protein–protein interactions, for example to observe when second messenger molecules bind to a specific receptor.

Classical microscopy can poorly detect the vertical localisation of the samples. To overcome this, the reflection from outside the focal plane need to be eliminated, as is for example achieved by confocal microscopes. There, the images correspond to virtual slices of approx. 0.5–1 nm. Using piezo actuators, the focal plane can be moved vertically. This enables the scanning of several images over the vertical axis, thus giving 3D images. A better resolution in the vertical space, compared to confocal microscopy, as well as deeper scanning within a tissue can be obtained via

Fig. 5.13 Localisation of STYX tagged by the *yellow* fluorescent protein, YFP, (*green*) in the unstimulated, i.e. serum-starved case (*top row*) and serum-stimulated (*bottom row*) [40]. The nuclei have been stained in *red* (Alexa Fluor 647). An overlay of the images on one row allows for quantifying how much STYX is in the nucleus, and how much is in other cell compartments. Figure reprinted from [40]



two photon microscopy [33], at the price of more expensive equipment and slower image acquisition.

For a more detailed overview of microscopy techniques than this section, see e.g. [52].

5.4.1 Live Cell Imaging

Microscopy allows for obtaining time series measurements in a natural cell culture environment. This so-called live cell imaging can for example be used to observe translocation of secondary signalling molecules into the nucleus. Figure 5.14 shows NF- κ B marked by GFP, allowing to measure not only the total amount per cell, but also its localisation. The translocation of NF- κ B from the cytoplasm to the nucleus and back is clearly visible by comparing Fig. 5.14a–c.

5.4.2 Photobleaching

A challenge in live cell fluorescence imaging is photobleaching, i.e. the reduction of the fluorescence over time due to photochemical destruction of the fluorophore. This can be reduced by lowering the light intensity at the price of higher noise, and by lowering the sampling rate. Fast sampling rates, as for example 1 Hz in [39], require additional measurements to estimate the time frame of interest, as a prolonged stimulation of the fluorophore would result in its bleaching and thus in a decreasing signal intensity, see e.g. Fig. 5.15.

Photobleaching can also be utilised on purpose. In fluorescence recovery after photobleaching (FRAP), the fluorophores in the region of interest (ROI) are turned

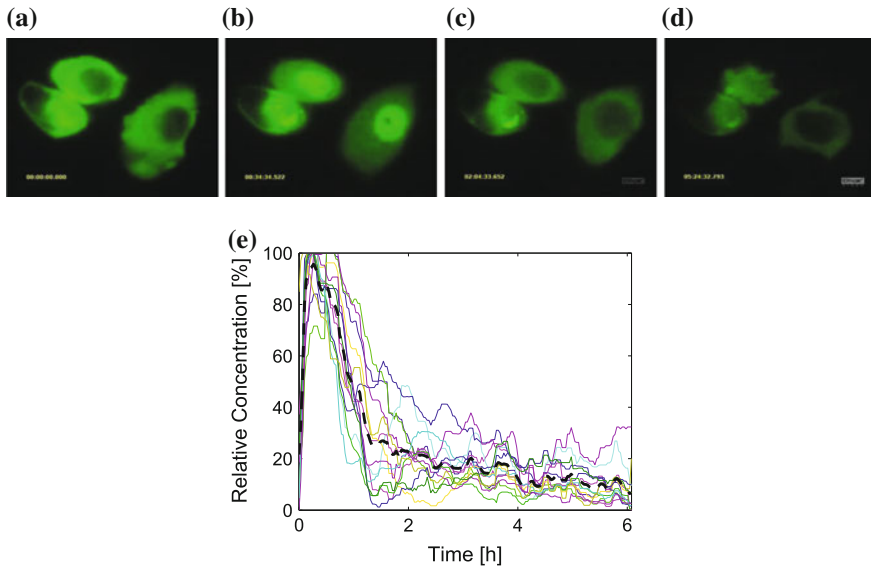


Fig. 5.14 Single cell fluorescence measurements of GFP-marked NF- κ B. Fluorescence live cell imaging of single cells with continuous stimulation of 10 ng ml^{-1} TNF. **a-d** Microscopy images of three cells at different time points post stimulation. **e** Quantification of the fluorescence signal in the nucleus relative to the cytoplasm for 14 individual cells. The *thick black dashed line* shows their average. Figure reprinted from [44]




Fig. 5.15 Fluorescent microscopy at relatively high sampling rates. At first, the GFP tagged cytochrome-c is mostly localised inside the mitochondria of the HeLa cells. Several minutes after stimulation with apoptotic agents, the permeabilisation of the membrane of individual mitochondria occurs within about half a minute. However, the individual mitochondria do not permeabilise simultaneously. **a** Fluorescent microscopy images, **b** quantification of the intensity of the fluorescence in the mitochondria in the two regions d_1 and d_2 shown in the *top left* image. Reprinted by permission from Macmillan Publishers Ltd: Cell Death and Differentiation [39], copyright 2009




Fig. 5.16 Dynamical measurements induced by photobleaching in the cytoplasm. HeLa cells were transfected with cDNA encoding GFP-ERK2. Between $t = -20$ s and $t = 0$ s, the cytosole was bleached. The recovery of the fluorescence could then be observed, as exemplified at $t = 60$ s. Figure reprinted from [40]

off by prolonged excitation. Then, time-lapse microscopy reveals how fluorophores from the neighbourhood can move into the ROI, or are newly expressed. Thus, FRAP experiments allow for the estimation of kinetic parameters involved in diffusion, active transport or expression rates. See Fig. 5.16 for an example. Similar are fluorescence loss in photobleaching (FLIP) experiments, where the ROI is a non-bleached region and the decrease of fluorescence is measured, see e.g. [52].

5.4.3 Förster Resonance Energy Transfer

Förster resonance energy transfer (FRET), often called fluorescence resonance energy transfer, is useful for measuring the amount of two fluorophores that are close to each other. The stimulation of the fluorophore with higher energy (lower wavelength) can lead to an energy transfer to the fluorophore of higher wavelength. For example, the transfer from CFP (cyan, 480 nm) to YFP (yellow, 535 nm) is possible if donor and acceptor are within a few nanometres [35]. FRET can be used to observe protein–protein binding by marking each type of protein with one fluorophore. It can also serve to make visible the activity of protease: a probe linking two fluorophores sends a FRET signal as long as it is intact, but not any more after having been cut by its protease. For example, Fig. 5.17 shows FRET activity of a caspase-3 substrate containing a cyan (CFP) and a yellow (YFP) fluorophore. The FRET from cyan to yellow sharply diminishes in individual cells within a few minutes between 2 and 7 h post stimulation with TNF. This corresponds to a cell-to cell variability of when caspase-3 gets activated.

5.4.4 Microinjection

In fluorescence microscopy, the fluorophore is usually added externally on an antibody, via transfection or via genetic modification. In microinjection, the fluorescent protein or dye is injected directly into the cells with a micropipette. The injected

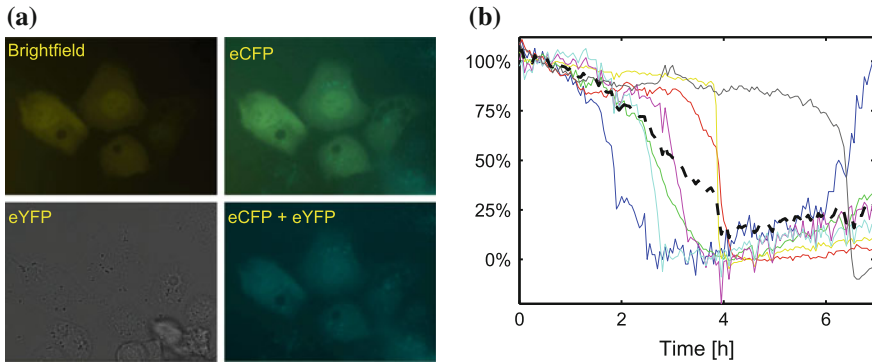


Fig. 5.17 Förster resonance energy transfer (FRET) at the example of a caspase-3 FRET probe [44]. **a** Microscopy images of cells in standard illumination (bright field, *top left*), cyan illumination and observation (*top right*), yellow illumination and observation (*bottom left*) and cyan illumination and yellow observation (*bottom right*). **b** Quantification of the FRET signal over time, normalised by its initial value. Each line corresponds to the FRET intensity integrated over an individual cell. Figure reprinted from [44]

amount cannot be fully controlled as it depends on the resistance in the cell. Therefore, the injected solution is coloured with an additional dye, thus allowing for an indirect quantification. Now the cells can be followed over time and a dose-response relation measured. Microinjection is therefore useful for quantitative perturbation experiments in various fields such as cell signalling, development and neurobiology.

Calcium propagation was one of the first systems in which microinjection experiments [42] were utilised in the development of mathematical models, see for example [51, 61]. Other systems where microinjection experiments were used are transport across the cell membrane [4] and apoptosis signalling, see [45] and Fig. 5.18, which measures the dependence of the time of death on the injected amount of caspase-3 in KYM-1 cells.

5.4.5 Tissue Visualisation

The development of quantitative mathematical models on the tissue scale as for example modelling of *in vivo* tumor growth [53] requires appropriate experimental techniques to quantify localisation, differentiation and death of cells. Different imaging approaches can be utilised for obtaining quantitative data in tissues with spatial resolution. For example, 3D confocal microscopy already presented above can be used at appropriate resolution to quantify cell migration [9], possible is also the utilisation of biomarkers, for example of proliferation or apoptosis, as in [58]. A key limitation is that light is the restricted penetration of light in tissues. This issue can partly be overcome using dynamic contrast-enhanced computed tomography,




Fig. 5.18 Time of death versus the amount of microinjected activated caspase-3. Single cells were microinjected with a mixture of caspase-3 and FITC-Dextran dye to quantify the relative microinjected amount in relative fluorescence units (RFU). **a–d** A number of cells where the injected *green* FITC-Dextran dye is visible. The microscopy images are taken at 00:04:36 **a**, 00:14:36 **b**, 00:30:36 **c** and 01:30:36 **d** after the microinjection. Morphological changes corresponding to cell death are already visible 14 min after injection of caspase-3. **e** For each single cell, a *star* represents the time to death (development of morphological signs of apoptosis) estimated from time lapse video versus the integrated initial FITC-Dextran fluorescence intensity. The colour of the stars differentiate data from three individual experiments. Cells that survived for more than 3 h were classified as survivors (infinite time to death). The *solid curve* shows the corresponding model simulation, superimposed by *box-plots* for the spread in the simulated cell population. **f** The percentage of dying cells as a function of fluorescence intensity is shown on the *left*, with the colour code shown on the *right*. Figure based on [44]

as used in [31] for quantifying thioacetamide-induced acute liver injury in rats. Such imaging approaches will help developing quantitative, not only phenomenological models of spatial phenomena on the tissue scale.

5.5 Genomics

High-throughput experimental techniques have revolutionised biology by significantly speeding up biological experiments. This is particularly the case in genome sequencing, which has seen in the past decades simultaneously a massive speed increase and cost reduction [37]. To illustrate this, genomics papers now compare routinely full genomes of different organisms [37] and a clinical usage of individual genome sequencing is almost within reach due to high-throughput sequencing, also called next-generation sequencing. For dynamical modelling, sequencing is useful in providing information on which proteins can be found in a cell type, though their amount and usually their chemical properties cannot be directly determined from the sequencing.

Due to DNA amplification techniques and cell sorting capabilities, it is possible to use a single cell as raw material for the sequencing [65]. For cells with high mutation rates like cancer cells or populations of heterogeneous cells, single cell sequencing could be informative for modelling by uncovering which proteins and even protein domains actually vary within the population. This small size of required DNA material is not only an advantage as contaminations by human DNA has been found in every sixth database of genome sequences of non-primates [27]. Contamination of human DNA is more difficult to spot, even though this is of high interest not only for biological and medical research, but also in forensic DNA analysis.

There is a distinct advantage of mRNA sequencing. Unlike microarrays, quantitative polymerase chain reaction (qPCR), or western blots for that matter, RNA sequencing allows to compare the mRNA levels of one gene to another. Usually the values obtained from mRNA sequencing are normalized to the entire RNA content of the sample (called the library).

5.5.1 DNA Microarrays

The deciphering of the genome combined with microprinting capabilities has made possible the fabrication of DNA microarray chips that can measure the expression of a number of genes, up to the full genome. This data is rarely useful for estimating kinetic parameters [11]. The main reason for this is that the cost per experiment is still so high that only few time points can reasonably be obtained. For example, the dynamic response of 3429 genes was measured at seven time points in response to oxidative stress [12], 12 time points for 5003 genes in oestrogen-induced signalling




Fig. 5.19 Gene expression levels measured by DNA microchips of butterfly genes (*vertical axis*) in 32 different spatial or temporal samples (*horizontal axis*) [16]. This figure shows 14 genes (out of 8000) involved in melanin synthesis. A two colour microarray hybridisation was used such that a *green* spot corresponds to a lower level of mRNA compared to the control, while *red* stands for a higher one, see scale on the *right*. Figure reprinted from [16]

[66], or 32 samples from butterflies were taken to measure the expression of 8000 genes [16], see Fig. 5.19.


5.5.2 Synthetic Biology

A direct consequence of the advances in genomics is synthetic biology, which aims at building novel, “engineered” systems using biological parts. The first examples were merely proof of concepts, building switches, clocks or logic elements. Synthetic biology examples have become more elaborate, targeting a wide area of applications such as sensing, production of macromolecules and therapeutics [25]. Synthetic biology can also be utilised as an experimental technique for unravelling complex dynamical systems such as clocks [21], or gaining a deeper understanding in signalling modules such as the MAP kinase [15]. As depicted in Fig. 5.20, signalling modules are not only modelled mathematically, but also build synthetically. With extra markers or tuning knobs, this makes more advanced experiments possible. In other words, synthetic biology allows for a “proof by construction”.

5.6 Discussion

The previous sections presented several experimental techniques frequently used for identifying models in systems biology and computational models of biological processes in general, see also Chap. 9 by Sunnaker et al. While each has its own advantages and disadvantages, there are a few global issues that affect most of them.

Fig. 5.20 Synthetic biology approach to studying system properties of a signalling modules such as the MAPK kinase system. The system is build as a synthetic biology module and modelled mathematically. Then, the influence of inputs signals and regulatory factors can be studied experimentally, serving in validating and improving the mathematical model as well as in gaining a deeper understanding of the system. Reprinted by permission from Elsevier Inc.: Cell [34], copyright 2011



Among these are the heterogeneity in a population of cells, the concentration ranges of macromolecules and practical identifiability, all discussed below.

5.6.1 Single Cell versus Population Measurements in Heterogeneous Populations

The experimental techniques presented in Sects. 5.2 and 5.3 inherently measure the average of approx. 1×10^4 – 1×10^6 cells. This is fine if the cells are not only identical, but also synchronous. If not, population measurements are misleading. Two comprehensible examples are oscillatory and switch-like systems, see Fig. 5.21. While the average of sinusoidals is almost constant, Fig. 5.21a, the average of steps is a smooth transition (Fig. 5.21b). In both cases, the single cell and the average correspond to fundamentally different behaviours. The same phenomenon can be observed in Figs. 5.14e and 5.17b, which show single cell data as well as a calculated average.

Cell-to-cell heterogeneity requires on the experimental side single cell measurements as for example those discussed in Sect. 5.4. On the modelling side, appropriate model structures are needed to enable the comparison of experimental data and simulations both on the single cell level as well as on the cell population level.

The identification of a single cell model directly from population measurements can be misleading as oscillations or quick, step-like responses might not be visible in the averaging measurement. One possible model structure that enables the merging of single cell and cell population data are cell ensemble models [20], as




Fig. 5.21 Illustration of single cell versus population measurements. **a** Oscillatory cell responses are represented by sinusoids. 10 sinusoidal samples with identical frequency and random phase shift (*colour lines*). The population measurement corresponds to an average, here the average of 100 such sinusoids (*thick black line*), which is again a sinusoidal with the same frequency, but with significantly smaller amplitude. **b** Switching phenomena can be represented by steps, whose rise time varies from cell to cell. 10 steps with random rise time between 0 and 1 (*colour lines*). The population measure is represented by the average of 100 such steps (*thick black line*). The population average is a smooth transition from low to high, in contrast to the step-wise switching of individual cells

depicted in Fig. 5.22. A cell ensemble model is an ensemble of single cell models. The individual cell models can for example be obtained by estimating the parameters of a nominal model and distributing some or all of the kinetic parameters across the ensemble, e.g. using log-normal distributions. Such a log-normal distribution was for example measured for TNF receptor 1 in (Fig. 5.9). Thus, assuming homogeneous degradation rates, the production rates are then also log-normally distributed.

The cell ensemble model of TNF-induced signalling [45] is able to produce both single cell and cell population responses as illustrated by the oscillating response of NF- κ B and the step-like one of caspase-3 in Fig. 5.23.

Heterogeneous cell populations are naturally present in tissues, due to cells of different type, age or environmental condition. This is for example the case of




Fig. 5.22 Illustration of single cell and population experiments together with the corresponding model simulations. While some experiments yield single cell data (*top left part*), others give average measurements over the population (*lower left*). To reproduce this in simulation, single cell models are generated and simulated (*top right*). The simulations can then be combined to obtain averages, i.e. simulation of the population. Both on the single cell and on the cell population level can model and experiments be compared [45]. Figure reprinted from [44]




Fig. 5.23 Illustration of single cell and population model simulations. The simulated responses to TNF of the proteins NF-κB **a** and caspase-3 **b** in individual cells of the cell population model in [45] (*grey lines*). The calculated population average is shown in *red*

bone remodelling in biomechanics [36], in neuroscience where the neuronal energy metabolism relies on astrocytes [62] or in embryogenesis [60].

5.6.2 Amounts of Proteins

Abundance measurements show that proteins concentrations span six orders of magnitude in *S. cerevisiae* [18] and seven orders of magnitude in human cells [6]. Covering this full dynamic range is very difficult, in particular by a single experimental technique. In particular, low abundant and only transiently abundant proteins or molecules are very difficult to quantify experimentally.

Fluorescent tags can be added to an intracellular protein by modifying its genome or by adding a vector containing this modified genome. In the latter case, two different versions of the protein exist: the original one as well as the one with the tag. This needs to be taken into account when estimating concentrations, for example by modelling this transfection step [50]. Furthermore, the number of vectors per cell can significantly vary across the population, thus resulting in a heterogeneity. Additionally, care needs to be taken that the modification does not significantly alter the protein's functionality or that the overexpression alters the qualitative behaviour.

5.6.3 Practical Identifiability

Identification of kinetic parameters requires that they impact on the measured signals. In systems biology models, quite a few systems are insensitive to a large number of their parameters, see e.g. [14]. This is for example the case in models that are on purpose close to bifurcation points, as for example to achieve input dependent switching delays [59]. Therefore, the process of developing a mathematical model will usually iterate between experiments, modelling and model analysis. The latter often contains parameter sensitivity analyses, see also Chap. 13 by van Schepdael et al. and Chap. 15 by Mengoni et al. Even though the link between sensitivity and identifiability seems obvious, they are not exactly antonymous [13]. Thus, additional parameter identifiability analyses are necessary for designing optimal experiments [38].

5.6.4 Databases

There is a quickly growing number of databases providing useful information for developing quantitative, dynamical models. The largest database of models is Bio-models [8], with approx. 550 manually curated models. Values of kinetic parameters can be found in SABIO-RK, a curated database with information on biochemical reactions, their kinetic rate equations with parameters and experimental conditions

[41], see also [64], or in KENDA (Kinetic ENzyme DATA), an extension to the BRENDA database, which provides an additional overview on functional kinetic data of enzymes [23], see also [46].

The decay rates of almost 20,000 mRNAs was experimentally measured and can be queried in [30], see also [48]. In mammalian cells, mRNA have a much faster turnover, as their life times are in the order of minutes, compared to hours or days for proteins.

Other databases possibly useful in this context are the ADME Database with data on interactions of substances with drug metabolising enzymes and with drug transporters [1], and the NIST Chemical Kinetics Database [32], which lists kinetic parameters of more than 38,000 thermal gas-phase chemical reactions.

5.7 Summary and Outlook

Dynamic phenomena require time series measurements. This chapter has presented several experimental methodologies for acquiring such data in intracellular signalling systems. These range from biochemical characterisation such as western blot and enzyme assays over flow cytometry and DNA microarrays to microscopy imaging.

Protein–protein interaction data can help inferring the model structure of signal transduction pathways [54]. High throughput experiments such as DNA microarrays, RNA sequencing and proteomics (mass spectroscopy) can also give a global overview of what is changing in a few different conditions or time points. Then, one can move on by focussing on these changes and generate better quality, higher resolution and dynamic data for mathematical modelling and parameter estimation. These experiments can also serve in calibrating kinetic parameters playing a role in steady-states.

Mass spectrometry and DNA microarray are starting to be used for measuring multiple time points, as shown by some examples (Figs. 5.11 and 5.19). Flow cytometry has some aspects of a high-throughput technology and is regularly employed for quantifying kinetic parameters as this technique can easily and relatively cheaply be used to measure several temporal samples, with sampling rates of up to 40 wells per minute [10].

Most experiments for modelling signalling networks rely on low throughput technologies such as assays, gels and microscopy in order to achieve the required data resolution and quality. The advancement of automation in biological and pharmacological laboratories, as for example in high-throughput screening, will lower the effort and increase the reliability of time series measurements.

Microscale experimental setups will permit for improved single cell experiment by manipulating cells, see e.g. [24, 28, 49]. New developments also permit to down-scale experiments like qPCR [57], or improve spatial configuration of cells or an organism as e.g. in [56, 67]. Further advances are to be expected due to the increased employment of robotics, control and data processing in biological experiments.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

1. ADME: ADME database. <http://jp.fujitsu.com/group/kyushu/en/services/admedatabase/> (2014). Accessed 01 Oct 2014
2. Aebersold, R., Burlingame, A.L., Bradshaw, R.A.: Western blots versus selected reaction monitoring assays: Time to turn the tables? *Mol. Cell. Proteomics* **12**(9), 2381–2382 (2013). doi:[10.1074/mcp.E113.031658](https://doi.org/10.1074/mcp.E113.031658)
3. Aksamitiene, E., Hoek, J.B., Kholodenko, B., Kiyatkin, A.: Multistrip Western blotting to increase quantitative data output. *Electrophoresis* **28**(18), 3163–3173 (2007). doi:[10.1002/elps.200700002](https://doi.org/10.1002/elps.200700002)
4. Almquist, J., Lang, P., Prätzel-Wolters, D., Deitmer, J.W., Jirstrand, M., Becker, H.M.: A kinetic model of the monocarboxylate transporter MCT1 and its interaction with carbonic anhydrase II. *J. Comput. Sci. Syst. Biol.* **3**(5), 107–116 (2010). doi:[10.4172/jcsb.1000066](https://doi.org/10.4172/jcsb.1000066)
5. Ball, G., Parton, R.M., Hamilton, R.S., Davis, I.: A cell biologist's guide to high resolution imaging. In: Conn, P.M. (ed.) *Imaging and Spectroscopic Analysis of Living Cells Optical and Spectroscopic Techniques, Methods in Enzymology*, vol. 504, pp. 29–55. Academic Press (2012). doi:[10.1016/B978-0-12-391857-4.00002-1](https://doi.org/10.1016/B978-0-12-391857-4.00002-1)
6. Beck, M., Schmidt, A., Malmstroem, J., Claassen, M., Ori, A., Szymborska, A., Herzog, F., Rinner, O., Ellenberg, J., Aebersold, R.: The quantitative proteome of a human cell line. *Mol. Syst. Biol.* **7**, 549 (2011). doi:[10.1038/msb.2011.82](https://doi.org/10.1038/msb.2011.82)
7. Benachour, H., Bastogne, T., Toussaint, M., Chemli, Y., Sève, A., Frochet, C., Lux, F., Tillement, O., Vanderesse, R., Barberi-Heyob, M.: Real-time monitoring of photocytotoxicity in nanoparticles-based photodynamic therapy: A model-based approach. *PLoS ONE* **7**(11), e48617 (2012)
8. Biomodels: Biomodels. <http://www.biomodels.org> (2014). Accessed 01 Oct 2014
9. Blacher, S., Ercicum, C., Lenoir, B., Paupert, J., Moraes, G., Ormenese, S., Bullinger, E., Noël, A.: Cell invasion in the spheroid sprouting assay: A spatial organisation analysis adaptable to cell behaviour. *PLoS ONE* **9**(5), e97,019 (2014). doi:[10.1371/journal.pone.0097019](https://doi.org/10.1371/journal.pone.0097019)
10. Black, C.B., Duensing, T.D., Trinkle, L.S., Dunlay, R.T.: Cell-based screening using high-throughput flow cytometry. *Assay Drug Dev. Technol.* **9**(1), 13–20 (2011). doi:[10.1089/adt.2010.0308](https://doi.org/10.1089/adt.2010.0308)
11. Csikász-Nagy, A.: Computational systems biology of the cell cycle. *Briefings Bioinform.* **10**(4), 424–434 (2009). doi:[10.1093/bib/bbp005](https://doi.org/10.1093/bib/bbp005)
12. Deferme, L., Briedé, J.J., Claessen, S.M.H., Jennen, D.G.J., Cavill, R., Kleinjans, J.C.S.: Time series analysis of oxidative stress response patterns in HepG2: a toxicogenomics approach. *Toxicology* **306**, 24–34 (2013). doi:[10.1016/j.tox.2013.02.001](https://doi.org/10.1016/j.tox.2013.02.001)
13. Dobre, S., Bastogne, T., Profeta, C., Barberi-Heyob, M., Richard, A.: Limits of variance-based sensitivity analysis for non-identifiability testing in high dimensional dynamic models. *Automatica* **48**(11), 2740–2749 (2012). doi:[10.1016/j.automatica.2012.05.004](https://doi.org/10.1016/j.automatica.2012.05.004)
14. Erguler, K., Stumpf, M.P.H.: Practical limits for reverse engineering of dynamical systems: A statistical analysis of sensitivity and parameter inferability in systems biology models. *Mol. BioSyst.* **7**, 1593–1602 (2011). doi:[10.1039/c0mb00107d](https://doi.org/10.1039/c0mb00107d)
15. Furukawa, K., Hohmann, S.: Synthetic biology: Lessons from engineering yeast MAPK signalling pathways. *Mol. Microbiol.* **88**(1), 5–19 (2013). doi:[10.1111/mmi.12174](https://doi.org/10.1111/mmi.12174)
16. Futahashi, R., Shirataki, H., Narita, T., Mita, K., Fujiwara, H.: Comprehensive microarray-based analysis for stage-specific larval camouflage pattern-associated genes in the swallowtail butterfly. *Papilio Xuthus*. *BMC Biol.* **10**, 46 (2012)

17. Gassmann, M., Grenacher, B., Rohde, B., Vogel, J.: Quantifying Western blots pitfalls of densitometry. *Electrophoresis* **30**(11), 1845–1855 (2009). doi:[10.1002/elps.200800720](https://doi.org/10.1002/elps.200800720)
18. Ghaemmaghami, S., Huh, W.K., Bower, K., Howson, R.W., Belle, A., Dephoure, N., O’Shea, E.K., Weissman, J.S.: Global analysis of protein expression in yeast. *Nature* **425**(6959), 737–741 (2003). doi:[10.1038/nature02046](https://doi.org/10.1038/nature02046)
19. Ghosh, R., Gilda, J.E., Gomes, A.V.: The necessity of and strategies for improving confidence in the accuracy of western blots. *Expert Rev. Proteomics* **11**(5), 549–560 (2014). doi:[10.1586/14789450.2014.939635](https://doi.org/10.1586/14789450.2014.939635)
20. Henson, M.A., Müller, D., Reuss, M.: Cell population modelling of yeast glycolytic oscillations. *Biochem. J.* **368**(Pt 2), 433–446 (2002). doi:[10.1042/BJ20021051](https://doi.org/10.1042/BJ20021051)
21. Hogenesch, J.B., Ueda, H.R.: Understanding systems-level properties: Timely stories from the study of clocks. *Nat. Rev. Genet.* **12**, 407–416 (2011). doi:[10.1038/nrg2972](https://doi.org/10.1038/nrg2972)
22. Jorda, J., Suarez, C., Carnicer, M., ten Pierick, A., Heijnen, J., van Gulik, W., Ferrer, P., Albiol, J., Wahl, A.: Glucose-methanol co-utilization in *Pichia pastoris* studied by metabolomics and instantaneous ¹³C flux analysis. *BMC Syst. Biol.* **7**(1), 17 (2013). doi:[10.1186/1752-0509-7-17](https://doi.org/10.1186/1752-0509-7-17)
23. KENDA database: Database of mrna life times. <http://www.brenda-enzymes.org> (2014). Accessed 01 Oct 2014
24. Khademhosseini, A., Langer, R., Borenstein, J., Vacanti, J.P.: Microscale technologies for tissue engineering and biology. *Proc. Natl. Acad. Sci. USA* **103**(8), 2480–2487 (2006). doi:[10.1073/pnas.0507681102](https://doi.org/10.1073/pnas.0507681102)
25. Khalil, A.S., Collins, J.J.: Synthetic biology: Applications come of age. *Nat. Rev. Genet.* **11**, 367–379 (2010). doi:[10.1038/nrg2775](https://doi.org/10.1038/nrg2775)
26. Lange, H.C., Eman, M., van Zuijlen, G., Visser, D., van Dam, J.C., Frank, J., de Mattos, M.J., Heijnen, J.J.: Improved rapid sampling for in vivo kinetics of intracellular metabolites in *Saccharomyces cerevisiae*. *Biotechnol. Bioeng.* **75**(4), 406–415 (2001). doi:[10.1002/bit.10048](https://doi.org/10.1002/bit.10048)
27. Longo, M.S., O’Neill, M.J., O’Neill, R.J.: Abundant human DNA contamination identified in non-primate genome databases. *PLoS ONE* **6**(2), e16,410 (2011). doi:[10.1371/journal.pone.0016410](https://doi.org/10.1371/journal.pone.0016410)
28. Marques, M.P.C., Fernandes, P.: Microfluidic devices useful tools for bioprocess intensification. *Molecules* **16**, 8368–8401 (2011). doi:[10.3390/molecules16108368](https://doi.org/10.3390/molecules16108368)
29. Miller, L.: Analyzing western blots with Image Studio Lite. <http://lukemiller.org/index.php/2013/02/analyzing-western-blots-with-image-studio-lite/> (2013). Accessed 01 Oct 2014
30. mRNA database: Database of mrna life times (2014). Accessed 01 Oct 2014
31. Murase, K., Kobayashi, S., Kitamura, A., Matsushita, T., Saito, S., Nishiura, M.: An empirical mathematical model applied to quantitative evaluation of thioacetamide-induced acute liver injury in rats by use of dynamic contrast-enhanced computed tomography. *Radiol. Phys. Technol.* **6**(1), 115–120 (2013). doi:[10.1007/s12194-012-0177-y](https://doi.org/10.1007/s12194-012-0177-y)
32. NIST: NIST chemical kinetics database. <http://kinetics.nist.gov/kinetics/welcome.jsp> (2014). Accessed 01 Oct 2014
33. Oheim, M., Michael, D.J., Geisbauer, M., Madsen, D., Chow, R.H.: Principles of two-photon excitation fluorescence microscopy and other nonlinear imaging approaches. *Adv. Drug Delivery Rev.* **58**(7), 788–808 (2006). doi:[10.1016/j.addr.2006.07.005](https://doi.org/10.1016/j.addr.2006.07.005)
34. O’Shaughnessy, E.C., Palani, S., Collins, J.J., Sarkar, C.A.: Tunable signal processing in synthetic MAP kinase cascades. *Cell* **144**(1), 119–131 (2011). doi:[10.1016/j.cell.2010.12.014](https://doi.org/10.1016/j.cell.2010.12.014)
35. Patterson, G.H., Piston, D.W., Barisas, B.G.: Förster distances between green fluorescent protein pairs. *Anal. Biochem.* **284**(2), 438–440 (2000). doi:[10.1006/abio.2000.4708](https://doi.org/10.1006/abio.2000.4708)
36. Peterson, M.C., Riggs, M.M.: A physiologically based mathematical model of integrated calcium homeostasis and bone remodeling. *Bone* **46**(1), 49–63 (2010). doi:[10.1016/j.bone.2009.08.053](https://doi.org/10.1016/j.bone.2009.08.053)
37. Pettersson, E., Lundeberg, J., Ahmadian, A.: Generations of sequencing technologies. *Genomics* **93**(2), 105–111 (2009). doi:[10.1016/j.ygeno.2008.10.003](https://doi.org/10.1016/j.ygeno.2008.10.003)
38. Raue, A., Becker, V., Klingmüller, U., Timmer, J.: Identifiability and observability analysis for experimental design in nonlinear dynamical models. *Chaos: An Interdisciplinary J. Nonlinear Sci.* **20**(4), 045105 (2010). doi:[10.1063/1.3528102](https://doi.org/10.1063/1.3528102)

39. Rehm, M., Huber, H.J., Hellwig, C.T., Anguissola, S., Dussmann, H., Prehn, J.H.M.: Dynamics of outer mitochondrial membrane permeabilization during apoptosis. *Cell Death Differ.* **16**(4), 613–623 (2009). doi:[10.1038/cdd.2008.187](https://doi.org/10.1038/cdd.2008.187)
40. Reiterer, V., Fey, D., Kolch, W., Kholodenko, B.N., Farhan, H.: Pseudophosphatase STYX modulates cell-fate decisions and cell migration by spatiotemporal regulation of ERK1/2. *Proc. Natl. Acad. Sci. USA* **110**(31), E2934–E2943 (2013). doi:[10.1073/pnas.1301985110](https://doi.org/10.1073/pnas.1301985110)
41. Sabio-rk: Sabio-rk. <http://sabio.villa-bosch.de> (2014). Accessed 01 Oct 2014
42. Sanderson, M.J., Charles, A.C., Dirksen, E.R.: Mechanical stimulation and intercellular communication increases intracellular ca^{2+} in epithelial cells. *Cell Regul.* **1**, 585–596 (1990)
43. Sato, S., Arita, M., Soga, T., Nishioka, T., Tomita, M.: Time-resolved metabolomics reveals metabolic modulation in rice foliage. *BMC Syst. Biol.* **2**(1), 51 (2008). doi:[10.1186/1752-0509-2-51](https://doi.org/10.1186/1752-0509-2-51)
44. Schliemann, M.: Heterogeneity in Tumour Necrosis Factor induced Pro- and Anti-Apoptotic Signalling from a Systems Biology Perspective. Shaker Verlag, Aachen, Germany (2013)
45. Schliemann, M., Bullinger, E., Borchers, S., Allgöwer, F., Findeisen, R., Scheurich, P.: Heterogeneity reduces sensitivity of cell death for TNF-stimuli. *BMC Syst. Biol.* **5**, 204 (2011). doi:[10.1186/1752-0509-5-204](https://doi.org/10.1186/1752-0509-5-204)
46. Schomburg, I., Chang, A., Placzek, S., Söhngen, C., Rother, M., Lang, M., Munaretto, C., Ulas, S., Stelzer, M., Grote, A., Scheer, M., Schomburg, D.: BRENDA in 2013: Integrated reactions, kinetic data, enzyme function data, improved disease classification: New options and contents in BRENDA. *Nucleic Acids Research* 41(Database issue), D764–D772 (2013). doi:[10.1093/nar/gks1049](https://doi.org/10.1093/nar/gks1049)
47. Schulze, W.X., Usadel, B.: Quantitation in mass-spectrometry-based proteomics. *Annual Rev. Plant Biol.* **61**(1), 491–516 (2010). doi:[10.1146/annurev-arplant-042809-112132](https://doi.org/10.1146/annurev-arplant-042809-112132)
48. Sharova, L.V., Sharov, A.A., Nedorezov, T., Piao, Y., Shaik, N., Ko, M.S.: Database for mRNA half-life of 19977 genes obtained by dna microarray analysis of pluripotent and differentiating mouse embryonic stem cells. *DNA Res.* **16**, 45–58 (2009)
49. Sinha, A., Jebraïl, M.J., Kim, H., Patel, K.D., Branda, S.S.: A versatile automated platform for micro-scale cell stimulation experiments. *J. Visualized Exp.* (78) (2013). doi:[10.3791/50597](https://doi.org/10.3791/50597)
50. Smith, E.M., Mueller, J.D.: The statistics of protein expression ratios for cellular fluorescence studies. *Eur. Biophys. J.* **41**(3), 341–352 (2012). doi:[10.1007/s00249-012-0792-x](https://doi.org/10.1007/s00249-012-0792-x)
51. Sneyd, J., Charles, A.C., Sanderson, M.J.: A model for the propagation of intercellular calcium waves. *Am. J. Phys. Cell Phys.* **266**, C293–C302 (1994)
52. Spiller, D.G., Wood, C.D., Rand, D.A., White, M.R.H.: Measurement of single-cell dynamics. *Nature* **465**, 736–745 (2010). doi:[10.1038/nature09232](https://doi.org/10.1038/nature09232)
53. Stapleton, S., Milosevic, M., Allen, C., Zheng, J., Dunne, M., Yeung, I., Jaffray, D.A.: A mathematical model of the enhanced permeability and retention effect for liposome transport in solid tumors. *PLoS ONE* **8**(12), e81157 (2013). doi:[10.1371/journal.pone.0081157](https://doi.org/10.1371/journal.pone.0081157)
54. Supper, J., Spangenberg, L., Planatscher, H., Dräger, A., Schroder, A., Zell, A.: BowTieBuilder: Modeling signal transduction pathways. *BMC Syst. Biol.* **3**(1), 67 (2009). doi:[10.1186/1752-0509-3-67](https://doi.org/10.1186/1752-0509-3-67)
55. Taylor, S.C., Berkelman, T., Yadav, G., Hammond, M.: A defined methodology for reliable quantification of western blot data. *Mol. Biotechnol.* **55**(5), 217–226 (2013). doi:[10.1007/s12033-013-9672-6](https://doi.org/10.1007/s12033-013-9672-6)
56. Toetsch, S., Olwell, P., Prina-Mello, A., Volkov, Y.: The evolution of chemotaxis assays from static models to physiologically relevant platforms. *Integr. Biol.* **1**(2), 170–181 (2009)
57. Toumazou, C., Shepherd, L.M., Reed, S.C., Chen, G.I., Patel, A., Garner, D.M., Wang, C.J.A., Ou, C.P., Amin-Desai, K., Athanasiou, P., Bai, H., Brizido, I.M.Q., Caldwell, B., Coomber-Alford, D., Georgiou, P., Jordan, K.S., Joyce, J.C., La Mura, M., Morley, D., Sathyavrudhan, S., Temelso, S., Thomas, R.E., Zhang, L.: Simultaneous DNA amplification and detection using a pH-sensing semiconductor system. *Nat. Methods* **10**(7), 641–646 (2013). doi:[10.1038/nmeth.2520](https://doi.org/10.1038/nmeth.2520)
58. Tran, P.T., Bendapudi, P.K., Lin, H.J., Choi, P., Koh, S., Chen, J., Horng, G., Hughes, N.P., Schwartz, L.H., Miller, V.A., Kawashima, T., Kitamura, T., Paik, D., Felsner, D.W.: Survival

- and death signals can predict tumor response to therapy after oncogene inactivation. *Sci. Transl. Med.* **3**(103), 103ra99 (2011). doi:[10.1126/scitranslmed.3002018](https://doi.org/10.1126/scitranslmed.3002018)
59. Trotta, L., Bullinger, E., Sepulchre, R.: Global analysis of dynamical decision-making models through local computation around the hidden saddle. *PLoS ONE* **7**(3), e33110 (2012). doi:[10.1371/journal.pone.0033110](https://doi.org/10.1371/journal.pone.0033110)
 60. Umulis, D.M., Othmer, H.G.: The importance of geometry in mathematical models of developing systems. *Genetics of system biology. Current Opin. Genet. Devel.* **22**(6), 547–552 (2012). doi:[10.1016/j.gde.2012.09.007](https://doi.org/10.1016/j.gde.2012.09.007)
 61. Warren, N.J., Tawhai, M.H., Crampin, E.J.: Mathematical modelling of calcium wave propagation in mammalian airway epithelium: Evidence for regenerative ATP release. *Exp. Physiol.* **95**(1), 232–249 (2010). doi:[10.1113/expphysiol.2009.049585](https://doi.org/10.1113/expphysiol.2009.049585)
 62. Wellstead, P., Cloutier, M.: An energy systems approach to Parkinson's disease. *Wiley Interdisc. Rev.: Syst. Biol. Med.* **3**(1), 1–6 (2011). doi:[10.1002/wsbm.107](https://doi.org/10.1002/wsbm.107)
 63. Wiechert, W., Nöh, K.: From stationary to instationary metabolic flux analysis. In: Kragl, U. (Ed.) *Technology Transfer in Biotechnology, Advances in Biochemical Engineering*, Vol. 92, pp. 145–172, Springer, Berlin (2005). doi:[10.1007/b98921](https://doi.org/10.1007/b98921)
 64. Wittig, U., Kania, R., Golebiewski, M., Rey, M., Shi, L., Jong, L., Algae, E., Weidemann, A., Sauer-Danzwith, H., Mir, S., Krebs, O., Bittkowski, M., Wetsch, E., Rojas, I., Müller, W.: SABIO-RK-database for biochemical reaction kinetics. *Nucleic Acids Res.* **40**(D1), D790–D796 (2012). doi:[10.1093/nar/gkr1046](https://doi.org/10.1093/nar/gkr1046)
 65. Yilmaz, S., Singh, A.K.: Single cell genome sequencing. *Curr. Opin. Biotechnol.* **23**(3), 437–443 (2012). doi:[10.1016/j.copbio.2011.11.018](https://doi.org/10.1016/j.copbio.2011.11.018)
 66. Zhang, P., Mourad, R., Xiang, Y., Huang, K., Huang, T., Nephew, K., Liu, Y., Li, L.: A dynamic time order network for time-series gene expression data analysis. *BMC Syst. Biol.* **6**(Suppl 3), S9 (2012). doi:[10.1186/1752-0509-6-S3-S9](https://doi.org/10.1186/1752-0509-6-S3-S9)
 67. Zhao, X., Xu, F., Tang, L., Du, W., Feng, X., Liu, B.F.: Microfluidic chip-based *C. elegans* microinjection system for investigating cell-cell communication in vivo. *Biosens. Bioelectron.* **50**, 28–34 (2013)

Chapter 6

Statistical Data Analysis and Modeling

Millie Shah, Zeinab Chitforoushzadeh and Kevin A. Janes

Abstract The availability of large structured datasets has prompted the need for efficient data analysis and modeling techniques. In systems biology, data-driven modeling approaches create models of complex cellular systems without making assumptions about the underlying mechanisms. In this chapter, we will discuss eigenvalue-based approaches, which identify important characteristics (information) of big datasets through decomposition and dimensionality reduction. We intend to address singular value decomposition (SVD), principle component analysis (PCA), and partial least squares regression (PLSR) approaches for data-driven modeling. In multi-linear systems (that share characteristics such as time points, measurements, etc.), tensor decomposition becomes particularly important for understanding higher-order datasets. Therefore, we will also discuss how to scale up these methods to tensor decomposition using an example dealing with host-cell responses to viral infection.

Keywords Data-driven modeling · Statistical modeling · Eigenvalue-based approaches · Singular value decomposition (SVD) · Principal component analysis (PCA) · Partial least squares regression (PLSR)

6.1 Introduction

A recurring uncertainty in biology is the molecular underpinnings of an observed cellular or tissue-level phenomenon. With sufficient knowledge about the relevant molecular mechanisms, it is now possible to model large systems of biochemical reactions accurately by simulation [1–3, 48]. However, even for well-studied pathways such as receptor tyrosine kinases [55], these hypothesis driven models, including ODE (ordinary differential equation)/PDE (partial differential equation) and stochastic models, quickly uncover gaps in our understanding [43]. Often, the phenomenon of interest

M. Shah · Z. Chitforoushzadeh · K.A. Janes (✉)
Department of Biomedical Engineering, University of Virginia,
Charlottesville, VA 22908, USA
e-mail: kaj5f@eservices.virginia.edu

© Springer International Publishing Switzerland 2016
L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_6

is so poorly characterized that we only really have a sense of the pathways that are important and a rudimentary rule set for how they could interact [3, 46].

In these circumstances, it can be advantageous to pursue statistical models that do not prescribe mechanisms but allow the data to define the system of interest [23, 25]. In statistical modeling, one must first collect a systematic dataset that has been designed to capture as many relevant variations and covariations as possible among genes, proteins, and cellular phenotypes [2, 16, 31]. Although not absolutely required, it is strongly recommended that the statistical approach be chosen conceptually before the data acquisition. Each class of models has its own set of strengths and weaknesses [23], and ideally the dataset should be tailored to exploit a model's strengths and avoid its weaknesses. Bayesian modeling [21, 50] is most effective for inferring the network structure of the phenomena being studied while techniques like partial least squares are ideal for predicting new behaviors [23]. Statistical models may be "mechanism free", but it is possible to guide models toward identifying new mechanisms by selecting the right biomolecular measurements and designing the experiments appropriately [25, 53].

With current technologies in molecular biology, any laboratory can now generate datasets that are highly multivariate [31]. Statistical modeling serves as a powerful way to extract as much information as possible from these often expensive and difficult-to-conceptualize datasets. The resulting patterns and relationships identified by statistical models are not always apparent when analyzing the full spectrum of the dataset, as it often contains measurements not significant to the system. Thus, the class of statistical models that we will discuss in this chapter center around those that build simplified representations of data to give a clearer picture of possible mechanisms underlying the system.

Usually, in modern biological datasets, we have many more variables per observation than observations of each variable. These "short and fat" data tables (or matrices) are inherently underconstrained; in frequentist statistics, it is equivalent to having fewer than zero degrees of freedom. Consequently, many of the dimensions are redundant with one another, in that they can be expressed as linear combinations of other variables. This redundancy allows the data matrix to be "reduced" in interesting and useful ways, depending on the type of statistical model and the overall goals of the study.

Here, we will review three main categories of statistical models that reduce the dimensions of multivariate datasets. We begin with singular value decomposition (SVD), which draws on the concept of eigenvalues and eigenvectors to decompose a matrix according to its eigenvalue spectrum. Then, we will discuss principal components analysis (PCA), which is conceptually akin to SVD but yields a factorized model that is more directly interpretable with respect to the starting dataset. Finally, we will link reduced dimensions to the concept of predictive statistical modeling through partial least squares regression (PLSR). As case studies, we include more modern implementations of SVD, PCA and PLSR, including tensor decomposition of data cubes or hypercubes, in anticipation for the types of structured datasets that will be forthcoming in molecular and cellular biology.

6.2 Singular Value Decomposition

Before going into detail about SVD computation, it is important to introduce some basic concepts from vector and matrix algebra. Most datasets can be organized as matrices with the rows indicating experimental observation, such as treatments and time points and the columns indicating variables, such as enzymatic activity and phosphoprotein levels.

One way to simplify multidimensional data is to focus on parts of the data that show the most variation. Linear algebra serves this purpose by finding orthogonal or linearly independent vectors in the data matrix. Since orthogonal vectors have zero projections into one another, they can act as latent variables onto which the data can be mapped [28]. Orthogonal vectors of a data matrix can be identified by calculating eigenvectors. The nonzero *eigenvector* (x) of matrix A satisfies the equation:

$$Ax = \lambda x \tag{6.1}$$

where A is a square matrix and λ is a scalar called an “eigenvalue”.

An eigenvector can serve as a new dimension along which the data can be projected. By definition, matrix A is an $n \times n$ square matrix. However, typical biological datasets have fewer observations than variables and thus are rarely square matrices with full rank. One way to solve this problem is by factorizing the data matrix using singular value decomposition.

6.2.1 Mathematical Framework

6.2.1.1 Singular Value Decomposition

Suppose that we define an $m \times n$ data matrix A that can be broken down into the product of three other matrices U , S , and V . This factorization results in the following equation:

$$A_{m \times n} = U_{m \times l} S_{l \times l} V_{l \times n}^T \tag{6.2}$$

where U is an $m \times l$ left-singular matrix, S is a square $l \times l$ diagonal matrix, V is an $l \times n$ right matrix, and U and V^T are orthogonal matrices. The diagonal entries in S are the singular values of A (square roots of non-zero eigenvalues of U and V^T) descending in magnitude from top left to bottom right, the columns in V^T are right-singular vectors and the columns in U are left-singular vectors [4] (Fig. 6.1). Once singular vectors are extracted, the significant ones can be determined and used for visualizing the data.




Fig. 6.1 SVD Decomposition Schematic. Decomposition of the yeast elutriation data from Spellman et al. [49] into a left singular-value matrix, a square matrix of eigenvalues (four eigenvalues shown), and a right singular-value matrix

6.2.2 Application of SVD to Gene Expression Data Analysis

Gene expression data is a good candidate for singular value decomposition based analysis due to the inherent noise in the measurements that makes the detection of small signals rather difficult. Alter et al. performed SVD analysis on the budding yeast elutriation gene microarray data published in [49]. The elutriation dataset used by Alter et al. contained 5,981 genes ($n=5,981$ genes) captured over the course of one yeast cell cycle (fourteen time points; $m=14$). The dataset can be tabulated to an $n \times m$ matrix with each row reflecting the expression of a single gene in 14 different time points (14-arrays) and each column showing the expression of n -genes in a single array (timepoint). SVD transforms this dataset from an $n \times m$ space to a reduced l -eigengenes \times l -eigenarrays subspace where $l = [\min m, n]$ (Fig. 6.1). The diagonals in the $l \times l$ matrix ϵ are eigenvalues here called “eigenexpression levels” $[\epsilon_l]$ which can be used to calculate “fractions of eigenexpression” for the l th eigenvalue from the equation below:

$$p_l = \frac{\epsilon_l^2}{\sum_{k=1}^l \epsilon_k^2} \tag{6.3}$$

Alter et al. used fractions of eigenexpression as a mean to infer the significance of eigengenes and their corresponding eigenarrays (singular vectors). Once the significance of singular values (SVs) was determined, the relationship between these mathematical concepts and biological processes or cellular states, in this case cell cycle, were investigated. To this end, the authors visualized individual singular values by plotting the expression level of each eigengene over time. Since the authors were interested in gene programs involved in a specific cellular state, they filtered out the first singular vector because it followed a steady state expression pattern. The next three SVs showed biologically meaningful oscillations during cell cycle. The oscillations of the second and fourth SVs at early time points corresponded to a transient response to elutriation. Thus SVD naturally decomposed the dynamical patterns of gene expression in the yeast cell cycle.

6.3 Principal Components Analysis

Following SVD, principal components analysis (PCA) can be used to compress a dataset to relevant measurements that approximate the data. Both computational and visual analysis is often hard to do in higher order datasets as each measurement (observation) constitutes its own dimension in space and the value of each sample (variable) constitutes a point in each of these dimensions. By transforming the data using PCA, we can identify important relationships in the data.

First, eigenvalues and eigenvectors are derived from the data covariance matrix [22, 29]. These eigenvectors make up an orthogonal basis set, or set of linearly independent vectors that, when combined, can describe the data. The eigenvectors paired with the smallest eigenvalues are eliminated to yield a compressed basis set. This basis set of eigenvectors is then used to generate a transformed data matrix, the dimensions of which are called latent dimensions or principle components (PCs) [22, 29, 45]. A principal component is by analogy a singular vector in SVD.

A latent dimension is a new dimension created to capture the majority of information in multiple the original dimensions [40]. Mathematically, a principle component is a linear combination of the original data dimensions, weights for which are determined by the magnitude of the eigenvector corresponding to that principle component [26, 29]. The eigenvector paired with the largest eigenvalue defines the first principal component and captures the greatest amount of variance in the data [29, 45]. In this component, the original dimensions with the most variance in variable data will have the largest weighting. Because the PCs are orthogonal, the second principal component will point in a direction perpendicular to the first component and capture the majority of the leftover variance. This iteration continues for all subsequent PCs. Thus, the transformed dataset is usually only made up of a handful of latent dimensions because they can capture the majority of the data variance eliminating any statistical noise from subsequent PCs. This filtering makes relevant relationships between samples more readily apparent.

Further, one can create predictive models with latent dimensions by searching for relationships between PCs using principal components regression (PCR). This method utilizes established regression techniques to find the relationship between several variables (predictor variables) and dependent variables not included in the predictors [26]. PCR uses the first few principal components to simplify the analysis of many variables to linear or multilinear regression between the components (predictors) and the desired measurements [26, 29]. Resulting coefficients of the PCs, fitted using least-squares approaches, can be decomposed to regression coefficients of each of the original variables in the component. The variable with the largest magnitude coefficient is the most correlated to the desired dependent variable while the sign of the coefficient indicates positive or negative correlation [26]. In this way, decomposition by PCR can be used to extract relationships between different variables in the dataset.

Thus, PCA and PCR can be used not only to generate hypotheses about sample relationships but also to generate data-driven predictions. PCA is often used to

analyze DNA (or cDNA) microarrays by clustering observational data to expose relevant coregulations of genes or relevant similarities or disparities between cellular samples such as different cancer tumors [12, 20, 42].

6.3.1 Mathematical Framework

6.3.1.1 Principal Component Analysis (PCA)

First, the dataset should be mean-centered so that the mean of each variable across all observations is zero. This adjustment greatly simplifies the covariance matrix calculation as well as eigenvector determination. For centering, the means of each variable (column) should be subtracted from each observation of that variable (row) in an element-wise manner as shown in Eq. (6.4).

$$\begin{bmatrix} M_{1,1} & M_{1,2} \\ M_{2,1} & M_{2,2} \end{bmatrix} - \begin{bmatrix} \overline{M_1} & \overline{M_2} \\ \overline{M_1} & \overline{M_2} \end{bmatrix} = \begin{bmatrix} A_{1,1} & A_{1,2} \\ A_{2,1} & A_{2,2} \end{bmatrix} \quad (6.4)$$

Here M is a representative 2×2 data matrix and another matrix of the same size, containing the means of each sample (columns), is subtracted to generate the adjusted data matrix A . Notice that the mean of all columns should now be zero.

Now the covariance matrix of the dataset can be found from A . While in SVD the original dataset was used for decomposition, in PCA the chief interest is in the covariance of the data not the absolute magnitude [40]. Thus, the sample covariance matrix is used for decomposition as shown in Eq. (6.5).

$$C = \begin{bmatrix} \text{cov}(1, 1) & \dots & \text{cov}(1, N) \\ \vdots & \ddots & \vdots \\ \text{cov}(M, 1) & \dots & \text{cov}(M, N) \end{bmatrix} = \frac{1}{N-1} \sum_{i=1}^N (A_i - \overline{A})(A_i - \overline{A})^T \quad (6.5)$$

which simplifies to $C = AA^T/(N-1)$.

Here C is a symmetric sample covariance matrix, where the elements of the matrix are the covariances of each observation (row) with every other variable dimension, M denotes the number of observations, and N is the number of variables (columns) in A . Because A is mean-centered ($\overline{A} = 0$), this equation simplifies to $AA^T/(N-1)$. Using this notation, we can find the eigenvectors of C by decomposing it into a diagonal matrix D [47].

We can rewrite Eq. (6.2) as,

$$C = VDVT^T \quad (6.6)$$

such that the columns of V are the eigenvectors of A which correspond to the eigenvalues in the diagonal matrix D . Here, eigenvalues correspond to the contribution of that eigenvector to the reconstruction of C from the decomposition. For the covariance matrix, small eigenvalues correspond to eigenvectors that contain a small amount of

the variance in the data. Thus, columns corresponding to low-magnitude eigenvalues can be eliminated from V to yield a compressed eigenvector matrix (B) that will make up the basis set of the data A [47].

Multiplying the compressed eigenvector matrix B with A transforms the adjusted data into principle component space as given by Eq. (6.7).

$$P = B^T A \quad (6.7)$$

Here P is the approximated data matrix where the rows correspond to latent dimensions or principle components and the columns correspond to samples. The elements of the matrix are the values of samples in each component. As previously mentioned, the eigenvectors are ordered from greatest corresponding eigenvalue to smallest. Therefore, the first principle component (first eigenvector) accounts for the most variance in the data. If P is composed of three or fewer principal components, the sample values can be plotted in a 2D or 3D fashion to group covarying samples.

Implementation of this method and example plots will be discussed later in the context of host-cell responses to viral infection [27].

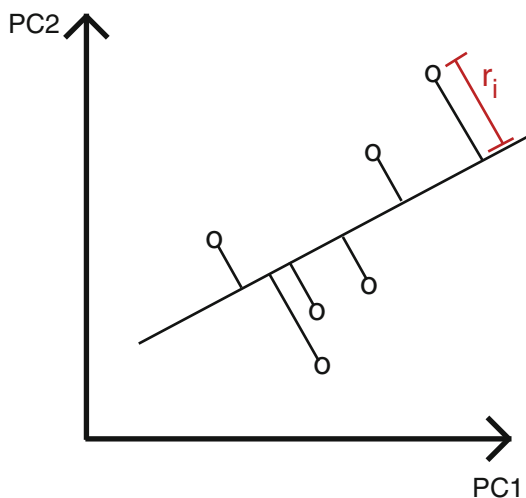
6.3.1.2 Principal Component Regression (PCR) Using Total Least Squares

After the principal components have been defined, there may be instances in which knowing the relationship between principle components or principle components and an independent observation dimension are useful. Linear or planar orthogonal regression techniques can be used to determine these relationships [26, 29, 45]. In this section we focus on total least squares regression (TLSR).

As opposed to ordinary least squares regression, TLSR aims to minimize the perpendicular residual error from the regression fit [45]. This is an important distinction as it implies variance or measurement error in all the dimensions. Measurement inaccuracies create uncertainty or associated variance in the position of each data point in principal component space. Therefore, regression models should take this into account when minimizing residual error to create an unbiased fit. For PCA, all the observation dimensions used to create latent dimensions are subject to measurement error or variance [45].

First, appropriate PCs must be chosen as predictor variables. In most cases choosing the first one or two principle components is the most relevant [26, 29]. However this is not always the case and a more in-depth discussion of choosing appropriate PCs can be found in Jolliffe [29].

Fig. 6.2 TSLR Orthogonal Residual Schematic. TSLR uses orthogonal residuals (red— r_i) to fit a regression line to data (o) displayed in PC space. For color figure please refer to online version



Once predictor variables have been chosen, iterative computational optimization algorithms, in environments like Matlab, can be used to identify the best-fit line or plane. In general these computational methods attempt to minimize Eq. (6.8) (Fig. 6.2).

$$E = \sum_{i=1}^N |r_i^2| \quad (6.8)$$

where E is the residual error and r_i is the orthogonal distance of a data point (o) from the regression. The schematic Fig. 6.2 illustrates the orthogonal distance (r_i) of a representative data point (o) from the linear regression line.

While PCA is an unsupervised decomposition method that does not take into account the inherent variance between variables, it can extract valuable information from multidimensional datasets. This information can be used to generate simplified regression models by using the principal components themselves as predictors rather than the original observations.

6.3.2 Application of PCA: Decomposition of Experimental Data

In this section we will use data published in [27] quantifying host-cell viral infection responses to demonstrate the utility of multidimensional decomposition in systems modeling.

6.3.2.1 Biological Introduction

Viral myocarditis including heart failure and dilated cardiomyopathy are most commonly attributed to coxsackievirus B3 (CVB3) infection of cardiomyocytes. Acute infections cause host-cell damage and viral progeny release. These progeny then infect neighboring cells and advance tissue damage until the immune system can clear the infection.

CVB3 induced myocarditis is the most common cause of heart failure in young adults, children, and immunocompromised patients. The only late-stage treatment available in these cases is heart transplantation suggesting that once the virus has begun to successfully replicate in host-cells, there is little that can be done other than remove the diseased tissue. Thus, early-stage treatments remain the best hope for reducing tissue damage.

However, no successful treatment currently exists because initial host-cell responses to infection are poorly understood. CVB3 must interact with several host-cell intracellular components to replicate successfully [54]. Such interactions include degrading certain machinery while taking over others, suggesting that numerous cellular signaling pathways are disrupted [15]. Treatments developed to recover normal function of these pathways before viral progeny are released remain the best hope for attenuation of disease severity.

6.3.2.2 Experimental Measurements and Generated Data

Because host-cell signaling consists of highly connected networks of proteins, perturbations in these pathways must be studied simultaneously to truly understand, at a systems level, how CVB3 affects the host. To accomplish this, Jensen et al. [27] surveyed several kinase activities in response to viral infection. The activities of eight canonical stress and inflammatory response kinases were quantified by phosphoprotein abundance in a phospho (p)-ELISA format. Each of these measurements was made in response to five different CVB3 viral doses (MOI) at several time points. In addition, six host-cell responses were quantified in a parallel experimental set-up. The resulting datasets form tensor structures (three-dimensional matrices) as shown in Fig. 6.3.

To draw parallels between the temporal patterns of kinase activation and cellular output, the data can be plotted in several 2D plots and then visually compared to each other. An example of this method is displayed in Fig. 6.4.

From visual inspection we can see that the data does not have correlated or simple activation patterns with respect to time or MOI. Some of the kinases show an increase then decrease in activation with time like Hsp27, ATF2, AKT, GSK3 and p38. Others have more complex bimodal temporal patterns like CREB, ERK, and $I\kappa B\alpha$. Further, some kinases show increased activation with respect to MOI that decreases after a certain MOI threshold is reached (Akt and $I\kappa B\alpha$, Fig. 6.4) while others show less generalizable patterns. Additionally, while the cellular outputs all show increases in expression (VP1), activation (casp-8, casp-9, casp-3), or abundance (RVP, Death)

with respect to both MOI and time, their relationship to the kinases cannot be inferred. Taken as a system, there is no simple way to explain the relationships between kinases, cellular outputs, time, and MOI using this type of data processing.

6.3.2.3 Tensor Decomposition

In such cases, tensor decomposition can illuminate important correlations in the data. Higher order data structures can be decomposed and modeled with either of two generalized SVD and PCA methods. Tucker decomposition [52] and CANDCOMP/PARAFAC (CP) [10, 32] are established methods that were first developed in the fields of psychometrics and chemometrics.

CP expresses an N-mode data tensor as a sum of rank one tensors, while Tucker decomposition factorizes the data into a core tensor and N-corresponding matrices [8, 32]. Thus, for a three mode tensor, CP factorizes the data into a sum of rank one tensors expressed as the outer product (\otimes) of vectors from three matrices (A, B, and C) (Eq. 6.9, Fig. 6.5a). This operation is also known as the Khatri-Rao product of A, B, and C. Tucker decomposition yields one core tensor W) and three matrices (A, B, and C) (Eq. 6.10, Fig. 6.5b). In both cases, the matrices A, B, and C correspond to mode 1, 2, and 3 respectively and contain the loadings or eigenvectors that define each principal component within that mode.

$$X_{P \times Q \times R} = \sum_{n=1}^N a_n \otimes b_n \otimes c_n \tag{6.9}$$

where N is the number of components or rank one tensors being used to approximate the data tensor X and a_n , b_n , and c_n are the vectors corresponding to the n th component.

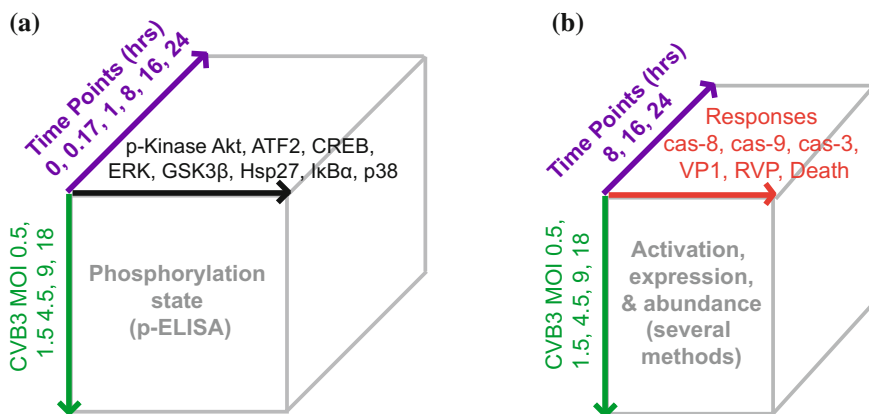


Fig. 6.3 Tensor structure of host-cell response data. Tensor data structure of phosphoprotein ELISA signal (a) and cell response data (b)




Fig. 6.4 Visualization of host-cell response data. Heat map visualization of phosphoprotein signals (a) and cell response data (b) given 5 different viral doses (MOIs) measured at 6 (a) or 3 (b) time points post-infection. Data are reprinted from Jensen et al. [27]

$$X_{P \times Q \times R} = \sum_{p=1}^P \sum_{q=1}^Q \sum_{r=1}^R g_{pqr} a_p \otimes b_q \otimes c_r \quad (6.10)$$

where g_{pqr} is an element in the core tensor G and a_p , b_q , and c_r are column vectors from the matrices A , B , and C .

It is important to note these generalized methods differ from PCA in that they usually result in unique solutions. In two-way PCA, the PCs can be rotated in space and still maintain the amount of variance captured by changing the scaling within the PCs. This is not true for higher order PCA as rotation will cause a loss of fit [8]. Consequently, CP and Tucker models differing in the number of components calculated must be done iteratively; each component cannot be found serially. In this way, CP and Tucker decomposition expand on the concepts of bilinear PCA to model important covariations in the data while maintaining the information-rich tensor structure.

To apply these methods to our CVB3 host-cell kinase signaling dataset (Fig. 6.3a) we must first mean center and standardize our tensor [8]. We want to center along the first mode or the observation mode, which, in this case, is the MOI dimension. It follows then that our variables are the different kinases (mode 2) through time (mode 3). Just as in bilinear PCA mean centering along the first mode gives us a zero

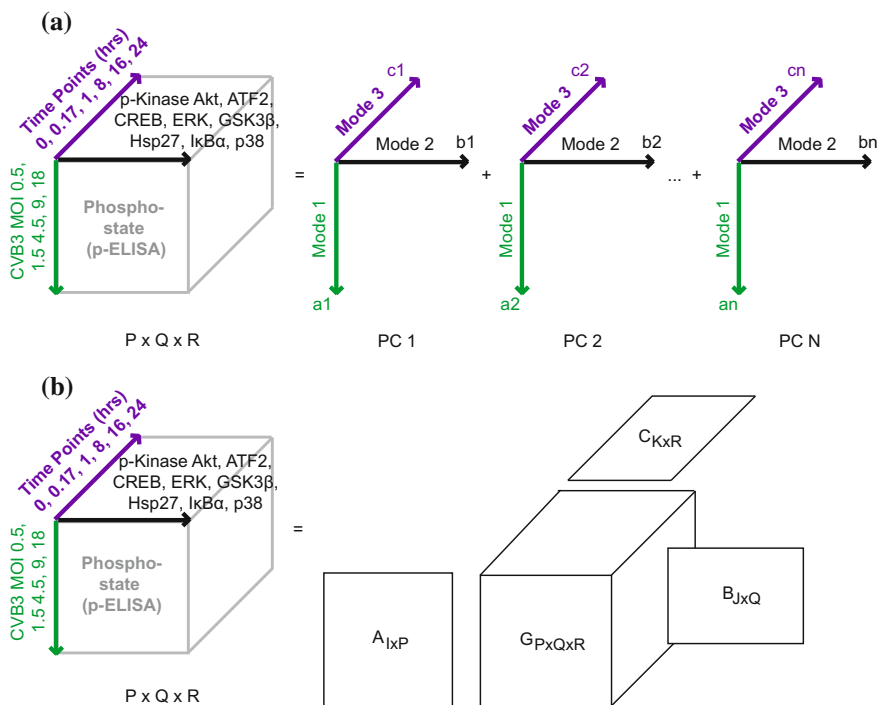


Fig. 6.5 Schematic of Tensor Decomposition Techniques. Schematic of CP (a) and Tucker (b) decomposition using experimental phosphoprotein data structure from Fig. 6.4. The data matrix and core matrix G have dimensions $P \times Q \times R$, PC = principal component, lowercase letters a, b, c denote vectors, and uppercase letters A, B, C denote matrices. For color figure please refer to online version

offset to calculate the variances in phosphosignal of each phosphoprotein. Next, standardization, or scaling by the standard deviation, in modes 2 (phosphoprotein) and 3 (time) allows us to compare these variances between phosphoproteins, time points, and MOIs.

Now we can decompose this adjusted tensor using the PARAFAC (CP) algorithm. Unique solutions or sets of principal components (A, B, and C) are fit using an alternating least squares (ALS) approach [8]. In short, this algorithm uses the adjusted data and an initial guess for B and C (mode 2 and 3 principal components respectively) to fit A by least squares regression. This is then repeated for each matrix A, B, and C until a convergence condition is reached.

As mentioned before, these steps have to be repeated for models with differing numbers of components. To determine the optimal number of components for the model, several different numbers of components can be surveyed, each followed by inspection of certain fit and stability metrics. These metrics include percent variance of the original dataset captured, sum-of-squared residuals between the model and original data, crossvalidation or leave-one-out stability analysis, and concordance




Fig. 6.6 Visualization of phosphoprotein activity in principal component space. Phosphoprotein (mode 2, *black*) and time post-infection (mode 3, *purple*) loadings are mapped in principal component space to show clusters with similar loadings in each PC. Loadings and PCs were generated using PARAFAC analysis. For color figure please refer to online version

with the Tucker core tensor [40]. The last of these utilizes the idea that PARAFAC will yield the same answer as Tucker decomposition with a core tensor of ones on the superdiagonal and zeros elsewhere [9]. Thus, 100% concordance means the A, B, and C matrices are the best possible fit of the data and do not need to be scaled by a non-superidentical core matrix.

After mean centering and standardizing the host-cell phosphoprotein dataset, a two compartment PARAFAC model can be generated which captures 80% of the variance in the original data. Two components were considered optimal as larger models did not meet concordance requirements and smaller models did not capture as much variance. To interpret this simplified model, we can visualize both the kinase and time loadings in PC1 and PC2 (Fig. 6.6). The magnitudes of the loadings indicate the amount it contributes to the respective principal component. The greater the magnitude of the loading, the greater contribution that variable has to the model generated from the Khatri-Rao product of the principal components.

Thus, when interpreting Fig. 6.6, we want to focus on clusters that lie far from zero on at least one PC, because those signaling proteins and times that lie close to the origin have little to no influence on the model. As labeled in the figure, the

PARAFAC model exposes two interesting clusters: (1) ERK and time 0.17h and (2) p38 and time 24h. This suggests that p38 signaling is responsible for most the variance associated with MOI at later time points while ERK signaling is responsible for most of the variance associated with MOIs at early time points. Biologically, we can now hypothesize that the ERK and p38 pathways are relaying the most information about CVB3 mediated viral infection at early and late times respectively [27]. How the viral genome interacts with these pathways would still need to be studied in carefully designed experiments that isolate the role of these pathways in vitro or in vivo. However, this demonstrates the ability of higher order decomposition techniques to uncover new hypothesis that lead to more targeted systems-biology studies.

6.3.3 *Partial Least Squares Regression (PLSR)*

As mentioned in the previous section with data matrices, PCA defines principal components that are optimized to capture the overall variance in the data matrix A . However, this does not mean that the resulting principal components are optimally interpretable, nor that they are the best regressors for predicting another data matrix. In such circumstances, it is preferred to rotate the leading principal components [28], which is easily achieved in two dimensions with the following linear operator:

$$\begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \quad (6.11)$$

(Similar operators can be defined for rotations in three dimensions.) A key point is that this “subspace rotation” does not affect the overall variance captured by the PCA model, because the solution is rotationally degenerate. Rather, it rebalances the variance among the retained principal components. Subspace rotation is commonly employed when building statistical models of biological processes [5, 22].

For statistical modeling of signal transduction, PLSR has proved widely useful and informative. Successful models have been built to link signaling to cell death [22, 34, 38, 39, 41], cell-cycle progression [39, 51], proliferation [7, 36, 37], and cytokine secretion [30, 34, 36]. More-recent theoretical work has suggested that, because of the fundamental chemical-reaction kinetics of biochemical networks, PLSR is virtually guaranteed to reduce a signaling circuit down to a handful of principal components for follow-on analysis [13]. Of course, there are caveats about framing a proper $X \rightarrow Y$ hypothesis [25], but it is reassuring to know that the approach is fundamentally sound and highly versatile. Consequently, PLSR has entered into the standard curricula for many systems-biology courses [33].

6.3.4 Mathematical Framework

For regression modeling within high-dimensional datasets, there is a more effective way of identifying correlated principal components than PCR followed by subspace rotation. In partial least squares regression (PLSR), principal components are identified numerically that maximize the covariance between an independent data matrix (X) and a dependent data matrix (Y). (Note the distinction from PCA, which simply maximizes capture of the overall variance of a single data matrix.) Computationally, PLSR arrives at a covariance model by jointly factorizing X and Y as follows:

$$X = TP^T \tag{6.12}$$

$$Y = UQ^T \tag{6.13}$$

where T and U are scores vectors and P and Q are loading matrices.

The regression between X and Y is linear between the “scores vectors” of the independent and dependent matrices:

$$U = TB \tag{6.14}$$

$$\text{Thus } Y = TBQ^T \tag{6.15}$$

The simplest protocol for building a PLSR model is by using the nonlinear iterative partial least squares (NIPALS) algorithm. In this algorithm, a row from Y is randomly chosen as the first guess for a scores vector (u), and then X is projected onto u to define the first guess at a “loadings vector” p . Here, the exchange of scores vectors (using u with X and t with Y) is critical for linking the two matrices together and building a PLSR model that maximizes the covariance between X and Y [17]. The first iteration of the loadings vector is then normalized and projected onto X to define a provisional t , which is subsequently projected onto Y to calculate the first iteration of its loadings vector, q . This loadings vector is normalized to unit length as done previously for p , and then the normalized q is projected onto Y to define the second iteration of u . This process continues until u converges to a fixed value within a specified tolerance. Software for building PLSR models is readily available in MATLAB, R, as well as independent commercial platforms [25].

6.3.5 Application of PLSR: Modeling Tensor Data Sets

In the previously discussed example from [27], two datasets were generated, a phosphoprotein signal tensor and a cell response tensor (Fig. 6.3). Unlike PCA, PLSR is a powerful technique that can help identify relationships between these two tensors and point to kinases that dictate viral infection responses in cardiomyocytes.

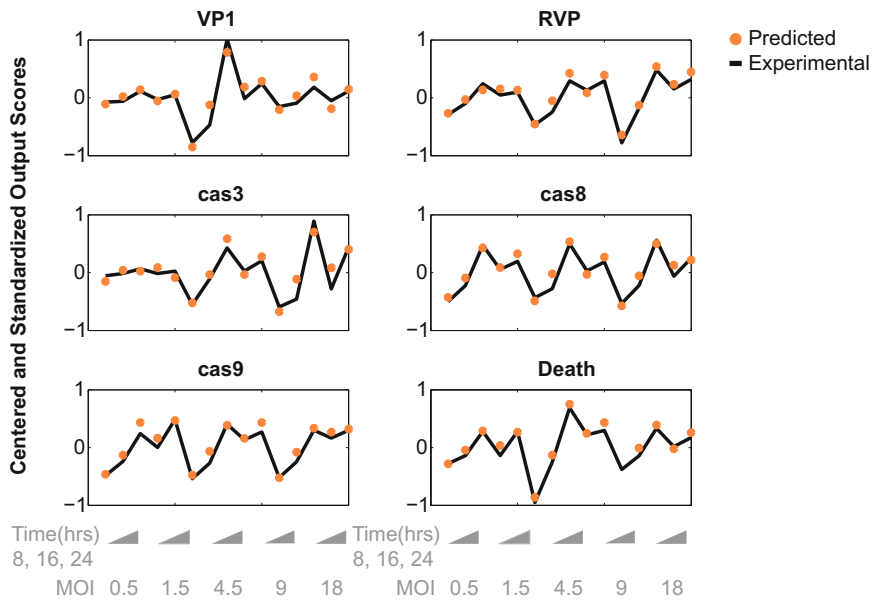


Fig. 6.7 PLSR model of experimental host-cell response data. All six of the cellular response measurements (*black*) can be successfully predicted by a four component PLSR model (*orange*). (*VP1* viral capsid protein 1, *RVP* release of viral progeny, *cas3* caspase 3, *cas8* caspase 8, *cas9* caspase 9, *Death* cellular death via MTS assay). For color figure please refer to online version

As mentioned previously, PLSR generates principal components that capture the maximal variances in both an input tensor and an output tensor. In this way, PLSR is a supervised decomposition and regression method that can be used as a predictive model. In this example, the predictor dataset is the signaling activity tensor (Fig. 6.6a) and the dependent data is the cellular response tensor (Fig. 6.3b). Note that the dimensions of the tensor datasets do not have to be equal except in mode 1 (observations). Before attempting PLSR, the tensor must be centered and standardized as described for the PARAFAC model.

Once each dataset is adjusted, a crossvalidated PLSR model can be developed by using a NIPALS algorithm to solve Eqs. 6.12 and 6.13 [6]. These equations are essentially the same as those for the PARAFAC model; the difference is that NIPALS solves them simultaneously, until a convergence criterion is met.

The four-component solution captures 89 % of the variance in the predictor tensor and 91 % variance in the output tensor. The resulting model can accurately predict all cellular outputs as shown in Fig. 6.7. To visualize the significant players in this multi-component model, we must choose which principal component combinations to look at. In general, we want to look at PCs that capture the most variance in the data, which are usually PC1 and PC2. However, in this case, PC3 and PC4 captured substantial amounts of non-redundant variance in the data. By analyzing Fig. 6.8 in a similar fashion to Fig. 6.6, we find that p38, Hsp27, CREB, GSK3B, and Akt cluster




Fig. 6.8 Visualization of phosphoprotein and cellular output data in principal component space. Phosphoprotein (mode 2 of predictor tensor, *black*) and cellular output (mode 2 of output tensor, *red*) loadings are mapped in principal component space to show clusters with similar loadings in each PC. Loadings and PCs were generated using a four component PLS model. For color figure please refer to online version

with four of the six cellular outputs (caspase 3 activation, viral progeny release, VP1 expression, and cell death).

We can then hypothesize that these signaling proteins are influential in deciding cellular responses to CVB3 infection and could be potential treatment targets to attenuate subsequent tissue damage. By utilizing the powerful extraction capabilities of tensor decomposition and PLSR modeling we have effectively eliminated half the potential targets we started with. We could eliminate more by looking at other PCs and isolating those that correlate with early time points. This example illustrates the utility of different tensor decomposition and modeling techniques that result in rational experiment design invaluable to systems biologists struggling to understand and manipulate immensely complex cellular networks.

6.4 Concluding Remarks

Will the age-old methods of SVD, PCA, and PLSR become obsolete in the midst of this new era of biomedical Big Data? We have many reasons to think not. First, the statistical models introduced in this chapter are among the simplest linear methods for

reducing complex datasets. Invoking Occam's razor, these models should be proved to be insufficient before more-complicated alternatives are sought. Simple models are more-easily interpretable—by generating principal components that can be immediately mapped back onto the primary data, the models stay grounded in what they were derived from. Statistical models therefore avoid the pitfalls of machine-learning approaches, such as support-vector machines and neural networks, which can make remarkable predictions but leave the user confused about how the predictions were made [35, 44].

Second, the iterative methods for PCA and PLSR are very scalable to large datasets, because they do not require calculating the covariance matrix as with older implementations of SVD. To circumvent problems interpreting PCs with many variables, there are sparse variants that use different constraints to generate PCs with more-limited mixtures of the original variables [56]. The computational simplicity allows model stability be easily assessed by standard numerical methods such as cross validation, bootstrapping, and permutation [14]. Statistical models thus provide a direct indication of their uncertainty based upon the data provided, which stands in contrast to physico-chemical models where uncertainty can be more cryptic [18].

Last and most important, even if the methods here are eventually supplanted by superior alternatives, they nonetheless serve as an important entry point for those interested in statistical modeling. Biologists are already accustomed to looking at their results—statistical models provide a more-formal way of inspecting complex data and illustrating the power of computation in real terms [24]. Just as it is difficult to imagine life now without a computer or a smartphone, biological research will soon become unfathomable without the aid of statistical models.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

1. Albeck, J.G., Burke, J.M., Spencer, S.L., Lauffenburger, D.A., Sorger, P.K.: Modeling a snap-action, variable-delay switch controlling extrinsic cell death. *PLoS Biol.* **6**(12), 2831–2852 (2008). doi:[10.1371/journal.pbio.0060299](https://doi.org/10.1371/journal.pbio.0060299) (07-PLBI-RA-3412 [pii])
2. Albeck, J.G., MacBeath, G., White, F.M., Sorger, P.K., Lauffenburger, D.A., Gaudet, S.: Collecting and organizing systematic sets of protein data. *Nat. Rev. Mol. Cell Biol.* **7**(11), 803–812 (2006)
3. Aldridge, B.B., Saez-Rodriguez, J., Muhlich, J.L., Sorger, P.K., Lauffenburger, D.A.: Fuzzy logic analysis of kinase pathway crosstalk in TNF/EGF/insulin-induced signaling. *PLoS Comput. Biol.* **5**(4), e1000340 (2009). doi:[10.1371/journal.pcbi.1000340](https://doi.org/10.1371/journal.pcbi.1000340)
4. Alter, O., Brown, P.O., Botstein, D.: Singular value decomposition for genome-wide expression data processing and modeling. *Proc. Natl. Acad. Sci. U.S.A.* **97**(18), 10101–10106 (2000)

5. Alter, O.: Genomic signal processing: from matrix algebra to genetic networks. *Methods Mol. Biol.* **377**, 17–60 (2007). doi:[10.1007/978-1-59745-390-5_2](https://doi.org/10.1007/978-1-59745-390-5_2)
6. Andersson, C.A., Bro, R.: The N-way toolbox for MATLAB. *Chemometr. Intell. Lab. Syst.* **52**(1), 1–4 (2000). doi:[10.1016/S0169-7439\(00\)00071-X](https://doi.org/10.1016/S0169-7439(00)00071-X)
7. Beyer, E.M., MacBeath, G.: Cross-talk between receptor tyrosine kinase and tumor necrosis factor- α signaling networks regulates apoptosis but not proliferation. [Research Support, N.I.H., Extramural]. *Mol. Cell. Proteomics* **11**(6), M111.013292. doi:[10.1074/mcp.M111.013292](https://doi.org/10.1074/mcp.M111.013292)
8. Bro, R.: PARAFAC. Tutorial and applications. *Chemometr. Intell. Lab. Syst.* **38**, 149–171 (1997)
9. Bro, R., Kiers, H.A.L.: A new efficient method for determining the number of components in PARAFAC models. *J. Chemometr.* **17**, 274–286 (2003)
10. Carroll, J.D., Chang, J.J.: Analysis of individual differences in multidimensional scaling via an N-way generalization of ‘Eckart-Young’ decomposition. *Psychometrika* **35**, 283–319 (1970)
11. Chen, W.W., Schoeberl, B., Jasper, P.J., Niepel, M., Nielsen, U.B., Lauffenburger, D.A., Sorger, P.K.: Input-output behavior of ErbB signaling pathways as revealed by a mass action model trained against dynamic data. *Mol. Syst. Biol.* **5**, 239 (2009)
12. Crescenzi, M., Giuliani, A.: The main biological determinants of tumor line taxonomy elucidated by a principal component analysis of microarray data. *FEBS Lett.* **507**(1), 114–118 (2001)
13. Dworkin, M., Mukherjee, S., Jayaprakash, C., Das, J.: Dramatic reduction of dimensionality in large biochemical networks owing to strong pair correlations [Research Support, N.I.H., Extramural Research Support, Non-U.S. Gov’t]. *J. R. Soc. Interface* **9**(73), 1824–1835 (2012). doi:[10.1098/rsif.2011.0896](https://doi.org/10.1098/rsif.2011.0896)
14. Efron, B., Tibshirani, R.J.: *An Introduction to the Bootstrap*. Chapman and Hall, London (1993)
15. Esfandiari, M., McManus, B.M.: Molecular biology and pathogenesis of viral myocarditis. *Annu. Rev. Pathol.* **3**, 127–155 (2008). doi:[10.1146/annurev.pathmechdis.3.121806.151534](https://doi.org/10.1146/annurev.pathmechdis.3.121806.151534)
16. Gaudet, S., Janes, K.A., Albeck, J.G., Pace, E.A., Lauffenburger, D.A., Sorger, P.K.: A compendium of signals and responses triggered by prodeath and prosurvival cytokines. *Mol. Cell Proteomics* **4**(10), 1569–1590 (2005)
17. Geladi, P., Kowalski, B.R.: Partial least-squares regression—a tutorial. *Anal. Chim. Acta.* **185**, 1–17 (1986)
18. Gutenkunst, R.N., Waterfall, J.J., Casey, F.P., Brown, K.S., Myers, C.R., Sethna, J.P.: Universally sloppy parameter sensitivities in systems biology models [Research Support, N.I.H., Extramural Research Support, U.S. Gov’t, Non-P.H.S.]. *PLoS Comput. Biol.* **3**(10), 1871–1878 (2007). doi:[10.1371/journal.pcbi.0030189](https://doi.org/10.1371/journal.pcbi.0030189)
19. Hoffmann, A., Levchenko, A., Scott, M.L., Baltimore, D.: The IkappaB-NF-kappaB signaling module: temporal control and selective gene activation. *Science* **298**(5596), 1241–1245 (2002)
20. Hu, J., Locasale, J.W., Bielas, J.H., O’Sullivan, J., Sheahan, K., Cantley, L.C., Vitkup, D.: Heterogeneity of tumor-induced gene expression changes in the human metabolic network. *Nat. Biotechnol.* **31**(6), 522–529 (2013). doi:[10.1038/nbt.2530](https://doi.org/10.1038/nbt.2530)
21. Hug, S., Schmidl, D., Li, W.B., Greiter, M.B., Theis, F.J.: Bayesian model selection methods and their application to biological ODE systems. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
22. Janes, K.A., Albeck, J.G., Gaudet, S., Sorger, P.K., Lauffenburger, D.A., Yaffe, M.B.: A systems model of signaling identifies a molecular basis set for cytokine-induced apoptosis. *Science* **310**(5754), 1646–1653 (2005)
23. Janes, K.A., Lauffenburger, D.A.: A biological approach to computational models of proteomic networks. *Curr. Opin. Chem. Biol.* **10**(1), 73–80 (2006)
24. Janes, K.A., Lauffenburger, D.A.: Models of signalling networks—what cell biologists can gain from them and give to them [Research Support, N.I.H., Extramural Research Support, Non-U.S. Gov’t]. *J. Cell. Sci.* **126**(Pt 9), 1913–1921 (2013). doi:[10.1242/jcs.112045](https://doi.org/10.1242/jcs.112045)
25. Janes, K.A., Yaffe, M.B.: Data-driven modelling of signal-transduction networks. *Nat. Rev. Mol. Cell. Biol.* **7**(11), 820–828 (2006)

26. Jeffers, J.: Two case studies in the application of principal component analysis. *J. Roy. Stat. Soc.* **16**(3), 225–236 (1967)
27. Jensen, K.J., Garmaroudi, F.S., Zhang, J., Lin, J., Boroomand, S., Zhang, M., Janes, K.A.: An ERK-p38 subnetwork coordinates host cell apoptosis and necrosis during coxsackievirus B3 infection [Research Support, N.I.H., Extramural Research Support, Non-U.S. Gov't]. *Cell Host Microbe* **13**(1), 67–76 (2013). doi:[10.1016/j.chom.2012.11.009](https://doi.org/10.1016/j.chom.2012.11.009)
28. Jensen, K.J., Janes, K.A.: Modeling the latent dimensions of multivariate signaling datasets. *Phys. Biol.* **9**(4), 045004 (2012). doi:[10.1088/1478-3975/9/4/045004](https://doi.org/10.1088/1478-3975/9/4/045004)
29. Jolliffe, I.T.: *Principal Component Analysis*. Springer Series in Statistic
30. Kemp, M.L., Wille, L., Lewis, C.L., Nicholson, L.B., Lauffenburger, D.A.: Quantitative network signal combinations downstream of TCR activation can predict IL-2 production response. *J. Immunol.* **178**(8), 4984–4992 (2007)
31. Kirk, P., Silk, D., Stumpf, M.P.H.: *Reverse engineering under uncertainty*, In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
32. Kolda T.G., Bader, B.B.: *Tensor Decompositions and Applications*. Sandia National Laboratories Report, SAND2007-6702 (2007)
33. Kreeger, P.K.: Using partial least squares regression to analyze cellular response data. *Sci. Signal* **6**(271), tr7 (2013). doi:[10.1126/scisignal.2003849](https://doi.org/10.1126/scisignal.2003849)
34. Kreeger, P.K., Mandhana, R., Alford, S.K., Haigis, K.M., Lauffenburger, D.A.: RAS mutations affect tumor necrosis factor-induced apoptosis in colon carcinoma cells via ERK-modulatory negative and positive feedback circuits along with non-ERK pathway effects. *Cancer Res.* **69**(20), 8191–8199 (2009). doi:[10.1158/0008-5472.CAN-09-1921](https://doi.org/10.1158/0008-5472.CAN-09-1921) (0008-5472.CAN-09-1921 [pii])
35. Krogh, A.: What are artificial neural networks? *Nat. Biotechnol.* **26**(2), 195–197 (2008)
36. Kumar, D., Srikanth, R., Ahlfors, H., Lahesmaa, R., Rao, K.V.: Capturing cell-fate decisions from the molecular signatures of a receptor-dependent signaling response. *Mol. Syst. Biol.* **3**, 150 (2007)
37. Kumar, N., Wolf-Yadlin, A., White, F.M., Lauffenburger, D.A.: Modeling HER2 effects on cell behavior from mass spectrometry phosphotyrosine data. *PLoS Comput. Biol.* **3**(1), e4 (2007)
38. Lau, K.S., Juchheim, A.M., Cavaliere, K.R., Philips, S.R., Lauffenburger, D.A., Haigis, K.M.: In vivo systems analysis identifies spatial and temporal aspects of the modulation of TNF-alpha-induced apoptosis and proliferation by MAPKs. *Sci. Signal* **4**(165), ra16 (2011). doi:[10.1126/scisignal.2001338](https://doi.org/10.1126/scisignal.2001338) (4/165/ra16 [pii])
39. Lee, M.J., Ye, A.S., Gardino, A.K., Heijink, A.M., Sorger, P.K., Macbeath, G., Yaffe, M.B.: Sequential application of anticancer drugs enhances cell death by rewiring apoptotic signaling networks. *Cell* **149**(4), 780–794 (2012). doi:[10.1016/j.cell.2012.03.031](https://doi.org/10.1016/j.cell.2012.03.031)
40. Martens, H., Martens, M.: *Analysis of one data table X: principal component analysis multivariate analysis of quality: an introduction*, 1st edn, pp. 93–110. Wiley, New York (2001)
41. Miller-Jensen, K., Janes, K.A., Brugge, J.S., Lauffenburger, D.A.: Common effector processing mediates cell-specific responses to stimuli. *Nature* **448**(7153), 604–608 (2007)
42. Misra, J., Schmitt, W., Hwang, D., Hsiao, L.L., Gullans, S., Stephanopoulos, G., Stephanopoulos, G.: Interactive exploration of microarray gene expression patterns in a reduced dimensional space. *Genome. Res.* **12**(7), 1112–1120 (2002). doi:[10.1101/gr.225302](https://doi.org/10.1101/gr.225302)
43. Nakakuki, T., Birtwistle, M.R., Saeki, Y., Yumoto, N., Ide, K., Nagashima, T., Kholodenko, B. N.: Ligand-specific c-Fos expression emerges from the spatiotemporal control of ErbB network dynamics [Research Support, N.I.H., Extramural Research Support, Non-U.S. Gov't]. *Cell* **141**(5), 884–896 (2010). doi:[10.1016/j.cell.2010.03.054](https://doi.org/10.1016/j.cell.2010.03.054)
44. Noble, W.S.: What is a support vector machine? *Nat. Biotechnol.* **24**(12), 1565–1567 (2006)
45. Pearson, K.: On lines and planes of closest fit to systems of points in space. *Philos. Mag.* **2**, 559–572 (1901)
46. Saez-Rodriguez, J., Alexopoulos, L.G., Epperlein, J., Samaga, R., Lauffenburger, D.A., Klamt, S., Sorger, P.K.: Discrete logic modelling as a means to link protein signalling networks with functional analysis of mammalian signal transduction. *Mol. Syst. Biol.* **5**, 331 (2009). doi:[10.1038/msb.2009.87](https://doi.org/10.1038/msb.2009.87) (msb200987 [pii])

47. Schlens, J.: A tutorial on principal component analysis: derivation, discussion, and singular value decomposition. http://www.cs.princeton.edu/picasso/mats/PCA-Tutorial-Intuition_1p.pdf. Accessed 1 Aug 2013
48. Schoeberl, B., Eichler-Jonsson, C., Gilles, E.D., Muller, G.: Computational modeling of the dynamics of the MAP kinase cascade activated by surface and internalized EGF receptors. *Nat. Biotechnol.* **20**(4), 370–375 (2002)
49. Spellman, P.T., Sherlock, G., Zhang, M.Q., Iyer, V.R., Anders, K., Eisen, M.B., Futcher, B.: Comprehensive identification of cell cycle-regulated genes of the yeast *Saccharomyces cerevisiae* by microarray hybridization. *Mol. Biol. Cell* **9**(12), 3273–3297 (1998)
50. Sunnåker, M., Stelling, J.: Model extension and model selection. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
51. Tentner, A.R., Lee, M.J., Ostheimer, G.J., Samson, L.D., Lauffenburger, D.A., Yaffe, M.B.: Combined experimental and computational analysis of DNA damage signaling reveals context-dependent roles for Erk in apoptosis and G1/S arrest after genotoxic stress [Research Support, N.I.H., Extramural Research Support, U.S. Gov't, Non-P.H.S.]. *Mol. Syst. Biol.* **8**, 568 (2012). doi:[10.1038/msb.2012.1](https://doi.org/10.1038/msb.2012.1)
52. Tucker, L.R.: Some mathematical notes on three-mode factor analysis. *Psychometrika* **31**, 279–311 (1966)
53. Vilela, M., Danuser, G.: What's wrong with correlative experiments? *Nat. Cell. Biol.* **13**(9), 1011 (2011). doi:[10.1038/ncb2325](https://doi.org/10.1038/ncb2325)
54. Whitton, J.L., Cornell, C.T., Feuer, R.: Host and virus determinants of picornavirus pathogenesis and tropism. *Nat. Rev. Microbiol.* **3**(10), 765–776 (2005). doi:[10.1038/nrmicro1284](https://doi.org/10.1038/nrmicro1284)
55. Wiley, H.S., Shvartsman, S.Y., Lauffenburger, D.A.: Computational modeling of the EGF-receptor system: a paradigm for systems biology. *Trends Cell. Biol.* **13**(1), 43–50 (2003)
56. Zou, H., Hastie, T., Tibshirani, R.: Sparse principal component analysis. *J. Comput. Graph. Stat.* **15**(2), 265–286 (2006). doi:[10.1198/106186006x113430](https://doi.org/10.1198/106186006x113430)

Chapter 7

Optimization in Biology Parameter Estimation and the Associated Optimization Problem

Gunnar Cedersund, Oscar Samuelsson, Gordon Ball, Jesper Tegnér and David Gomez-Cabrero

Abstract Parameter estimation—the assignment of values to the parameters in a model—is an important and time-consuming task in computational biology. Recent computational and algorithmic developments have provided novel tools to improve this estimation step. One of these improvements concerns the optimization step, where the parameter space is explored to find interesting regions. In this chapter we review the parameter estimation problem, with a special emphasis on the associated optimization methods. In relation to this, we also provide concepts and tools to help you select the appropriate methodology for a specific scenario.

Keywords Parameter estimation · Optimization · Heuristic · Fitness function

7.1 Introduction

Mathematical models have been part of biological research for more than half a century. One of the starting points for this development was the now classical model for the action potential in an axon, developed by Hodgkin and Huxley [27, HHM]. HHM

G. Cedersund (✉)

Integrative Systems Biology, Department of Biomedical Engineering, Linköping University, 58185 Linköping, Sweden
e-mail: gunnar.cedersund@liu.se

G. Cedersund

Department of Clinical and Experimental Medicine, Linköping University, 58185 Linköping, Sweden

O. Samuelsson · G. Ball
Stockholm, Sweden

J. Tegnér · D. Gomez-Cabrero

Unit of Computational Medicine, Department of Medicine, Karolinska Institutet, Solna, Sweden
e-mail: david.gomezcabrero@ki.se

J. Tegnér · D. Gomez-Cabrero

Center for Molecular Medicine, Stockholm, Sweden

© Springer International Publishing Switzerland 2016

L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_7




Fig. 7.1 Schematic of the basic Hodgkin-Huxley model (HHM) of a neuron. The figure depicts the basic components of HHM models that represent the biophysical characteristics of cell membranes. In the model Capacitance (C_m) represents the lipid bilayer. Nonlinear (g_n) and linear (g_L) conductances represent voltage-gated and leak ion channels respectively. Batteries (E) represent the electrochemical gradients driving the flow of ions. Current sources (I_p) represent ion pumps and exchangers. Figure obtained from Wikimedia, under Creative Commons CC0 1.0 Universal Public Domain Dedication. Author: Krishnavedala

explains these action potentials using a set of coupled nonlinear ordinary differential equations (ODEs), which are associated to membrane channels with specific time- and voltage-dependent properties (Fig. 7.1). The HHM is of major relevance because (i) it was the culmination and integration of a large number of experiments [17, 52] where the model provided a mechanistically unified vision of the system; and because (ii) HHM generated specific predictions that were subsequently validated using new single-channel recording techniques [26]. For these reasons, Hodgkin and Huxley were awarded the Nobel Prize in 1963, the first Nobel Prize in physiology awarded for the development of a mathematical model. An illuminating account of the development of both experiments and model is available in [26].

However, those early successes by Hodgkin and Huxley were possible because they could bypass one of the biggest hurdles in most current biological modeling: the simultaneous determination of all the values of the parameters in the model from systems-level data. These parameters that need to be estimated appear in the ODEs, and in the case of HHM, these parameters could be experimentally characterized from initial targeted experiments. In other words, their only problem was that of forward simulation, which was possible even in the pre-computer era. For most other biological models, parameters cannot be determined directly in specific experiments, but instead need to be estimated simultaneously from systems-wide data. Particularly, one needs to search the space that is spanned by the unknown parameters in the model, in order to characterize which parts of the parameter space are consistent with experimental data and prior knowledge. This task of determining parameters from data is known as parameter estimation, and a key step in this task is to optimize the agreement between model and data. In other words, optimization research enters as a natural component in research on parameter estimation for biological models

[2, 22, 23, 36]. This optimization step is not trivial, and there are a wide variety of methods to choose from.

In this chapter, we will describe the basics of the parameter estimation problem, look in depth at the associated optimization problem, and at the available tools to solve it. In Sect. 7.2, the parameter estimation problem is introduced using an intuitive example and in Sect. 7.3, the problem is treated from a more mathematical viewpoint. Section 7.4 describes several meta-heuristic optimization approaches and Sect. 7.5 compares these approaches and provides useful inputs for selecting among them. Finally, Sect. 7.6 provides a closing overview of the chapter and provides links to open challenges, most of which are discussed in other chapters of this book.

7.2 The Parameter Estimation Problem


Let us start by introducing the parameter estimation problem from an intuitive point of view, through a specific biological model. The example is formulated as an ODE model, which is the most widespread formalism in systems biology [10, 31, 37, 48]. An ODE model can be formulated in state-space form, which is a mathematically well explored representation [1, 6], written as:

$$\dot{x}(t) = f(x(t), u(t), P) \tag{7.1}$$

$$y(t) = h(x(t), u(t), P) \tag{7.2}$$

where $x(t)$ are the states, \dot{x} are the time-derivate of the states, f and h are smooth nonlinear functions, P are the parameters, and $u(t)$ and $y(t)$ are input and output signals, respectively. All symbols are vectors.

Let us now consider the specific model in Fig. 7.2. This figure depicts an interaction-graph for a simple model of the insulin signaling network, taken from [43]. As can be seen, the model describes how insulin binds to the insulin receptor (IR), which thus becomes phosphorylated (IRp), and able to phosphorylate the protein IRS, which phosphorylates the protein PKB, which finally stimulates transport of Glut4 to the membrane, which together with Glut1 transports glucose into the cell. There is a straightforward approach to go from an interaction-graph to a state-space formulation [50], and the full ODEs are given in [43]. A short summary of this process is as follows. The states correspond to the concentrations (or, alternatively, the amounts) of the different molecules, i.e. $[IR]$, $[IRSp]$, etc. The concentrations are affected by the reactions, and the rate of these reactions is the first thing that needs to be defined. In this model, mass-action kinetics are assumed, which means that the rates are given by the product of the concentrations of the substrates and the regulators multiplied by an unknown rate constant [3]. The unknown rate constants are called kinetic rate constants, here denoted k_i , and they thus regulate the speed of each individual reaction. It is often these rate constants that need to be determined from experimental data. These rates are summed together to form the ODEs, so that rates corresponding to ingoing reactions appear with a plus sign, and rates corresponding to outgoing reactions appear with a minus sign. In other words, the ODE for $[IRS]$



$$u = [\text{insulin}]$$

$$y = (k_{y1}[\text{IRp}], k_{y2}[\text{IRSsp}], k_{y3}[\text{PKBp}])$$

Fig. 7.2 A schematic description of an insulin-signaling network. The input and output signals are the concentrations of insulin, and the measurements of the phosphorylated proteins, *IRp*, *IRS_p* etc. This schematic description is referred to as an interaction-graph, and it corresponds one-to-one with the ODEs, as is described for instance in [43], which also is the paper from which this model is taken

is given by:

$$[\dot{IRS}] = \frac{d[IRS]}{dt} = \dot{x}_{IRS} = k_{2b}[IRSsp](t) - k_{2f}[IRS][IRp](t) \quad (7.3)$$

In the model specification, each parameter p_i (usually, but not always, corresponding to some rate constant k_j) is specified to have values in a specified range. This range is here defined by l_b below and u_b above

$$l_{b,i} < P_i < u_{b,i} \quad (7.4)$$

These boundaries are referred to as box-constraints. The box-constraints should include all physically possible parameter values. If the range of physiologically likely parameters is unknown, a box-constraint would still be required in order to restrict the parameter space. Many physiologically relevant models are at least bounded by 0 below, and the time-scale of the observed dynamics can often give a suggestion for an upper bound as well. The box-constrained parameter space may also be referred to as the search space, and a parameter set is a point in the (box-constrained) parameter space.

The parameter estimation problem is to find those parameter sets that both are biologically feasible and that display a satisfactory agreement with experimental observations. An experimental observation may refer to any set of values that are available from experimental measurements. One example is a so-called perturbation experiment, where a cell is perturbed from steady-state, and where the transient response of the system is observed. In the insulin example above, insulin is such a perturbation, i.e. the levels of all the states are assumed to be in steady-state before the perturbation, and one then follows how the signaling intermediates ($[IRp]$, $[IRSp]$, etc.) are changing over time. Such perturbation experiments can also be done in other systems, e.g. gene regulatory networks. In gene regulatory networks the availability of perturbation experiments is growing as several technological advances [39] have allowed the generalization of gene perturbation screens of the systems. One such example is RNA interference experiments, which are being used as a tool for both association discovery and for validation purposes [53].

The agreement between model and data can be of two types: qualitative and quantitative. A qualitative agreement can often be assessed by mere inspection and reasoning. For instance, the initial response to an insulin stimulation is that $[IRp]$ and $[IRSp]$ goes up. Another qualitative observation may be that the system oscillates, i.e. that it changes periodically up and down over time. In these two cases, the model will qualitatively agree with the data if it goes up in response to insulin, or if it oscillates, respectively. However, experimental data that have been collected often contains fundamentally more information than that: it usually contains quantitative information. Such quantitative information is typically measured as the average deviation from the data points (Eq. 7.5 below). Loosely, parameters that give an acceptable agreement with the data are called feasible, and other parameters are referred to as infeasible. Whereas the question of whether or not simulations for a specification parameter agree qualitatively with data can be assessed by mere inspection, the quantitative agreement requires a more formal treatment. This parameter characterization—which simulations agree, and which do not agree, with the data—is the task of parameter estimation.

Originally parameter estimation was done “by hand”, by combining physiological knowledge with reasoning and manual tests. Such an approach is time-consuming, requires expert knowledge on the exact role of each parameter, and is therefore infeasible for most biological models. With the invention and development of computers, this estimation “by hand” could be replaced by more exhaustive searches in the parameter space, since computers allow for thousands or millions of parameter sets to be tested. This development of computers has gone hand-in-hand with the developments in the field of optimization. More specifically, the application of methods from the field of optimization has been essential for the exploration of extremely large parameter spaces. This is the case, since the task of covering large parameter spaces seldom can be solved by “brute-force” approaches, even with the rapid advances in computer power. Let us therefore now turn to a more formal treatment of these concepts, using the language of mathematical optimization.

7.3 The Optimization Problem in Parameter Estimation

In Sect. 7.2, we learned that parameter estimation is concerned with the search for feasible parameter sets in the box-constrained parameter space, and that “feasible” measures the agreement between our experimental observations and the model simulations, either qualitatively or quantitatively. Now the challenge is to transform the concepts “feasible” and “agreement” into mathematical terms.

Let us start by considering a specific set of data and simulation possibilities, again taken from insulin signaling [5]. The example is depicted in Fig. 7.3. The data reflects the amount of IRSp, and at time $t = 0$ the system (isolated fat cells) have been stimulated with 100 nM insulin. The collected data points are depicted in blue, where the error bars depict the uncertainties (top to bottom reflects 2 standard deviations, 2σ). As can be seen, these data have a big uncertainty, but one can nevertheless see certain features clearly. First, the system responds by going up, but then around one $t = 1$ min, the response decreases, and from somewhere between 5 and 10 min and onwards, the system has settled at an intermediate steady-state level. Now, when a model tries to mimic these data, there are two aspects that can be taken into account. First, one may capture the presence of the overshoot (that the response goes up and then down), and this is captured by the blue but not the green line. Second, one may capture the quantitative features of the data, e.g. the location of the final steady state. This steady-state is, conversely, captured by the green but not the blue line. These two lines correspond to simulations with the studied model, done using different parameter values, and this example illustrates that quantitative and qualitative aspects




Fig. 7.3 Example of data and model simulations. This example is also from insulin signaling, as described in [5]. In blue you see the data points with uncertainty, and the other lines show simulations for different parameters values: one that is qualitatively wrong but quantitatively right (green), one with the opposite problem (blue), and one from a feasible parameter (red). A cost function is formed by summing together the normed and squared residuals, as described in Eq. (7.5)

do not always coincide. The task is therefore, to decide upon how to quantify the features (quantitative or qualitative) that are important in the data, and then optimize this quantification, to find the best parameters (red line).

In other words, we need an automatic way to evaluate how “feasible” a parameter set is. Let us start by denoting the experimental data point at time t_k with $y(t_k)$. Similarly, let $\hat{y}(t_k|P)$ denote the simulated data point, at time t_k and using the parameters P . Now, we are interested in the differences between these two values, the so called residuals $r(t_k) = y(t_k) - \hat{y}(t_k|P)$. In general, the larger the residuals, the worse the model, so we are interested in the average size of the residuals. However, we want them to be always positive (since it does not matter if the model is above or below the data). Furthermore, if the data point is highly uncertain (there is a big $\sigma(t_k)$), that residual should be less important, when taking the average. All of these requirements are met by studying the following cost function

$$V(P) = \sum_{k=1}^N \frac{(y(t_k) - \hat{y}(t_k|P))^2}{\sigma(t_k)^2} \quad (7.5)$$

As can be seen the residuals have been squared (to make the residual always positive), normalized with the standard deviation of the measurement uncertainty (divided with σ), and then summed together (giving the average size, considering all residuals). This cost function is the most common choice, and it is usually denoted the chi-square cost function. Independently of which cost function, $V(P)$, that is used, the final parameters are in the end defined by

$$\hat{P} = \underset{P}{\operatorname{argmin}} V(P) \quad (7.6)$$

Let us now consider a general optimization problem, and see if we have succeeded with the reformulation of the parameter estimation problem. In mathematical optimization, the problem is defined by three elements: (a) the objective function—in our case the evaluation function described in (5); (b) a set of decision variables that can vary during the search—in our case the parameters P ; and (c) the constraints that limit the value of the decision variables or the relations between them—in our case the box constraints (4). The problem we have described is thus a mathematical optimization problem. More specifically, our problem is a continuous optimization problem, because the decision variables are all real numbers. A systems biology oriented description of optimization is available at [2].

Let us now consider a complication of this general scenario that often happens: additional constraints. Such constraints may stem from prior knowledge, not contained in the data. Consider for instance a system that contains the variables C and D , where expert knowledge assures us that D cannot be more than twice the value of C at any time. Then this extra information should be added to the problem setting. The information can be added to the system in two ways. The first way is to include the information as a constraint in the objective function (in the same way the boxed-parameter space defines inequalities):

$$D(p, t) < 2C(p, t) \forall t \quad (7.7)$$

The second way (which we favor here, as it is easier to integrate with the heuristic methodologies shown in the next section) is to construct a new objective function $V_{tot}(P)$ that is the weighted combination of the normal “data-based” objective function $V(P)$ (e.g. that in Eq. (7.5)) and a new one, $eval2(P)$

$$V_{tot}(P) = V(P) + w \cdot eval2(P) \quad (7.8)$$

$$eval2(p) = \sum_{k=1}^N dist(C, D, P, t_k) \quad (7.9)$$

$$dist(C, D, P, t) = \begin{cases} (2C(p, t) - D(p, t)) & \text{if } (2C(p, t) > D(p, t)) \\ 0 & \text{otherwise} \end{cases} \quad (7.10)$$

As can be seen, $eval2(P)$ will be zero if the relation based on expert knowledge between C and D is fulfilled by the simulation, and if the constrain is violated, $eval2(P)$ will grow with the size of the violation. The parameter describes the weighting between the two sources of knowledge: the experimental data ($V(P)$) and the expert knowledge ($eval2(P)$). These kind of *ad hoc* expansions of the cost function can be used to incorporate also other types of knowledge, or preferences regarding how the data should behave. In the example in Fig. 7.3, one could for instance decide that the model should have an overshoot. One could then add a similar punishment as $eval2$ in (8) to such simulations that do not produce such an overshoot.

An alternative approach to a weighted combination of different elements in a cost function is to compute Pareto Optimal (PO) solutions. Considering two evaluation functions, A and B, a PO solution is such that there is no other solution with better evaluation for both A and B simultaneously. The study of Pareto Optimal is an active field in multi-objective optimization but we do not consider it further in this chapter. Relevant examples of multi-objective optimization in biology can be found in [8, 57, 58].

Finally, optimization has been used in the study of many different types of biological systems, and the details of the formulations may differ from setting to setting. Common examples include flux balance analysis [51], metabolic engineering, signaling pathways [43, 54], reverse engineering [24], and stochastic modelling [60]. However, a common feature, usually appearing in biological systems is that there are many local optima. This means that global optimization algorithms should be used to search the parameter space. Such global methods are usually described using meta-heuristics.

7.4 Meta-Heuristics: A Tool for Optimization

7.4.1 Introduction to Meta-Heuristics

The term meta-heuristics is a composition of the Greek words *meta* (beyond) and *heureskin* (to find) or *heuristic* (rules of thumb). Meta-heuristics combine basic heuristic methods in higher-level frameworks to efficiently explore the box constrained objective function to be optimized. Introductions to meta-heuristics can be found at [19].

Meta-heuristics require little or no information about the problem to be solved, and they are sometimes referred to as black-box optimization methods. Interestingly, the same meta-heuristic algorithm can usually be applied to a wide range of different problems. Meta-heuristics do not make explicit use of Hessians or gradients as opposed to gradient descent methods, and can therefore be used also on discontinuous problems. More importantly, most parameter estimation problems are multi-modal, i.e., they contain several optima; hence, traditional gradient-based local search methods usually fail at locating the global optimum in such problems. In the search for the global optimum, one is therefore directed to use global optimization methods where one alternative is to use meta-heuristic methods.

Within the area of global optimization there are two different approaches, deterministic and stochastic methods, and they differ in their ability to guarantee convergence towards the global optimum. The deterministic methods can in theory provide a level of assurance of finding the global optimum. However, the computational time often increases exponentially with problem size, which often makes it impossible to locate the global optimum in a reasonable time. The stochastic methods, including meta-heuristics, use random components to locate the vicinity of the optimum. The trade-off is that one cannot guarantee the identification of the global optimum. However, the computational cost is significantly lower for stochastic methods compared to deterministic methods and this is an important advantage in large problems.

Many different meta-heuristic methods exist and in this chapter only a few of the most common methods are described. It should also be noted that the way of implementing the algorithms can differ widely, and here, only the pseudo-code is shown in order to illustrate the basic structure of the algorithms. Let us now turn to the specific algorithms.

7.4.2 Simulated Annealing

Simulated annealing was originally developed for solving combinatorial optimization problems [35]. However, the scope of use for the algorithm has grown and it is currently used in various optimization areas for numerous types of optimization problems. The algorithm resembles the internal energy of atoms at different temperature levels. The analogy is as follows. At a high temperature, multiple states for the atoms are possible and a random state is adopted. A reduced temperature restricts

the number of possible states to ultimately converge to a single state at the absolute zero temperature. Finally, if the atoms are allowed to equilibrate at each temperature level, the final state is also the optimal state in terms of energy. In optimization, the energy function, which describes the states of the atoms, is replaced by the objective function $V(P)$ as described in [35]. Thus, a state of the atoms is replaced by a parameter vector in the search space.

In practice, the search space may be explored by a geometrical figure, e.g. a simplex. In a D -dimensional search space, a simplex is a hyper-triangle with $D+1$ vertexes, where each vertex consists of a parameter vector. The simplex is programmed to update the parameter vector (vertex) with the worst objective function value with a better one. The updating procedure of the simplex is iterated, which results in a simplex that deforms and moves across the search space. At high temperatures, the simplex performs more or less a random search. However, the probability to include a parameter vector with higher cost than those already included in the simplex becomes smaller as the temperature decreases. Thus, at low temperatures, the simplex performs a local search.

To be able to guarantee convergence to the optimal parameter vector, an impractically slow cooling scheme must be used [55]. In practice, pre-defined faster cooling schemes are used, which do not guarantee global convergence. In this chapter, a slightly modified Simulated annealing procedure based on the Nelder-Mead Downhill Simplex algorithm (NMDS) is shown, for which a description can be found in [46]. The modification makes it possible to search for several optima simultaneously by using several simplexes, see “Algorithm 4” on page 27 in [44].

In this modified version, the starting point for each simplex is found by a clustering technique, which makes sure that the simplexes are started with a sufficient Euclidian distance between them. The temperature is decreased after a specified number of iterations for each one of the simplexes. The magnitude of the temperature decrease is set by a reduction factor which reduces the temperature in a step-wise manner. As the temperature reaches the defined end temperature it is set to zero, and no uphill moves are allowed. This is similar to a local search. A short pseudo-code description is given in Fig. 7.4, to illustrate the algorithm. For further information, see [44].

7.4.3 Scatter Search

The meta-heuristic method Scatter search was originally introduced by Glover in the late 1970s [18]. Scatter search involves strategic combinations of parameter vectors to generate improved parameter vectors.

As originally proposed in [18], the core of Scatter search is the Reference Set (RefSet) which contains a limited number of promising parameter vectors. The number of parameter vectors, denoted R_s , in the RefSet is an option specified by the user. Different heuristic methods interact with the RefSet during the search. Most scatter search implementations follow *the five step template*, originally described by Glover. These five steps, or sub-methods, are: *Diversification Generation Method*,

Pseudo-code Simulated Annealing

Requirements: Start temperature T_s , end temperature T_e , reduction factor α , start guess P_0 , number of iterations I at each temperature, number of iterations at absolute zero temperature I_z , and maximum number of simplexes S .

- 1: Initiate parameters and set $k = 1$
 - 2: Set $T_k = T_s$
 - 3: **while** $T_k > T_e$ **do**
 - 4: Perform I iterations of Nelder-Mead Downhill Simplex algorithm for the S simplexes
 - 5: Set $T_{k+1} = \alpha T_k$
 - 6: Calculate R new restart points with single-linkage-clustering where $R \leq S$
 - 7: Set $k = k + 1$
 - 8: **end while**
 - 9: Set $T_k = 0$
 - 10: Perform I_z iterations of Nelder-Mead Downhill Simplex algorithm at the absolute zero temperature
 - 11: **return** Best parameter vectors found during the last temperature iteration
-

Fig. 7.4 Pseudo-code: simulated annealing

Improvement Method, Reference Set Update Method, Subset Generation Method, and a Solution Combination Method. The methods will be briefly explained in the way they are implemented in SSm (Scatter Search for Matlab). For a detailed description, the reader is referred to the original description of SSm [15].

- *Diversification Generation Method* generates a set d consisting of D_p diverse trial parameter vectors. The parameter values for the D_p parameter vectors with D parameters are selected from ND cells, where the rows are N intervals for each parameter and the D columns are the different parameters. The intervals for each parameter are defined by dividing the region between the lower and upper boundaries in N intervals. Thus, all parameters have N intervals regardless of range between lower and upper boundary. Since the D_p parameters in d are selected from all of the ND cells, one can be certain that the trial parameter vectors are diverse. The probability $prob_{i,n}^{p+1}$ of choosing a new parameter vector $p + 1$ with one of its parameters in one of the N intervals is inversely proportional to the number of previous parameters in the specific interval, i.e.,

$$prob_{i,n}^{p+1} = \frac{\frac{1}{f_i}}{\sum_{k=1}^N \frac{1}{f_k}} \tag{7.16}$$

where f_i is the number of previous values of parameter i in the interval n . When the boundaries of one or several parameters are of different orders of magnitude, the intervals can be defined according to a logarithmic distribution providing a spread of the D_p parameter vectors across the entire search space.

- *Improvement Method* is optional to include in a Scatter Search design but is normally required to obtain high quality parameter vectors (in terms of cost function value). One option is to use the local search method *Dynamic hill climbing* [41] as an improvement method. The improvement method can be restricted by activating a merit filter which ensures that no local search is performed from a worse initial point than previously found.
- *Reference Set Update Method* updates the R_s parameter vectors according to quality and diversity. The first half of the RefSet is updated with high quality parameter vectors (low cost function value) which are selected from the D_p diverse parameter vectors such that

$$V(P_j) \leq V(P_{j+1}) \leq \dots \leq V(P_{R_s}) \quad (7.17)$$

where $V(P_j)$ is the cost function value for one of the R_s parameter vectors j . In order to obtain a RefSet with both parameter vectors with low cost and parameter vectors which are spread in the search space, the second half of the RefSet contains spread or *diverse* parameter vectors. Diversity is measured either by the spread in the parameter vectors in terms of Euclidian distance, or by the spread in the parameter vectors in terms of direction. For further information, see pages 57–60 in [14].

- *Subset Generation Method* creates a combination of parameter vectors and in this case, pairs of all P_1, P_2, \dots, P_{R_s} parameter vectors.
- *Solution Combination Method* explores the distance between all paired parameter vectors to find new parameter vectors. Different approaches exist on how to do the search, but in this implementation only linear or hyper-rectangle searches are used [14]

How the different methods are combined in a basic search is shown in the pseudo-code, in Fig. 7.5. It should be pointed out that it is the way of implementing the five methods that decides the sophistication of the algorithm rather than the specific methods [14].

7.4.4 Genetic Algorithms

While many bio-researchers rely on computational methodologies for the analysis of large data-sets, it is interesting to observe how many methodologies have been designed in the recent decades by considering the principles observed in biological systems. For instance, machine-learning applications have used neurons as a reference to develop neural network classifiers [4, 30]. A second example, Genetic

Pseudo-code Scatter Search

Requirements: R_s (size of *RefSet*) and number of D_p diverse parameter vectors in the set d . Set remaining options to default.

```

1: Start with  $d = \emptyset$ 
2: Repeat
3:     Use the Diversification generation method to construct a
       parameter vector and apply the Improvement Method Local solver
4:     Let  $P$  be the resulting parameter vector
5:     if  $P \in d$  then
6:          $d = d \cup \{P\}$ 
7:     Else
8:         Discard  $P$ 
9:     end if
10: until  $|d| = D_p$ 
11: Use the Reference set update method to build  $RefSet = \{P_1, \dots, P_{R_s}\}$  with
       the best  $\frac{R_s}{2}$  quality parameter vectors and  $\frac{R_s}{2}$  diverse parameter vectors in  $d$ 
12: Sort the parameter vectors in RefSet according to their objective function
       value such that  $P_1$  is the best parameter vector and  $P_{R_s}$  is the worst
13: Set  $NewSolutions = TRUE$ 
14: while  $NewSolutions$  do
15:     Generate  $NewSubsets$  with the Subset Generation Method
16:     Set  $NewSolutions = FALSE$ 
17:     while  $NewSubsets = \emptyset$  do
18:         Select the next subset  $s$  in  $NewSubsets$ 
19:         Apply the Solution Combination Method to  $s$  to obtain new
           trial parameter vectors
20:         Apply the Improvement Method to the  $s$  trial parameter
           vectors
21:         Apply the RefSet Update Method to the  $R_s \cup s$  parameter
           vectors
22:         if RefSet has changed then
23:             Set  $NewSolutions = TRUE$ 
24:         end if
25:         Delete  $s$  from  $NewSubsets$ 
26:     end while
27: end while

```

Fig. 7.5 Pseudo-code: scatter search

Algorithms (GA), arguably the most well-known optimization procedure inspired by biology, is based in the studies of DNA sequence evolution [20] (pseudo-code in Fig. 7.6). GA mimics the process of natural selection in order to find quality solutions in optimization problems. Briefly, GA mimics the natural selection process by considering an initial population of solutions (that can be generated randomly)

Pseudo-code Genetic Algorithm

Requirements: maximum number of iterations (It_{max}), number of elements in a population (pop), maximum CPU, Mutation, Crossover and Selection.

```

1:   Initiate randomly the solution set: pool.
2:   Set  $k = 1$ 
3:   while  $It_{max} > k$  do
4:       for in 1:pop
5:           Select randomly two solutions a,b.
6:           Crossover(a,b): generates newsol.
7:           Mutation(newsol)
8:           Add newsol to pool
9:       end for
10:      Selection(pool)
11:      Set  $k = k + 1$ 
12  end while
11:  return Best solution found during the process

```

Fig. 7.6 Pseudo-code: genetic algorithm

and iterating over the population in a survival-of-the-fittest approach. The iterations make use of the following three elements:

Selection Criteria periodically the population of solutions is evaluated, and those that score poorly are discarded. A selection criterion, based on the evaluation function, is used to make the decision. Random discards or weight-based discards are also considered as options.

Genetic Operators solutions in a population will be used to generate new solutions that will be added to the population pool. The new solutions are generated through “cross-over” (using two solutions and generating a novel, “child” solution from them) and “mutation” (modifications of a solution). There are many ways to combine cross-over and mutations to generate new solutions. Once a new solution is generated it will also be evaluated and added to the pool of solutions.

Termination Criteria it is necessary to define when the algorithm will stop searching for new solutions. Termination criteria usually considered are (i) number n of iterations without improvement of a solution, (ii) finding a solution that reaches a lower boundary, (iii) maximum number of iterations, (iv) maximum running time or (v) a combination of any of the criteria (i)–(iv).

GA, as many other meta-heuristic algorithms, faces the challenge of setting a good trade-off between exploration and local optimization. Exploring is necessary when finding optimal solutions and it can be stressed by (a) including selection criteria that do not consider uniquely “best evaluated” solutions but enforce a heterogeneity in the solution pool; a second option, among others, is to (b) consider mutations that may lower or increase the evaluation score of a solution. The burden with high

heterogeneity is that the computational time required in the search will increase dramatically; for this reason it is important to limit exploration, and for this reason local optimization may help guide the search more rapidly. Importantly, “naïve GA” has been shown to have a premature convergence that prevents a global optimization [28]. Interestingly, there are hybrid methodologies combining GA with local optimum search exploiting the properties of GA without limiting its exploring ability [11, 25, 40].

It is worth considering that GA was originally conceived to solve discrete (or integer-based) optimization problems [28]. In the context of parameter estimation, binary representations of parameter values may solve the problem of representing parameter sets containing real numbers, but this option will also increase the number of solutions to investigate. Finally, GAs have developed into a wider area of optimization algorithms known as *Evolutionary computation* (EC) that are able to deal with continuous values [12].

7.4.5 Particle Swarm Optimization

EC algorithms find solutions through continuous optimization of a population of solutions. Within this setting Particle Swarm Optimization (PSO) is amongst the most well-known algorithms [32, 33] other relevant representatives are Ant Colony Optimization [13] or Genetic Programming [34].

Given a real parameter space, PSO first places a set of particles (agents) randomly in the parameter space, assigning them a random initial velocity. Every particle receives a fitness-score (the cost, $V(P)$, delivered by the objective function) associated to the parameter set visited; those values are stored. Additionally the parameter set with the globally best fitness-score is also stored in **GBest**. Particle positions are updated iteratively based on a (randomly) weighted combination of: (1) the inertia of the particle, (2) a vector pointing from the present particle position to the best known position and (3) a vector pointing from the present particle position towards **GBest**. In this way particles share information of the best areas to search (social knowledge) but they also keep their best-visited solutions (cognitive knowledge). The iterative movement of particle i is explained by the following two equations:

$$vel(i) = \sigma vel(i) + r_1 * c_1 (pos(i) - best(i)) + r_2 * c_2 (pos(i) - GBest) \quad (7.18)$$

$$pos(i) = pos(i) + \chi vel(i) \quad (7.19)$$

$pos(i)$ and $vel(i)$ respectively denote the current position and the current velocity of particle i . The symbols c_1 and c_2 are integer non-negative values, named cognitive and social respectively, r_1 , r_2 are real values drawn from $[0,1]$, σ and χ are non-negative real values, named inertia weight and constriction factor, respectively.

Pseudo-code Particle Swarm Optimization

Requirements: maximum number of iterations (It_{max}), number of particles ($nparticles$), maximum CPU, cognitive weight (c), social weight (s), inertia weight (in), constriction factor (con).

```

1:   Initiate set of particles: randomly assign positions.
2:   for  $i$  in  $particle\_set$  do
3:      $evaluate(i)$ 
4:      $best(i)=position(i)$ 
5:   end for
6:   Set  $k = 1$ 
7:    $GBest: position(i)$ , that  $evaluate(i)=\min(evaluate(i), i \in particle\_set)$ .
8:   while  $It_{max} > k$  do
9:     for  $i$  in  $1:nparticles$  do
10:      calculate  $velocity(i)$ 
11:       $update\_position(i,GBest, best(i), velocity(i), c, s, in, con)$ .
12:      if  $evaluate(i)<evaluate(best(i))$  then
13:         $best(i)= position(i)$ 
14:      end if
15:      if  $evaluate(i)<evaluate(GBest)$  then
16:         $GBest = position(i)$ 
17:      end if
18:    end for
19:     $k=k+1$ 
20:  end while
21:  return  $GBest, evaluate(GBest)$ 

```

Fig. 7.7 Pseudo-code: particle Swarm optimization

PSO has become a widely used methodology mainly because it is easy to implement, it is flexible and can be applied to most continuous-based problems [45] and requires little fine-tuning (few parameters to adjust) [33]. The caveat of PSO is that it is mainly limited to continuous problems (despite discrete to continuous mapping approaches [47]). The pseudo-code for PSO is found in Fig. 7.7.

7.5 Performance of Meta-Heuristics

Defining the difference between good and bad performance is essential when comparing algorithms. In accordance to the introduction, several parameter sets can give the same model output. Depending on the search purpose, different criteria for the performance exist. Denoting by a *run* the “*completion of the selected optimization*

algorithm”, four possible criteria (among many others) to evaluate meta-heuristics are:

- (1) **BS**: Given n runs, with limited time x , what is the quality of the best solution found in all the runs, i.e. what is the best cost found? BS provides a measure of how good is an algorithm searching for good solutions; however if the best solution is found once in 100 runs, that would require running the algorithm many times in order to find such a good solution.
- (2) **BSa**: Given n runs, with limited time x , what is the average quality (average cost) of the solutions found, when pooling the best solution per run. BSa is useful when we may require algorithms that do not provide the best solutions (or lower BS) but on average perform better.
- (3) **BSff**: similar to *BS* but instead of fixing to a limited time x , the run is fixed to a maximum number of objective-function evaluations.
- (4) **BSaff**: similar to *BSa* but instead of fixing to a limited time x , the run is fixed to a maximum number of objective-function evaluations.

Several studies have approached the comparison of methodologies in different optimization problems. Unfortunately, there is no perfect optimization algorithm; given a performance evaluation, for each optimization methodology there is always a class of problems where the methodology will be worse than for other methodologies. That is, there is no method that will be optimal for all problems. This is a direct consequence of the No-Free-Lunch theorems presented in [59]. However, for making a selection among possible optimization algorithms, we recommend to follow two pieces of advice:

- *Hybrid version* while there is no “best-ever” algorithm for all types of optimization problems, it has been observed that hybrid methodologies, those combining ideas/methods from several optimization methodologies, tend to work robustly in all optimization problems. We recommend the implementation of hybrid versions.
- *Widely used* We can observe that there exist several optimization methods that are continuously implemented in parameter-fitting tools. For instance **Neurofitter** [56], developed to find parameter sets that able to reproduce experimental data in neuronal models, includes Random search, Particle Swarm Optimization, Evolution strategies, Multi-Start Local Optimization and combinations of them. Similarly, **COPASI** [29], an open source package that allows the generation and simulation of biological processes, includes Evolutionary Strategy, Genetic Algorithm(s), Particle Swarm Optimization, Random Search, and Simulated Annealing. This may reflect (1) the flexibility of those methods, (2) the low complexity to implement them, but also (3) that there are many successful examples of their applications.

7.6 Conclusions

The present chapter has provided a brief introduction to the field of optimization and its applications in the challenge of parameter estimation in biology. We have provided introductory explanations to some of the most widely used concepts and optimization algorithms. The focus of this chapter has been on meta-heuristic methods, within the field of stochastic global optimization. The benefit of these methods is that they can work on any optimization problem, and that they can deal with multi-modal systems: systems with multiple local optima. An outstanding question remains what algorithms are best in a specific situation. Herein, we have not really taken a stand on this issue, but nevertheless provide some tools and concepts that may be useful when comparing methods in different situations.

There are many implementations of these methods, and there are also many ways to implement models. Some of these model-formulation alternatives are standardized markup languages, such as SBML [9] and CellML [16], and some of these also have resources for parameter estimation. For instance, fitMatlabCellML has been developed for CellML. Similarly, sloppyCell [42], COPASI, and SBtoolbox [49] have been developed for SBML. However, for independent development and testing of these methods, it will always remain a tractable alternative to work with your own implementations of these algorithms.

In the following chapters, optimization tools are used to find parameter sets [38], and modified optimization algorithms are also used to generate a pool of “good-quality parameter sets”, used for prediction uncertainty analysis [7, 21]. These slightly different settings will put new demands on future developments of optimization algorithms. All in all, it is therefore plausible that the development and evaluation of optimization algorithms will remain an important part of modelling in biology for the foreseeable future.

References

1. Allman, E.S., Rhodes, J.A.: *Mathematical Models in Biology*. Cambridge University Press, Cambridge (2003)
2. Banga, J.R.: Optimization in computational systems biology. *BMC Syst. Biol.* **2**, 47 (2008). doi:[10.1186/1752-0509-2-47](https://doi.org/10.1186/1752-0509-2-47)
3. Berg, J.M., Tymoczko, J.L., Stryer, L.: *Biochemistry*, 5th edn. W. H. Freeman, New York (2002)
4. Bishop, C.M.: *Neural Networks for Pattern Recognition*, 1st edn. Oxford University Press, New York (1996)
5. Brännmark, C., Palmér, R., Glad, S.T., Cedersund, G., Strålfors, P.: Mass and information feedbacks through receptor endocytosis govern insulin signaling as revealed using a parameter-free modeling framework. *J. Biol. Chem.* **285**(26), 20171–20179 (2010)
6. Britton, N.F.: *Essential Mathematical Biology*, p. 335. Springer, New York (2005)
7. Cedersund, G.: Prediction uncertainty estimation despite unidentifiability: an overview of recent developments. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)

8. Cedersund, G.: Conclusions via unique predictions obtained despite unidentifiability-new definitions and a general method. *FEBS J.* **279**(18), 3513–3527 (2012). doi:[10.1111/j.1742-4658.2012.08725.x](https://doi.org/10.1111/j.1742-4658.2012.08725.x)
9. Chaouiya, C., Berenguier, D., Keating, S.M., Naldi, A., van Iersel, M.P., Rodriguez, N., Dräger, A., Büchel, F., Cokelaer, T., Kowal, B., Wicks, B., Gonçalves, E., Dorier, J., Page, M., Monteiro, P.T., von Kamp, A., Xenarios, L., de Jong, H., Hucka, M., Klamt, S., Thieffry, D., Le Novère, N., Saez-Rodriguez, J., Helikar, T.: SBML Qualitative Models: a model representation format and infrastructure to foster interactions between qualitative modelling formalisms and tools *BMC Syst. Biol.* **7**(135), (2013)
10. De Jong, H.: Modeling and simulation of genetic regulatory systems: a literature review. *J. Comput. Biol.* **9**(1), 67–103 (2002)
11. De Jong K.: Genetic algorithms: a 30 year perspective. In: Booker, L., Forrest, S., Mitchell, M., Riolo, R. (eds) *Perspectives on Adaptation in Natural and Artificial Systems*, Oxford University Press, New York (2005)
12. De Jong, K.A.: *Evolutionary computation: a unified approach*. MIT Press, Cambridge (2006)
13. Dorigo, M., Birattari, M., Stützle, T.: Ant colony optimization—artificial ants as a computational intelligence technique. *IEEE Comput. Intell. Mag.* **1**, 28–39 (2006)
14. Egea, J.: *New heuristics for global optimization of complex bioprocesses*. Dissertation, University de Vigo (2008)
15. Egea, J., Rodriguez-Fernandez, M., Banga, J., Martí, R.: Scatter search for chemical and bio-process optimization. *J. Glob. Optim.* **37**(3), 481–503 (2007)
16. Garny, A., Nickerson, D., Cooper, J., Weber dos Santos, R., Miller, A.K., McKeever, S., Nielsen, P., Hunter, P.J.: CellML and associated tools and techniques. *Philos. Trans. R. Soc. A* **366**(1878), 3017–3043 (2008)
17. Gerstner, W., Sprekeler, H., Deco, G.: Theory and simulation in neuroscience. *Science (New York, N.Y.)* **338**(6103), 60–65 (2012). doi:[10.1126/science.1227356](https://doi.org/10.1126/science.1227356)
18. Glover, F.: Heuristics for integer programming using surrogate constraints. *Decis. Sci.* **8**(1), 156–166 (1977)
19. Glover, F., Kochenberger, G.A.: *Handbook of metaheuristics 57*. International Series in Operations Research & Management Science, Springer, New York (2003)
20. Goldberg, D.E.: *Genetic Algorithms in Search, Optimization and Machine Learning*, 1st edn. Addison-Wesley Longman Publishing Co., Inc, Boston (1989)
21. Gomez-Cabrero, D., Ardid, S., Cano-Colino, M., Tegnér, J., Compte, A.: Neuroswarm: a methodology to explore the constraints that function imposes on simulation parameters in large-scale networks of biological neurons. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
22. Gomez-Cabrero, D., Compte, A., Tegnér, J.: Workflow for generating competing hypothesis from models with parameter uncertainty. *Interface Focus* **1**(3), 438–449 (2011). doi:[10.1098/rsfs.2011.0015](https://doi.org/10.1098/rsfs.2011.0015)
23. Greenberg, H.J., Hart, W.E., Lancia, G.: Opportunities for combinatorial optimization in computational biology. *INFORMS J. Comput.* **16**(3), 211–231 (2004)
24. Gustafsson, M., Hörnquist, M., Lundström, J., Björkegren, J., Tegnér, J.: Reverse engineering of gene networks with LASSO and nonlinear basis functions. *Ann. N.Y. Acad. Sci.* **1158**, 265–275 (2009). doi:[10.1111/j.1749-6632.2008.03764.x](https://doi.org/10.1111/j.1749-6632.2008.03764.x)
25. Hart, W.E.: *Adaptive global optimization with local search*. Doctoral Dissertation, University of California, San Diego (1994)
26. Hodgkin, A.L.: Chance and Design in Electrophysiology: An informal account of certain experiments on nerve carried out between 1934 and 1952. *J. Physiol.* **263**(I), 1–21 (1976)
27. Hodgkin, A.L., Huxley, A.F.: Currents carried by sodium and potassium ions through the membrane of the giant axon of *Loligo* This information is current as of January 29, This is the final published version of this article; it is available at: this version of the article may not be. *J. Physiol. (Paris)* **116**, 449–472 (1952)
28. Holland, J.H.: *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. University of Michigan Press, Ann Arbor (1975)

29. Hoops, S., Sahle, S., Gauges, R., Lee, C., Pahle, J., Simus, N., Singhal, M., Xu, L., Mendes, P., Kummer, U.: COPASI—a COMplex PATHway SIMulator. *Bioinformatics* **22**(24), 3067–3074 (2006)
30. Hopfield, J.J.: Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U.S.A.* **79**, 2554–2558 (1982)
31. Hübner, K., Sahle, S., Kummer, U.: Applications and trends in systems biology in biochemistry. *FEBS J.* **278**(16), 2767–2857 (2011)
32. Kennedy, J., Eberhart, R.C.: Particle swarm optimization. In: *Proceedings IEEE international conference on neural networks*, Piscataway, pp. 1942–1948 (1995)
33. Kennedy J, Eberhart R.C.: (2001) *Swarm Intelligence*. Morgan Kaufmann Publishers, Massachusetts (2001)
34. Koza, J.R.: *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge (1992)
35. Kirkpatrick, S., Gelatt Jr, C.D., Vecchi, M.: Optimization by simulated annealing. *Science* **220**(4598), 671–680 (1983)
36. Larranaga, P.: Machine learning in bioinformatics. *Briefings Bioinf.* **7**(1), 86–112 (2006). doi:[10.1093/bib/bbk007](https://doi.org/10.1093/bib/bbk007)
37. Li, Z., Li, P., Krishnan, A., Liu, J.: Large-scale dynamic gene regulatory network inference combining differential equation models with local dynamic Bayesian network analysis. *Bioinformatics (Oxford, England)* **27**(19), 2686–2691 (2011). doi:[10.1093/bioinformatics/btr454](https://doi.org/10.1093/bioinformatics/btr454)
38. Mannakee, B.K., Ragsdale, A.P., Transtrum, M.K., Gutenkunst, R.N.: Sloppiness and the geometry of parameter space. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
39. Markowitz, F.: How to understand the cell by breaking it: network analysis of gene perturbation screens. *PLoS Comput. Biol.* **6**(2), e1000655 (2010). doi:[10.1371/journal.pcbi.1000655](https://doi.org/10.1371/journal.pcbi.1000655)
40. Martinez-Estudillo, A., Hervás-Martinez, C., Martinez-Estudillo, F., Garca-Pedrajas, N.: Hybrid method based on clustering for evolutionary algorithms with local search. *IEEE Trans. Syst. Man Cybernetics* **34**(1) (2004)
41. Maza, M., Yuret, D.: Dynamic hill climbing: overcoming the limitations of optimization techniques. In: *Proceedings of the 2nd Turkish Symposium on Artificial Intelligence and ANN*, 1993
42. Myers, C.R., Gutenkunst, R.N., Sethna, J.P.: Python unleashed on systems biology. *Comput. Sci. Eng.* **9**, 34 (2007)
43. Nyman, E., Brännmark, C., Palmér, R., Brugård, J., Nyström, F.H., Strålfors, P., Cedersund, G.: A hierarchical whole-body modeling approach elucidates the link between in Vitro insulin signaling and in Vivo glucose homeostasis. *J. Biol. Chem.* **286**(29), 26028–26041 (2011)
44. Pettersson, T.: *Global optimization methods for estimation of descriptive models*. Master’s thesis, Linköping University, Sweden (2008)
45. Poli, R.: Analysis of the publications on the applications of particle swarm optimisation. *J. Artif. Evol. Appl.* **2008**, 1–10 (2008)
46. Press, W., Teukolsky, S., Vetterling, W., Flannery, B.: *Numerical Recipes in C*, 2nd edn. Cambridge University Press, Cambridge (1992)
47. Roy, R., Dehuri, S., Cho, S.B.: A novel particle Swarm optimization algorithm for multi-objective combinatorial optimization problem. *Int. J. Appl. Metaheuristic Comput. (IJAMC)* **2**(4), 41–57 (2012)
48. Schlitt, T., Brazma, A.: Current approaches to gene regulatory network modelling. *BMC Bioinform.* **8**(Suppl 6), S9 (2007). doi:[10.1186/1471-2105-8-S6-S9](https://doi.org/10.1186/1471-2105-8-S6-S9)
49. Schmidt, H., Jirstrand, M.: Systems biology toolbox for MATLAB: a computational platform for research in systems biology. *Bioinformatics* **22**(4), 514–515 (2006)
50. Scholma, J., Schivo, S., Urquidí Camacho, R.,a, van de Pol, J., Karperien, M., Post, J.N.: Biological networks 101: computational modeling for molecular biologists. *Gene* **533**(1), 379–384 (2014). doi:[10.1016/j.gene.2013.10.010](https://doi.org/10.1016/j.gene.2013.10.010)
51. Segrè, D., Zucker, J., Katz, J., Lin, X., Haeseleer, P.D., Rindone, W.P., Church, G.M.: From Annotated Genomes to Metabolic Flux Models and Kinetic Parameter Fitting. *OMICS J. Integr. Biol.* **7**(3), 301–316 (2003)

52. Sejnowski, T.J., Koch, C., Churchland, P.S.: Computational Neuroscience. *Science* **241**, 1299–1306 (1987)
53. Skogsberg, J., Lundström, J., Kovacs, A., Nilsson, R., Noori, P., Maleki, S., Björkegren, J.: Transcriptional profiling uncovers a network of cholesterol-responsive atherosclerosis target genes. *PLoS Genet.* **4**(3), e1000036 (2008). doi:[10.1371/journal.pgen.1000036](https://doi.org/10.1371/journal.pgen.1000036)
54. Swameye, I., Müller, T.G., Timmer, J., Sandra, O., Klingmüller, U.: Identification of nucleocytoplasmic cycling as a remote sensor in cellular signaling by databased modeling. *Proc. Natl. Acad. Sci.* **100**, 1028–1033 (2003)
55. Triki, E., Collette, Y., Siarry, P.A.: Theoretical study on the behavior of simulated annealing leading to a new cooling schedule. *Eur. J. Oper. Res.* **166**(1 SPEC. ISS.), 77–92 (2005)
56. Van Geit, W., Achard, P., De Schutter, E.: Neurofitter: a parameter tuning package for a wide range of electrophysiological neuron models. *Front. Neuroinformatics* **1**, 1 (2007). doi:[10.3389/neuro.11.001.2007](https://doi.org/10.3389/neuro.11.001.2007)
57. Van Geit, W., De Schutter, E., Achard, P.: Automated neuron model optimization techniques: a review. *Biol. Cybern.* **99**(4–5), 241–251 (2008). doi:[10.1007/s00422-008-0257-6](https://doi.org/10.1007/s00422-008-0257-6)
58. Van Riel, N.W.: Dynamic modelling and analysis of biochemical networks: mechanism-based models and model-based experiments. *Briefings Bioinf.* **7**(4), 364–374 (2006). doi:[10.1093/bib/bbl040](https://doi.org/10.1093/bib/bbl040)
59. Wolpert, D.H., Macready, W.G.: No free lunch theorems for optimization. *IEEE Trans. Evol. Comput.* **1**, 67–82 (1997)
60. Zimmer, D.C.: Parameter estimation for stochastic models of biochemical reactions. *J. Comput. Sci. Syst. Biol.* **6**, 011–021 (2012). doi:[10.4172/jcsb.1000095](https://doi.org/10.4172/jcsb.1000095)

Chapter 8

Interval Methods

Warwick Tucker

Abstract We describe a modern approach to parameter estimation, based on set-valued computations combined with a branch and bound step. This allows us to examine entire sets of parameters, and thus to exhaust the global search within a finite number of steps. In addition, we show that the method can be accelerated by set-valued constraint propagation, which allows great improvement of its efficiency. To illustrate the applicability of our method we apply it to some networks of biochemical reactions modeled by a generic class of ODEs called Generalized Mass Action Models (GMAs).

Keywords Interval arithmetic · Set-valued methods · Directed rounding · Parameter estimation

8.1 Introduction

Interval methods [1, 11, 14, 17] refers to a framework designed to make numerical computations mathematically rigorous. This requires taking both rounding and discretization errors into account. The underlying arithmetic is set-valued; the size of the sets correspond to the current uncertainty in the computation at hand. Rather than computing approximations to sought quantities, the aim is to compute enclosures of the same. Today, the field has reached a high level of maturity, and the techniques have been used in solving many hard problems in mathematics in fields such as dynamical systems, control theory, and optimization.

W. Tucker (✉)

Department of Mathematics, Uppsala University, Uppsala, Sweden
e-mail: warwick@math.uu.se

© Springer International Publishing Switzerland 2016
L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_8

8.2 Interval Arithmetic

Here, we will briefly describe the fundamentals of interval analysis, which provides a simple way to extend real-valued computations to set-valued ones. For a concise reference on this topic, see e.g. [1, 11, 12].

Let \mathbb{IR} denote the set of closed intervals. For any element $\mathbf{a} \in \mathbb{IR}$, we adopt the notation $\mathbf{a} = [\underline{\mathbf{a}}, \overline{\mathbf{a}}]$. Thus “ $a \in \mathbf{a}$ ” means “the point a belongs to the interval \mathbf{a} ”. If \star is one of the operators $+$, $-$, \times , \div , we define the arithmetic on elements of \mathbb{IR} by

$$\mathbf{a} \star \mathbf{b} = \{a \star b : a \in \mathbf{a}, b \in \mathbf{b}\}, \quad (8.1)$$

except that $\mathbf{a} \div \mathbf{b}$ is undefined if $0 \in \mathbf{b}$. Working exclusively with closed intervals, we can describe the resulting interval in terms of the endpoints of the operands:

$$\begin{aligned} \mathbf{a} + \mathbf{b} &= [\underline{\mathbf{a}} + \underline{\mathbf{b}}, \overline{\mathbf{a}} + \overline{\mathbf{b}}] \\ \mathbf{a} - \mathbf{b} &= [\underline{\mathbf{a}} - \overline{\mathbf{b}}, \overline{\mathbf{a}} - \underline{\mathbf{b}}] \\ \mathbf{a} \times \mathbf{b} &= [\min(\underline{\mathbf{a}}\underline{\mathbf{b}}, \underline{\mathbf{a}}\overline{\mathbf{b}}, \overline{\mathbf{a}}\underline{\mathbf{b}}, \overline{\mathbf{a}}\overline{\mathbf{b}}), \max(\underline{\mathbf{a}}\underline{\mathbf{b}}, \underline{\mathbf{a}}\overline{\mathbf{b}}, \overline{\mathbf{a}}\underline{\mathbf{b}}, \overline{\mathbf{a}}\overline{\mathbf{b}})] \\ \mathbf{a} \div \mathbf{b} &= \mathbf{a} \times [1/\overline{\mathbf{b}}, 1/\underline{\mathbf{b}}], \quad \text{if } 0 \notin \mathbf{b}. \end{aligned} \quad (8.2)$$

When computing with finite precision, directed rounding must also be taken into account (see e.g. [11, 12]). It follows from (8.1)—or from (8.2)—that addition and multiplication are both associative and commutative: for $\mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{IR}$, we have

$$\begin{aligned} \mathbf{a} + (\mathbf{b} + \mathbf{c}) &= (\mathbf{a} + \mathbf{b}) + \mathbf{c}; & \mathbf{a} + \mathbf{b} &= \mathbf{b} + \mathbf{a}, \\ \mathbf{a} \times (\mathbf{b} \times \mathbf{c}) &= (\mathbf{a} \times \mathbf{b}) \times \mathbf{c}; & \mathbf{a} \times \mathbf{b} &= \mathbf{b} \times \mathbf{a}. \end{aligned}$$

Also, it is clear that the elements $[0, 0]$ and $[1, 1]$ are the unique neutral elements with respect to addition and multiplication, respectively. Note, however, that in general an element in \mathbb{IR} has no additive or multiplicative inverse. For example, we have $[1, 2] - [1, 2] = [-1, 1] \neq [0, 0]$, and $[1, 2] \div [1, 2] = [\frac{1}{2}, 2] \neq [1, 1]$. As a consequence of the arithmetic rules, the distributive law does *not* always hold. As an example, we have

$$[-1, 1]([-1, 0] + [3, 4]) = [-1, 1][2, 4] = [-4, 4],$$

whereas

$$[-1, 1][-1, 0] + [-1, 1][3, 4] = [-1, 1] + [-4, 4] = [-5, 5].$$

This unusual property is important to keep in mind when representing functions as part of an interval calculation. Interval arithmetic satisfies a weaker rule than the distributive law, which we shall refer to as *sub-distributivity*:

$$\mathbf{a}(\mathbf{b} + \mathbf{c}) \subseteq \mathbf{ab} + \mathbf{ac}. \quad (8.3)$$

This is a set theoretical property that illustrates one of the fundamental differences between real- and interval arithmetic. A key feature of interval arithmetic is that it is *inclusion monotonic*, i.e., if $\mathbf{a} \subseteq \mathbf{x}$, and $\mathbf{b} \subseteq \mathbf{y}$, then

$$\mathbf{a} \star \mathbf{b} \subseteq \mathbf{x} \star \mathbf{y}, \quad (8.4)$$

where we demand that $0 \notin \mathbf{y}$ for division. This property is what allows us to model real quantities with a finite amount of information, as described in the next section.

8.3 Set-Valued Numerical Methods

One of the main reasons for passing to the interval realm is that we want a simple way of enclosing the *range* $R(f; D) = \{f(x) : x \in D\}$ of a real-valued function $f : D \rightarrow \mathbb{R}$. Except for the most trivial cases, mathematics provides few tools to describe this set.

We begin by extending the real functions to *interval functions*. By this, we mean functions that take and return intervals rather than real numbers. Interval arithmetic (8.2) provides the theory of extending rational functions, i.e., functions on the form $f(x) = p(x)/q(x)$, where p and q are polynomials. Simply substituting all occurrences of the real variable x with the interval variable \mathbf{x} (and the real arithmetic operators with their interval counterparts) produces a rational interval function $F(\mathbf{x})$, called the *natural interval extension* of f . As long as no singularities are encountered, we have the inclusion

$$R(f; \mathbf{x}) \subseteq F(\mathbf{x}), \quad (8.5)$$

by property (8.4).

Example 1 Let $f(x) = 7/(3 + x)$, and consider the domain $x \in [0, 100]$. By the inclusion property (8.5), we have the following range enclosure:

$$R(f; [0, 100]) \subseteq F([0, 100]) = \frac{7}{3 + [0, 100]} = \frac{7}{[3, 103]} = \left[\frac{7}{103}, \frac{7}{3} \right].$$

In this particular case, the enclosure is *sharp*, i.e., we have $R(f; [0, 100]) = F([0, 100])$.

In fact, this type of range enclosure can be achieved for any reasonable function, e.g. $e^{\mathbf{x}} = [e^{\underline{x}}, e^{\overline{x}}]$ and so on. The only property we demand from the interval extension of a function f is that $R(f; \mathbf{x}) \subseteq F(\mathbf{x})$.

Higher-dimensional functions $f: \mathbb{R}^n \rightarrow \mathbb{R}$ can be extended to an interval function $F: \mathbb{IR}^n \rightarrow \mathbb{IR}$ in a similar manner. The function argument is then an *interval-vector* $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, which we also refer to as a *box*.

Example 2 Consider $f(x_1, x_2) = x_1 e^{x_2 - x_1}$, on the domain $(x_1, x_2) \in ([0, 1], [3, 4])$. By the inclusion property (8.5), we have the following range enclosure:

$$\begin{aligned} R(f; ([0, 1], [3, 4])) &\subseteq F([0, 1], [3, 4]) \\ &= [0, 1]e^{[3,4]-[0,1]} = [0, 1]e^{[2,4]} = [0, 1][e^2, e^4] = [0, e^4], \end{aligned}$$

which again happens to be sharp.

Interval analysis can be extended to an *exception-free* system by redefining the interval extension process. Without going into details, this makes *all* operations well-defined (including division by zero), see [23]. This extension is very useful in some situations where e.g. logarithms of sets containing negative numbers are repeatedly encountered.

There exist several open source programming packages for interval analysis [2, 4, 15], as well as commercial products such as [5, 7].

8.4 Interval Bisection

As a simple illustration of the powers of interval analysis, we will study the bisection method. This is a well-known algorithm for locating a zero of a continuous function. To be precise, let f be continuous on $[a, b]$, and suppose that $f(a)f(b) < 0$. Then, by the Intermediate Value Theorem, f has at least one root $\zeta \in (a, b)$. The bisection method proceeds as follows: Initially, we set $a_0 = a$ and $b_0 = b$. At stage k , we compute the midpoint $c_k = (a_k + b_k)/2$. Now there are three possibilities. If $f(c_k) = 0$, then we can set $\zeta = c_k$, and terminate the search. If $f(a_k)f(c_k) < 0$, we set $a_{k+1} = a_k$ and $b_{k+1} = c_k$. If $f(a_k)f(c_k) > 0$, we set $a_{k+1} = c_k$ and $b_{k+1} = b_k$. The search is guaranteed to converge to a zero of f since we have $|a_k - b_k| = 2^{-k}|a_0 - b_0|$.

There are, however, several flaws with the bisection method, when used as a zero-finding algorithm. One is the problem of finding points a and b satisfying the starting condition $f(a)f(b) < 0$. A second problem occurs when f has several zeros within the search domain. Suppose that f has N simple zeros $\zeta_1 < \zeta_2 < \dots < \zeta_N$ in $[a, b]$. Then the bisection method will find the even-labeled zeros with probability zero, whilst the odd-labeled zeros are located with uniform probability, see [3].

The interval bisection methods deals elegantly with both problems. Instead of aiming directly at finding a zero of f , it discards subsets of $[a, b]$ on which f is *guaranteed* to be non-zero. Such subsets are called *inconsistent*, and are removed from the search as early as possible. By using 8.5, it follows that $0 \notin F(\mathbf{x}) \Rightarrow 0 \notin R(f; \mathbf{x})$. Therefore, the strategy of the interval version bisection scheme is to

```

void bisect(pfcn F, interval x, double tol) {
    if ( subset(0.0, F(x)) ) { // If zero is contained in F(x)
        if ( width(x) < tol ) // ... and the tolerance is met
            cout << x << endl; // ... print the subinterval.
        else { // Otherwise, divide and conquer.
            bisect(F, interval(min(x), mid(x)), tol);
            bisect(F, interval(mid(x), max(x)), tol);
        }
    }
}

```

Fig. 8.1 A recursive implementation of the interval-valued bisection method

recursively bisect the search space, retaining only those subintervals \mathbf{x}_i satisfying $0 \in F(\mathbf{x}_i)$. Such intervals are called *indefinite*, since they *may* contain zeros of f . Once a indefinite subinterval has reached a width smaller than the tolerance tol , it is sent to the output. A C++ implementation of the interval-valued bisection method is presented in Fig. 8.1. When the search has been exhausted, we are left with a collection of indefinite intervals $\mathbf{x}_1, \dots, \mathbf{x}_M$ whose union contains *all* zeros of f within $[a, b]$:

$$\mathbb{S} = \{\zeta \in [a, b]: f(\zeta) = 0\} \subseteq \bigcup_{i=1}^M \mathbf{x}_i = \bar{\mathbb{S}}.$$

Of course, the set $\bar{\mathbb{S}}$ will overestimate \mathbb{S} . We can, however, expect a very good agreement between $\bar{\mathbb{S}}$ and \mathbb{S} , as long as the tolerance tol is kept reasonably small.

8.5 Set-Inversion


In the set-valued framework it makes a lot of sense to extend the concept of solving equations of the form $f(x) = y$ to include solving for inclusions $f(x) \in \mathbf{y}$, where \mathbf{y} is a given *set*. This is a very natural problem formulation in the presence of uncertainties; think of \mathbf{y} as representing noisy data with error bounds.

If \mathbf{y} is a set with non-empty interior, we should expect the solution set \mathbb{S} to share this property. Therefore, we should be able to approximate the solution set both from the inside and the outside. In other words, we would like to compute two sets $\underline{\mathbb{S}}$ and $\bar{\mathbb{S}}$ that satisfy

$$\underline{\mathbb{S}} \subseteq \mathbb{S} \subseteq \bar{\mathbb{S}}.$$

$\underline{\mathbb{S}}$ and $\bar{\mathbb{S}}$ are called the *inner* and *outer* approximations of \mathbb{S} , respectively. By measuring the size of their difference $\bar{\mathbb{S}} \setminus \underline{\mathbb{S}}$ we can obtain reliable information about how close we are to the solution set \mathbb{S} .

Fig. 8.2 An outer approximation $\bar{\mathbb{S}}$ of the solution set \mathbb{S} (shaded). All rectangles of width greater than 10^{-2} belong to the inner approximation $\underline{\mathbb{S}}$



Given a partition $\mathcal{P}(\mathbb{X})$ of the domain \mathbb{X} , the outer approximation $\bar{\mathbb{S}}$ contains all partition elements whose interval image has nonempty intersection with the range \mathbf{y} . The inner approximation contains all partition elements whose interval image is contained in the range \mathbf{y} . In other words, we have

$$\underline{\mathbb{S}} = \{\mathbf{x} \in \mathcal{P}(\mathbb{X}) : F(\mathbf{x}) \subseteq \mathbf{y}\} \quad \bar{\mathbb{S}} = \{\mathbf{x} \in \mathcal{P}(\mathbb{X}) : F(\mathbf{x}) \cap \mathbf{y} \neq \emptyset\}.$$

Example 3 Consider the nonlinear function

$$f(x) = \sin x_1 + \sin x_2 + \frac{2}{5}(x_1^2 + x_2^2),$$

and suppose we want to find the set

$$\mathbb{S} = \{x \in [-5, +5]^2 : f(x) \in [-0.5, 0.5]\}.$$

By a simple bisection procedure, we adaptively partition the domain $\mathbb{X} = [-5, 5]^2$ into subrectangles, and discard rectangles \mathbf{x} such that $F(\mathbf{x}) \cap [-0.5, 0.5] = \emptyset$. The remaining rectangles are classified according to whether they belong to $\underline{\mathbb{S}}$ and/or $\bar{\mathbb{S}}$. Note that, as soon as a subrectangle is determined to belong to $\underline{\mathbb{S}}$, it undergoes no further bisection. Only subrectangles whose interval images intersect $\partial\mathbb{S}$ are subdivided. This is illustrated in Fig. 8.2, where we used a stopping tolerance of 10^{-2} .

8.6 Parameter Estimation

The ability to solve nonlinear inclusions has great practical applications for parameter estimation. Given a finitely parameterized model function $f(x; p) = y$ together with a set of uncertain data $(x_1, \mathbf{y}_1), \dots, (x_n, \mathbf{y}_n)$ and a search space \mathbb{P} , the task is to solve for the set

$$\mathbb{S} = \{p \in \mathbb{P} : f(x_i; p) \in \mathbf{y}_i, i = 1, \dots, n\}.$$

This can be done by computing inner and outer approximations of \mathbb{S} , exactly as described above. The inner and outer approximations then become:

$$\begin{aligned}\underline{\mathbb{S}} &= \{\mathbf{p} \in \mathcal{P}(\mathbb{P}) : f(x_i; \mathbf{p}) \subseteq \mathbf{y}_i \text{ for all } i = 1, \dots, n\}, \\ \overline{\mathbb{S}} &= \{\mathbf{p} \in \mathcal{P}(\mathbb{P}) : f(x_i; \mathbf{p}) \cap \mathbf{y}_i \neq \emptyset \text{ for all } i = 1, \dots, n\}.\end{aligned}$$

Rather than assigning probability densities to model uncertainties, set-inversion methods use compact sets of numbers, e.g., intervals as its basic elements. This results in a less complex approach than the Bayesian one, as only boundedness of the errors is assumed—no prior information about probability distributions enter the computations. The ability to compute with entire sets of parameters means that any given search space can be (recursively) subdivided into smaller subsets—each of which can be examined independently of the others. The fact that the subsets form a partition of the original search space means that we can perform a *global* search: every single parameter is accounted for. This subdivision process is known as a branch-and-bound (BaB) method.

BaB search algorithms are deterministic global search algorithms that repeatedly partition (branch) the search space, and discard (bound) inconsistent parts. Set-valued computations allows safe identification of parameter sets that contain only *inconsistent* values and of those that contain only *consistent* values. Parameter regions that are not identified as consistent or inconsistent are called *indefinite*. Only the latter type are retained for further subdivision. We illustrate this procedure in the following example.

Example 4 Consider the model function $f(x; p_1, p_2) = 5e^{-p_1x} - 1 \times 10^{-6}e^{-p_2x}$ depending on two parameters p_1 and p_2 . We generate data at the nine sites $x_i = 5(i - 1)$, $i = 1, \dots, 9$ for the parameter values $p^* = (0.11, -0.32)$ with 90% relative noise added:

$$y_i = f(x_i; p_1^*, p_2^*), \quad \mathbf{y}_i = y_i(1 + [-0.9, +0.9]), \quad i = 1, \dots, 9.$$

This set-valued data is illustrated in Fig. 8.3a. Starting with the search space $\mathbb{P} = ([0, +100], [-100, 0])$, and using a stopping tolerance (the maximal side length of a parameter box) of 10^{-2} , we arrive at the inner and outer approximations illustrated in Fig. 8.3b.

While BaB methods work well for problems of low dimension, they are not practical for high-dimensional searches: the number of subsets grows large even for a small depth of the recursive subdivision process. To overcome this problem, *constraint propagation* (CP) techniques are used to shrink (or even entirely delete) subsets of the partition. CP-techniques use detailed structural information about the model function and data in order to propagate information between the parameter space and the data sets. This makes it possible to detect mathematical inconsistencies,

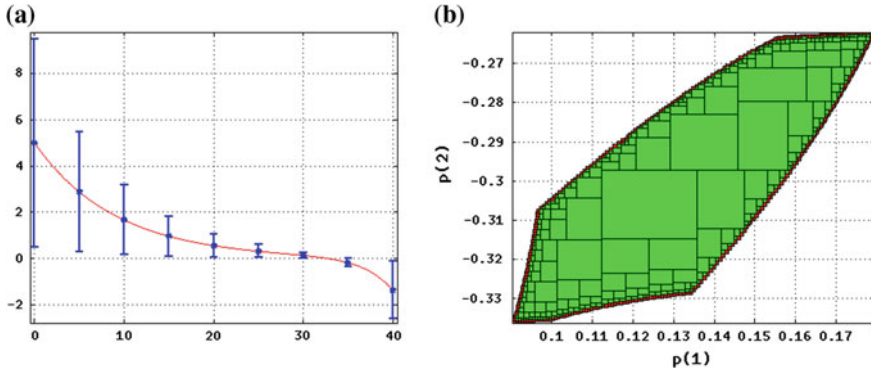


Fig. 8.3 **a** The set-valued data generated from the model function. **b** The corresponding set of consistent (green) and indefinite (red) parameters

which in turn can be used to shave off portions of parameters or data in an automated fashion. Combined with a BaB search, this produces a set-inversion method suited for large problems, see e.g. [9, 19, 20].

8.7 Constraint Propagation

The basic idea behind constraint propagation is to view a single equation $y = f(x)$ as a set of constraints that must be satisfied by the solution(s) of the original equation. Viewed as set-valued functions, these constraints often act as *contractors*, i.e., propagating the original search domain through the constraints often produces a smaller domain.

Example 5 As a first example, consider the equation $y = x_1^3 + x_2$, which can be recast as $x_1 = \sqrt[3]{y - x_2}$ and $x_2 = y - x_1^3$. Given a search domain $(\mathbf{x}_1, \mathbf{x}_2)$, we can impose the constraints $x_1 \in \mathbf{x}_1 \cap \sqrt[3]{y - \mathbf{x}_2}$ and $x_2 \in \mathbf{x}_2 \cap (y - \mathbf{x}_1^3)$ on all solutions. Taking $y = 2$, and looking for a solution with $(x_1, x_2) \in ([0, 1], [0, 1])$, the constraints produce

$$\begin{aligned}
 x_1 &\in \mathbf{x}_1 \cap \sqrt[3]{y - \mathbf{x}_2} = [0, 1] \cap \sqrt[3]{2 - [0, 1]} = [0, 1] \cap \sqrt[3]{[1, 2]} = [0, 1] \cap [1, \sqrt[3]{2}] = \{1\} \\
 x_2 &\in \mathbf{x}_2 \cap (y - \mathbf{x}_1^3) = [0, 1] \cap (2 - [0, 1]^3) = [0, 1] \cap (2 - [0, 1]) = [0, 1] \cap [1, 2] = \{1\}
 \end{aligned}$$

which actually happens to give the (unique) solution within the domain.

In general, we are not so fortunate as in this example, but the constraints do often contract the components of the search domain.

Example 6 As a second example, consider the equation $y = x_1 e^{x_2}$, which can be recast as $x_1 = y/e^{x_2}$ and $x_2 = \log(y/x_1)$. Given a search domain $(\mathbf{x}_1, \mathbf{x}_2)$, we can

impose the constraints $x_1 \in \mathbf{x}_1 \cap (y/e^{x_2})$ and $x_2 \in \mathbf{x}_2 \cap \log(y/\mathbf{x}_1)$ on all solutions. Taking $y = 1$, and looking for a solution with $(x_1, x_2) \in ([0, 2], [-1000, 0])$, the constraints produce

$$\begin{aligned} x_1 &\in \mathbf{x}_1 \cap (y/e^{x_2}) = [0, 2] \cap (1/[e^{-1000}, 1]) = [0, 2] \cap [1, e^{1000}] = [1, 2] \\ x_2 &\in \mathbf{x}_2 \cap \log(y/\mathbf{x}_1) = [-1000, 0] \cap \log(1/[1, 2]) = [-1000, 0] \cap [\log \frac{1}{2}, 0] = [\log \frac{1}{2}, 0]. \end{aligned}$$

This pass through the constraints contracted the domain by a factor ≈ 2885 , and happens to produce the smallest rectangle containing all solutions within the given domain.

In both examples, we have explicitly recast the equations into their inverse operations; in general this is not practical. Fortunately, the inversion of basic functions can easily be automated with a programming language supporting operator overloading. By iterating the contractions, in combination with a partitioning scheme, the solution set is rapidly obtained. Furthermore, if we somewhere in the process encounter an empty intersection, then we have proved that no solution exists within the given domain. For a thorough treatment of constraint propagation techniques, see [8] and references therein.

8.8 Examples

In this section we present two examples taken from modeling and simulation of biochemical systems. Both examples have model functions that are solutions to ordinary differential equations, and are as such not explicit. We overcome this complication by approximating the time-derivative \dot{x} by a suitable linear combination of the components of x . There are many ways to do this, resulting in different approximation schemes such as finite differences, collocation methods etc. We opt for a spline approach which works well when the data is not too noisy.

8.8.1 Example: A Generalized Mass Action Model

A *generalized mass action model* is a system of ordinary differential equations on the form

$$\dot{x}_i = \sum_{j=1}^{N_i} a_{ij} \prod_{k=1}^d x_k^{g_{ijk}} \quad (i = 1, \dots, d). \quad (8.6)$$

Each variable x_i (which assumes only positive values) represents the concentration of some reactant, and \dot{x}_i denotes the time derivative of x_i . The parameters a_{ij} are known as *rate constants*, whereas the parameters g_{ijk} are referred to as the *kinetic*




Fig. 8.4 **a** A branched pathway with feedback inhibitions. **b** The corresponding GMA-system

orders. The interactions corresponding to the system are usually summarized in the form of a network. We first consider the GMA-system presented in Fig. 8.4.

Here, the *topology* is assumed to be known, i.e., we know which parameters appear as non-zero quantities in the differential equation. Furthermore, we also use information regarding *dependencies*. More precisely, we know that the second term of the first component matches the first term of the second component, and that the third term of the first component matches the first term of the third component. Thus, all in all, we are to determine the values of 13 distinct parameters, arranged in two 5-dimensional problems, and one 3-dimensional problem.

For the computations, we used 10 sets of initial conditions, and each trajectory was sampled at 20 points in time. The search region for each kinetic order g_{ij} and rate constant a_i was formed by embedding each true parameter value in an interval with a radius proportional to the modulus of the true value, e.g., $\mathbf{a}_i = [(1 - \rho)a_i, (1 + \rho)a_i]$. The scaling factor ρ took values in $\{0, 0.1, 0.5, 1.0, 2.0, 3.0, 4.0, 5.0\}$. The stopping tolerance was set to 1×10^{-4} for the second and third component of the GMA-system. This produced parameter enclosures no wider than 4.61×10^{-4} . These were then inserted into the first component of the GMA-system, after which the remaining three parameters were solved for using the tolerance 1×10^{-3} . This setup produced the correct parameter values rounded to three significant digits.


In Table 8.1, we present the timings for the estimation, as well as the number of parameter boxes examined during the search.

We were surprised to note that the solution to the 3-dimensional subproblem (estimating the three leading parameters of the first component) was immediate. By this, we mean that the search only required one single pass through the constraints before satisfying the stopping tolerance. This resembles the situation presented in Example 5, illustrating the potential strength of using constraints as contraction mappings.

Table 8.1 The computational effort for the GMA-system

Relative radius (%)	0	10	50	100	200	300	400	500
CPU time (s)	0.02	3.99	5.08	5.68	6.88	8.50	10.77	13.63
Examined boxes	3	827	1067	1185	1463	1835	2341	3063

Fig. 8.5 **a** A branched pathway with activations and inhibitions. **b** The corresponding S-system



8.8.2 Example: An S-System

S-systems are a special class of GMAs having two terms per component; one corresponding to the production, and one corresponding to the degradation of the reactant.

An S-system can be considered as a condensed version of a GMA-system, obtained by aggregating individual reactions into the net processes of synthesis and degradation—see [16, 21] for a detailed account of this procedure. In particular, such a system consists of a system of ordinary differential equations of the form

$$\dot{x}_i = \alpha_i \prod_{k=1}^d x_k^{g_{ik}} - \beta_i \prod_{k=1}^d x_k^{h_{ik}} \quad (i = 1, \dots, d),$$

with non-negative rate constants α_i and β_i and real-valued kinetic orders g_{ij} and h_{ij} .

Our second example was described in [22], and falls into the category of S-systems (Fig. 8.5). Again, the topology is assumed to be known, i.e., we know which parameters appear as non-zero quantities in the differential equation. Contrary to the previous example, however, we no longer use any information regarding dependencies. That is, we do not use the fact that the second term of the second component matches the first term of the third component, and so on. Thus, all in all, we are to determine the values of 17 distinct parameters, arranged in three 4-dimensional problems, and one 5-dimensional problem.

For the computations, we used 5 sets of initial conditions, and each trajectory was sampled at 20 points in time. The search region for each kinetic order g_{ij} and rate constant a_i was formed by embedding each true parameter value in an interval with a radius proportional to the modulus of the true value, e.g., $\mathbf{a}_i = [(1 - \rho)a_i, (1 + \rho)a_i]$. The scaling factor ρ took values in $\{0, 0.1, 0.5, 1.0, 2.0, 3.0, 4.0, 5.0\}$. The stopping tolerance was set to 1×10^{-3} , which in all cases produced the correct parameter values rounded to three significant digits. In Table 8.2, we present the timings for the estimation, as well as the number of parameter boxes examined during the search.

Table 8.2 The computational effort for the S-system

Relative radius (%)	0	10	50	100	200	300	400	500
Time (in s)	0.00	2.22	2.94	3.24	3.55	3.83	4.03	4.36
Examined boxes	4	734	988	1072	1190	1292	1376	1494

In [18, 22], the search region for each of the kinetic orders g_{ij} was set to $[-1, +1]$, whereas the rate constants α_i were sought for within the domain $[0, 20]$. Using the same data set and stopping tolerance as above required a total running time of 4.09 s, assuming the correct topology, but using no prior information regarding dependencies. During the search a total of 1364 parameter boxes were examined. The estimated parameter values were, again, correct to at least three significant digits.

8.9 Discussion

Parameter estimation via set inversion and constraint propagation techniques differ in a fundamental way from the main-stream statistical-type parameter estimation methods in that we solve the problem by a pruning scheme based on a contraction principle, rather than recasting the estimation as a global minimization problem. This has several advantages: first, it is well-known that global minimization is an intractable problem, in the sense that numerical solutions often converge to a local, rather than a global, minimum, and there is no way of telling the two cases apart. Second, the quantity to be minimized is often chosen to be a (weighted) least-square error. This procedure distorts the search landscape, and may produce many seemingly good fits to the data. The set-valued methods described here simply contract/discard entire portions of the parameter space under scrutiny, according to their consistency with the underlying data, avoiding the problems mentioned above.

The limitation of interval methods is that they often are time-consuming compared to traditional numerical methods. This is perhaps not surprising considering that mathematical statements require more information—and therefore more processing of the available data—compared to local numerical solvers.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

1. Alefeld, G., Herzberger, J.: Introduction to Interval Computations. Academic Press, New York (1983)
2. CAPD—Computer Assisted Proofs in Dynamics Library, version 4.0. <http://capd.ii.uj.edu.pl/>
3. Corliss, G.: Which Root Does the Bisection Algorithm Find? *SIAM Rev.* **19**, 325–327 (1977)
4. CXSC—C++ eXtension for Scientific Computation, version 2.0. <http://www.math.uni-wuppertal.de/org/WRST/xsc/cxsc.html>
5. Forte Developer 7: C++ Interval Arithmetic Programming Reference. <http://docs.sun.com/app/docs/doc/816-2465>

6. Hammer, R., et al.: C++ toolbox for Verified Computing. Springer, Berlin (1995)
7. INTLAB—INTerval LABoratory, version 4.1.2. <http://www.ti3.tu-harburg.de/~rump/intlab/>
8. Jaulin, L., Kieffer, M., Didrit, O., Walter, E.: Applied Interval Analysis: With Examples in Parameter and State Estimation. Springer, Robust Control and Robotics (2001)
9. Johnson, T., Tucker, W.: Rigorous parameter reconstruction for differential equations with noisy data. *Automatica* **44**(9), 2422–2426 (2008)
10. Kulisch, U.W., Miranker, W.L.: Computer Arithmetic in Theory and Practice. Academic Press (1981)
11. Moore, R.E.: Interval Analysis. Prentice-Hall, Englewood Cliffs (1966)
12. Moore, R.E.: Methods and Applications of Interval Analysis. SIAM Studies in Applied Mathematics, Philadelphia (1979)
13. Moore, R.E., Kearfott, R.B., Cloud, M.J.: Introduction to Interval Analysis. SIAM, Philadelphia (2009)
14. Neumaier, A.: Interval Methods for Systems of Equations. Cambridge University Press, Cambridge (1990)
15. PROFIL/BIAS—Programmer’s Runtime Optimized Fast Interval Library/Basic Interval Arithmetic Subroutines. <http://www.ti3.tu-harburg.de/Software/PROFILEnglisch.html>
16. Torres, N.V., Voit, E.O.: Pathway analysis and optimization in metabolic engineering. Cambridge University Press, Cambridge (2002)
17. Tucker, W.: Validated numerics: A short introduction to rigorous computations. Princeton University Press, Princeton (2011)
18. Tucker, W., Moulton, V.: Reconstructing metabolic networks using interval analysis. *Lect. Notes Comp. Sci.* **3692**, 192–203 (2005)
19. Tucker, W., Moulton, V.: Parameter reconstruction for biochemical networks using interval analysis. *Reliab. Comput.* **12**(5), 389–402 (2006)
20. Tucker, W., Kotalik, Z., Moulton, V.: Estimating parameters for generalized mass action models using constraint propagation. *Math. Biosci.* **208**(2), 607–620 (2007)
21. Voit, E.O.: Computational analysis of biochemical systems. Cambridge University Press, (2000)
22. Voit, E.O., Almeida, J.: Decoupling dynamical systems for pathway identification from metabolic profiles. *Bioinformatics* **20**(11), 1670–1681 (2004)
23. Walster, W.G., Hansen, E.: Global Optimization Using Interval Analysis, vol. 264. CRC Press, Series in Pure and Applied Mathematics (2003)

Chapter 9

Model Extension and Model Selection

Mikael Sunnåker and Joerg Stelling

Abstract In this chapter we are concerned with the topic of construction, assessment, and selection of models in general, and of biochemical models in particular. Standard approaches to model construction and (automated) generation of candidate models are first discussed. We then present the most commonly used methods for model assessment, as well as the underlying concepts and ideas. In particular we focus on the information theoretic and Bayesian approaches to model selection. Information theoretic methods for model selection include the Akaike information criterion and the more recent deviance information criterion. Bayesian approaches include the computation of posterior ratios for relative model probabilities from Bayes factors as well as the approximate Bayesian information criterion. We also briefly discuss other methods such as cross-validation and bootstrapping techniques, and the theoretically appealing approach of minimum description length. We sketch how the most important results can be derived, emphasize distinctions between the methods, and discuss how model inference methods are employed in practice. We conclude that there is no generally applicable method for model assessment: a suitable choice depends on the specific inference problem, and to some extent also on the subjective preferences of the modeler.

Keywords Model selection · Bayesian information criterion · Akaike information criterion · Bayes factor · Sampling strategy

M. Sunnåker · J. Stelling (✉)

Department of Biosystems Science and Engineering, ETH Zurich, 4058 Basel, Switzerland
e-mail: joerg.stelling@bsse.ethz.ch

M. Sunnåker

e-mail: mikael.sunnaker@bsse.ethz.ch

M. Sunnåker · J. Stelling

Swiss Institute of Bioinformatics, ETH Zurich, 4058 Basel, Switzerland

M. Sunnåker

Competence Center for Systems Physiology and Metabolic Diseases, ETH Zurich,
8093 Zurich, Switzerland

© Springer International Publishing Switzerland 2016

L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_9

9.1 Introduction

Quantitative dynamical models constitute powerful tools for the detailed analysis of biochemical phenomena; this analysis may not be possible with more coarse-grained modeling strategies such as Boolean or Bayesian models. Examples of such phenomena include ultrasensitivity, multi-stability (bifurcations and hysteresis effects), oscillations, and signal amplification. The main purpose of constructing a biochemical model may in principle be to infer the correct structure of the modeled system, or to accurately predict the system response. However, accurate predictions require that the model incorporates the relevant system components, which can in turn only be inferred based on quantitative experimental observations of sufficient number and quality.

The inference of models that constitute valid representations of a given system is therefore one of the most important goals of computational systems biology (see also [31]). To infer a model we must be able to assess how well it represents the biological system, which requires a formal method for model assessment. Such methods are commonly based on an evaluation of the descriptive or predictive ability of the model, given a set of experimental data. The goal of the inference process is that the model should constitute the best possible representation of the studied system, for a given set of constraints. However, it may be difficult to check if this has been achieved in practice, or how far away we are from this goal. We therefore commonly contend ourselves with evaluating the relative quality of a predefined set of models. In a general scenario, therefore, we are concerned with close interactions between experimental and computational analysis using a set of mathematical models, with the (possible) aim of identifying the model that best represents the observed behavior of our biological system under study; this implies an iterative cycle as illustrated in Fig. 9.1.

This chapter serves as a guide for how to approach model inference from a theoretical and computational perspective. To make the chapter self-contained and to give the reader a head-start on the topic, we start with some basic notions and a brief discussion of the general model selection setting. We then discuss principles for how to convert verbal hypotheses about system mechanisms, which are usually available at the initial stages of a project, into a set of hypothetical mathematical models. This is followed by a discussion of different approaches and methods, as well as their most important technical aspects, for the assessment and selection of hypothetical models.

9.2 The Basic Setting

Before discussing the available methods for assessment and automated construction of biochemical models we will introduce some basic concepts that are useful. The components represented by a biochemical model are usually chemical species

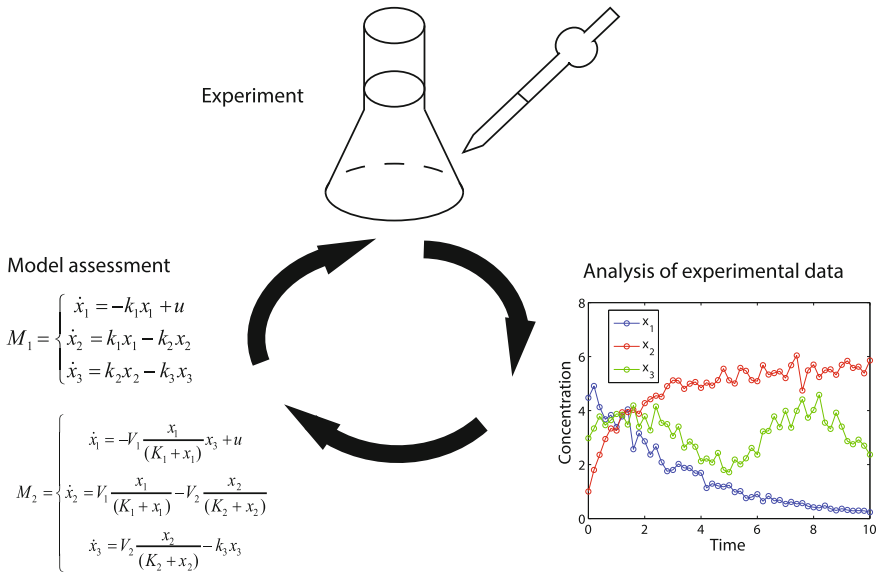


Fig. 9.1 Illustration of the ideal inference process for biochemical models. Given an initial set of data, a set of models, and an inference method, the models are assessed and compared. If the conclusions are ambiguous it may be necessary to produce additional experimental data, and the iterative process of experiments, data analysis, and model assessment is then continued until the results are satisfactory

(e.g., metabolites, proteins, transcripts) that may exist in different states (due to post-translational modifications, complex formations, etc.). The model represents the biochemical system by state variables that capture species abundances, and by the interactions between the state variables. The presence or absence of molecular interactions, referred to as biochemical reactions, defines the network structure and, in particular, the reaction stoichiometries. However, to simulate the time evolution of a model we also need to define the forms of the reaction terms. The reaction terms are typically parameterized (for example, when using enzyme kinetics to describe reaction rates), and the particular parameter values may strongly influence model predictions. The identification of representative reaction terms and parameter values requires experimental data containing information about the modelled entities. Importantly, the experimental data incorporated in the inference process must unambiguously correspond to model quantities, for example to certain combinations of the state variables.

The ingredients in the process of model assessment and selection are (see also Fig. 9.2):

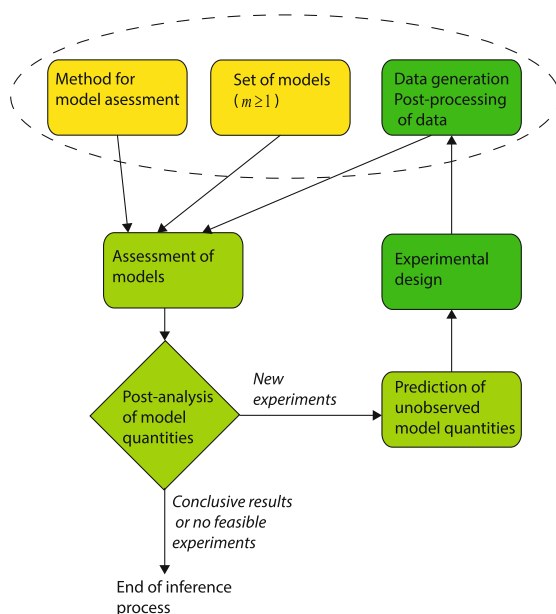


Fig. 9.2 Steps of the typical model assessment process. The three necessary initial components in the inference process are given encased in the dashed oval on *top*. The set of candidate models is then assessed, and the quantities of interest are analyzed. If the results are conclusive, or if no further (informative) experiments are feasible, the inference process comes to an end. Otherwise model predictions of unobserved model quantities are used to design informative experiments, and additional data is produced. We then iteratively follow the steps defined in the green boxes to incorporate, analyze, and generate new data until the outcome is satisfactory, or until no additional experiments are feasible

- A set of $m \geq 1$ testable dynamical models \mathcal{M} (the model class), corresponding to a set of m hypotheses.
- A set of experimental data \mathcal{D} , which takes the form of observations of one or more system entities over time (see also [45]).
- An inference method, that is, a formal procedure to draw conclusions about if a model constitutes a satisfactory system representation.

With these components at hand, we can assign a performance score to each of the models for the specified inference method and the available experimental data. System properties of interest, which are represented by the different models, may for example be the existence of hypothetical reactions, and the affinities or maximal velocities of reactions. There are two common scenarios resulting from the comparison of model behavior and experimental data. First, none of the models may reflect the system behavior accurately enough, for example, based on statistical tests of model fit. This indicates that important biochemical mechanisms are missing in the current set of hypotheses (models), and that the model class should be extended by new hypotheses. Note that it is a critical (and largely open) problem how to

incorporate such model extensions systematically. Second, the evidence from data may not provide conclusive support for one of the models, implying that we have a set of alternative hypotheses that can explain the data. In such cases we may investigate if the high-scoring models share certain properties, which are then assumed to well reflect the underlying system. If we are interested in model predictions, we may simulate more than one model and compute weighted predictions by employing the model scores. Another option, in particular if we want to find conclusive support for a single model, is to produce more experimental data. Ideally, computational analysis and data generation are then performed in an iterative fashion until the results are conclusive (see Fig. 9.1).

Example: Suppose that the number of transport proteins of a nutrient into a cell is unknown, and that we would like to use modeling to determine this number. The analysis is based on extra- and intracellular measurements of the nutrient. We first create a core model with well-established intracellular chemical species and reactions, and then a set of $m = 10$ extended models that each incorporates a different number (1–10) of transport reactions. To assess the models we decide to use the method of least squares in combination with a term that penalizes for the number of transport reactions (parameters). It turns out that the model with five transport reactions performs better than the models with fewer reactions, but that all models with at least five reactions can equally well explain the data. To discriminate between the remaining six models (with 5–10 reactions) we would need to collect more experimental data.

9.3 The Process of Inference in Biochemistry

9.3.1 Well Posed Modeling Problems

The purpose of a (biochemical) model should be to solve a well-posed problem, which may for example be a topology identification or prediction problem. A problem is well-posed in the sense of Hadamard if a unique solution to the problem exists such that the solution depends continuously on the data and parameters. The problem should be formulated in such a way that the investigator is in principle able to reach a conclusion by answering a sufficiently large number of questions with distinct alternatives (or equivalently with enough information). Examples of well posed problems in biochemistry are of the type: “is there a reaction between the chemical species X and Y?” and “will the average cytosolic concentration of metabolite Z reach above a given threshold upon addition of nutrients to a cell culture?”, whereas common examples of ill-posed problems are: “how does pathway P work?” and “what happens to metabolite Z upon nutrient addition?”. For an ill-posed problem

it is arbitrary, and depending on the subjective interpretation of the researcher, if a model is considered useful for solving the problem. An investigator should therefore strive to reformulate ill-posed problems, which are particularly common in the early stages of a project, into well-posed problems. The well-posed problems above give a hint to the areas of use of biochemical modelling. In general, biochemical models are tools that are used to infer the inner workings of systems from observational data, or to describe and predict the evolution of systems.

It is also important to define the scope of the model, since this may strongly influence the choice of modeling strategy. By the scope of the model, for example, we refer to the considered time span of the experiments, the granularity of the system representation (e.g., regarding the level of molecular resolution), and the type of system interventions covered (e.g., environmental or genetic conditions). It is important to discuss and communicate the model scope to all collaborators involved in the project—otherwise parts of the generated experimental data may not be useful.

9.3.2 Construction of Candidate Models

The initial step of the model construction process is to formulate, potentially still vague, ideas about solutions to the problem of interest as a set of (verbal) hypotheses. The set of hypotheses takes the form of a hypothesis class with m elements: $\mathcal{H} = \{H_1, H_2, \dots, H_m\}$. In the next step we reformulate the set of hypotheses into a set of mathematical descriptions that form a candidate model class: $\mathcal{M} = \{M_1, M_2, \dots, M_m\}$. Following the general guidelines proposed by Johnson et al. [28], there are in principle three steps in the conversion process between hypotheses and models:

- Specify the state variables of the model (e.g., chemical species in biochemical models).
- Specify the form of the interactions between variables (e.g., biochemical reactions) as well as the response variables (corresponding to experimental measurements).
- Define an error structure of the model (e.g., the form of the measurement noise).

Note that there should be a one-to-one mapping from a hypothesis to a model (i.e., a hypothesis should not be equally well described by two or more models), otherwise the hypothesis must be refined [28]. The selection of candidate hypotheses, and the corresponding candidate models to be evaluated, is an important process that should be done carefully. In fact, the difficulty to generate a suitable set of candidate models has commonly been criticized as a drawback of model selection methods by advocates of alternative approaches. For example, Steidl [50] argued that hypothesis testing is more informative than model selection for the analysis of ecological data when insufficient information is available to formulate strong candidate models, which generated some controversy in the literature [47, 49].

Here we focus on dynamical biochemical models in the form of systems of ordinary differential equations (ODEs). The ODEs of a model $M_i \in \mathcal{M}$ can be specified by:

$$\frac{d\mathbf{x}(t)}{dt} = f(\mathbf{x}(t), \boldsymbol{\theta}, \mathbf{u}(t)) \quad (9.1)$$

with (in general) a nonlinear function f , the state variables $\mathbf{x}(t) \in \mathbb{R}^n$, the parameters $\boldsymbol{\theta} \in \mathbb{R}_+^d$, and the inputs $\mathbf{u}(t) \in \mathbb{R}^q$. In biochemical models, the state variables typically represent the concentrations of chemical species, the parameters are for example used to specify the rates of reactions, and the inputs may for example represent the addition of a chemical compound to a cell culture. Note that the modeling target quantities \mathbf{Q} , which is the set of all available and potential types of experimental measurements to be used in the inference process, must have a clear interpretation and must be computable for each model in \mathcal{M} . We define a response function $\mathbf{y} \in \mathbf{Q}$ that relates the model components to experimental observations:

$$\mathbf{y}_k = h(\mathbf{x}_k) + \mathbf{e}_k, \quad \mathbf{e}_k \sim N(0, \mathbf{S}_k), \quad (9.2)$$

where the (potentially non-linear) function h maps system states to observables, the subscript $k = 1, \dots, N$ denotes observational time point t_k , $\mathbf{y}_k \in \mathbb{R}^m$, and where we assume that the measurement noise $\mathbf{e}_k \in \mathbb{R}^m$ is Gaussian distributed with covariance matrix $\mathbf{S}_k \in \mathbb{R}^{m \times m}$. In some cases it may be more appropriate to use a non-Gaussian measurement noise, although Gaussian distributions are commonly employed in practice (mainly for convenience reasons). In the following, we will omit the response function in the model definitions, since it is usually identical for all candidate models. Also, let the following observational data be available: $\mathcal{Y} = [\mathbf{Y}_1, \mathbf{Y}_2, \dots, \mathbf{Y}_N]$, where $\mathbf{Y}_k \in \mathbb{R}^m$.

The parts (components and interactions) of the biochemical models can be (operationally) divided into two categories: parts that are considered to be well-established and parts that are hypothetical. Typically we are interested in assessing the plausibility of the hypothetical parts, given the well-established core mechanisms. A model with only well-established, and no hypothetical mechanisms is referred to as the core model M_c : $\left\{ \frac{d\mathbf{x}^c(t)}{dt} = f^c(\mathbf{x}^c(t), \boldsymbol{\theta}^c, \mathbf{u}^c(t)) \right\}$. The model class \mathcal{M} comprises models that are constructed by combining the core model with a subset of the hypothetical model parts. The parts of model M_i may be decomposed into:

$$M_i: \left\{ \frac{d\mathbf{x}^i(t)}{dt} = \mathbf{f}^i(\mathbf{x}^i(t), \boldsymbol{\theta}^i, \mathbf{u}^i(t)) = \mathbf{f}^c(\mathbf{x}^c(t), \boldsymbol{\theta}^c, \mathbf{u}^c(t)) + \tilde{\mathbf{f}}^i(\mathbf{x}^i(t), \boldsymbol{\theta}^i, \mathbf{u}^i(t)) \right\}, \quad (9.3)$$

where $\tilde{\mathbf{x}}^i$, $\tilde{\boldsymbol{\theta}}^i$, and $\tilde{\mathbf{u}}^i$ are components in model M_i , but not in M_c , so that: $\mathbf{x}^i = \{\mathbf{x}^c, \tilde{\mathbf{x}}^i\}$, $\boldsymbol{\theta}^i = \{\boldsymbol{\theta}^c, \tilde{\boldsymbol{\theta}}^i\}$, and $\mathbf{u}^i = \{\mathbf{u}^c, \tilde{\mathbf{u}}^i\}$. Also note that $\tilde{\mathbf{f}}^i$ has the same arguments as \mathbf{f}^i in Eq.(9.3), since $\tilde{\mathbf{f}}^i$ may incorporate components of the core model. A number of $m + 1$ unique candidate models can be constructed if the hypothetical model parts are mutually exclusive, namely the core model and models M_1, M_2, \dots, M_m as defined

in Eq.(9.3). However, if the hypothetical model parts are not mutually exclusive, a number of: $\sum_{k=0}^m \binom{m}{k} = 2^m$ unique candidate models can be constructed of the form:

$$M_{\alpha(\Lambda)}: \left\{ \frac{d\mathbf{x}^\Lambda(t)}{dt} = \mathbf{f}^c(\mathbf{x}^c(t), \boldsymbol{\theta}^c, \mathbf{u}^c(t)) + \sum_{j=\Lambda} \tilde{\mathbf{f}}^j(\mathbf{x}^j(t), \boldsymbol{\theta}^j, \mathbf{u}^j(t)) \right. \quad (9.4)$$

where $\Lambda \subseteq \{1, 2, \dots, m\}$, $\alpha(\Lambda)$ maps the subset Λ to a specific integer between 1 and 2^m , and $\mathcal{M} = \{M_1, M_2, \dots, M_{2^m}\}$. The approach to combine a core model with a subset of hypothetical model extensions is commonly referred to as ensemble modeling [34]; for a recent application of ensemble modeling see [63].

9.3.2.1 Automated Construction of Dynamical Biochemical Models

The large effort required to manually construct a single model, the high computational cost connected to the evaluation of each model, and the potential combinatorial explosion in the number of hypothetical models make it difficult to cover a large part of the hypothesis space. Unfortunately, there is no generally accepted step-by-step strategy for model development, so the modeler must rely on experience and prior knowledge instead [14]. A sensible characterization of complex systems will most likely always necessitate a great deal of detective work and use of expert knowledge from the problem domain [53]. However, some recent approaches partially automate the model construction process, either by model reduction or by model extension.

To automate the generation of candidate models, one can start from a large model where parts are not necessary to explain the data. An interesting approach in this direction was proposed by Raue et al. [41]. This method is based on the computation of “profile likelihoods”, for which one of the model parameters is evaluated at a number of predefined values, and the rest of the parameters are optimized conditioned on the fixed parameter. This results in a confidence region for the evaluated parameter, and the parameter may then be eliminated if it is not significantly different from zero. However, this approach entails high computational costs because it requires a large number of optimizations to generate candidate models. To our knowledge this procedure has also not yet been tested in practice, but it was discussed as a potential model reduction approach [41]. A systematic heuristic approach for identifying reduced models based on exploring the parameter space has recently been proposed [52]. The effects of parameter elimination, or “evaporating parameters”, for gradient based optimization methods have been discussed by Transtrum et al. [56]. A method for the abstraction of detailed genetic regulatory network models was proposed by Kuwahara [35]. This method operates around a predefined set of model parameter values and applies different model reduction techniques (e.g., Michaelis-Menten approximation, dimerization reduction, and elimination of irrelevant nodes) to speed up stochastic simulations.

Alternatively, model extension approaches can be employed to generate candidate models. A recent method for automatic model extension by Schmidt et al. [37] uses a brute force approach based on symbolic regression (also see [3]) to search the combined space of model equations and parameter values with a genetic programming algorithm. Unfortunately, brute force approaches are often computationally expensive due to a combinatorial explosion in the number of terms to be evaluated, and as a consequence of high-dimensional parameter spaces. In addition, it may be difficult to interpret the resulting equations in terms of system components, and system constraints (e.g., mass balances of components) may be violated.

An interesting recent approach to automated model construction is that of [33], where stochastic model extensions are incorporated into models to represent uncertainties in model topologies ('system noise'). By introducing stochastic terms, one obtains models of the form:

$$M^s : \left\{ \frac{d\mathbf{x}(t)}{dt} = f(\mathbf{x}(t), \boldsymbol{\theta}, \mathbf{u}(t)) + G(\mathbf{x}(t), \boldsymbol{\theta}, \boldsymbol{\sigma}), \right. \quad (9.5)$$

where $G(\mathbf{x}(t), \boldsymbol{\theta}, \boldsymbol{\sigma})$ is a stochastic term, and $\boldsymbol{\sigma}$ is a nuisance parameter. The additional term can also be thought of as an extension to the structure of the ODE model in Eq. (9.1) [17]. Kristensen et al. proposed to let Eq. (9.5) take the form of a set of stochastic differential equations (SDEs) [33]. Equation (9.5) is then rewritten in differential form with $G(\mathbf{x}(t), \boldsymbol{\theta}, \boldsymbol{\sigma})dt = \boldsymbol{\sigma}d\boldsymbol{\omega}$:

$$M^{sde} : \{d\mathbf{x}(t) = f(\mathbf{x}(t), \boldsymbol{\theta}, \mathbf{u}(t)) dt + \boldsymbol{\sigma}d\boldsymbol{\omega} \quad (9.6)$$

By inferring the form of the additional stochastic term (that is, by estimation of $\boldsymbol{\sigma}$) it may be possible to identify missing deterministic model parts. However, the incorporation of this term may also improve the estimates of $\boldsymbol{\theta}$ and of other model components. A big challenge in SDE based model development is to automatize parts of the procedure. Similar tools for model diagnostics and improvements that are not based on SDEs have also been proposed [29], but those cannot be used to pinpoint model structure misspecifications [33].

9.3.2.2 Model Averaging and Model Expansion

There is no such thing as a correct, or true, model in the sense that the model captures every detail of the studied system. In fact, a model with this property would by definition not be a model, but equivalent to the system itself. A good model is useful because only characteristic features of the system, which have a significant impact on the model predictions, are incorporated. If the predictive power is of primary concern, it is therefore possible to average over candidate models, assuming that by averaging one will focus on the common, characteristic features.

Model averaging is a natural operation in the Bayesian setting. The posterior probability of a model M_i , among a set of models $\mathcal{M} = \{M_1, M_2, \dots, M_m\}$ and given a set of experimental data \mathcal{Y} , can be quantified by Bayes' theorem as:

$$p(M_i|\mathcal{Y}) = \frac{p(\mathcal{Y}|M_i)p(M_i)}{p(\mathcal{Y})} = \frac{p(\mathcal{Y}|M_i)p(M_i)}{\sum_{i=1}^m p(\mathcal{Y}|M_i)p(M_i)} \quad (9.7)$$

Equation(9.7) gives the posterior model probability in the simple case that M_i is not a parameterized model [30]. However, if M_i is a parameterized model we can integrate over the parameter space to compute:

$$p(M_i|\mathcal{Y}) = \frac{p(M_i)}{p(\mathcal{Y})} p(\mathcal{Y}|M_i) = \frac{p(M_i)}{p(\mathcal{Y})} \int_{\Theta} p(\mathcal{Y}|M_i, \boldsymbol{\theta}) p(\boldsymbol{\theta}|M_i) d\boldsymbol{\theta} \quad (9.8)$$

Note that $p(\mathcal{Y})$ is a constant that disappears when we compute the posterior ratio between models for model selection (see Sect.9.4.4), and it is commonly omitted from Eq.(9.8). To compute $p(M_i|\mathcal{Y})$ we therefore need to define the prior model probability $p(M_i)$, the prior distribution of the model parameters $p(\boldsymbol{\theta}|M_i)$, and the likelihood function $p(\mathcal{Y}|M_i, \boldsymbol{\theta})$.

We will now show how uncertainty can be incorporated in the model structure with arguments similar to those of Draper [17]. The inherent uncertainty in the model structure results in a model on the form: $M = (\boldsymbol{\theta}, \mathbf{S}, M_c)$, where M_c are fixed (core) parts of the model, $\boldsymbol{\theta} \in \Theta$ are model parameters, and $\mathbf{S} \in \Psi$ are nuisance parameters. Since M_c is fixed we will omit it in the following so that: $M = (\boldsymbol{\theta}, \mathbf{S})$. This gives us the following (continuous) set of models: $\mathcal{M}' = (\Theta, \Psi)$. The distribution for a target quantity $\mathbf{x} \in \mathbf{Q}(\mathcal{M}')$ can then be computed as:

$$\begin{aligned} p(\mathbf{x}|\mathcal{Y}) &= \int_{\mathcal{M}'} p(\mathbf{x}|\mathcal{Y}, M) p(M|\mathcal{Y}) dM \\ &= \int_{\Psi} \int_{\Theta} p(\mathbf{x}|\mathcal{Y}, \boldsymbol{\theta}, \mathbf{S}) p(\boldsymbol{\theta}, \mathbf{S}|\mathcal{Y}) d\boldsymbol{\theta} d\mathbf{S} \\ &\propto \int_{\Psi} \int_{\Theta} p(\mathbf{x}|\mathcal{Y}, \boldsymbol{\theta}, \mathbf{S}) p(\mathcal{Y}|\boldsymbol{\theta}, \mathbf{S}) p(\boldsymbol{\theta}, \mathbf{S}) d\boldsymbol{\theta} d\mathbf{S} \end{aligned} \quad (9.9)$$

Equation(9.9) can be used to compute statistical moments, for example the expected value of \mathbf{x} , for a continuous set of models. However, for most practical applications, the integrals in Eq.(9.9) cannot be solved analytically; we will discuss numerical approximations in Sect.9.6.

For a finite set of fixed model structures, $\mathcal{M} = \{M_1, \dots, M_m\}$, Eq.(9.9) takes the form:

$$p(\mathbf{x}|\mathcal{Y}) = \sum_{i=1}^m \int_{\Theta} p(\mathbf{x}|\mathcal{Y}, \theta_i, M_i) p(\theta_i, M_i|\mathcal{Y}) d\theta_i \propto \sum_{i=1}^m \int_{\Theta} p(\mathbf{x}|\mathcal{Y}, \theta_i, M_i) p(\mathcal{Y}|\theta_i, M_i) p(\theta_i, M_i) d\theta_i \quad (9.10)$$

Equation(9.10) shows how to compute the weighted (with posterior $p(M|\mathcal{Y})$) expected value of \mathbf{x} .

The posterior distribution of \mathbf{x} can be used to compute statistical properties of interest. For example, the first moment of \mathbf{x} , given model M , can be estimated as:

$$E[\mathbf{x}] = \int \mathbf{x} p(\mathbf{x}|\mathcal{Y}) d\mathbf{x} = \int \mathbf{x} \left(\frac{p(\mathcal{Y}|\mathbf{x})}{p(\mathcal{Y})} \right) p(\mathbf{x}) d\mathbf{x} \quad (9.11)$$

where $E[\mathbf{x}]$ denotes the expected value of \mathbf{x} (first moment). We can now compute a Monte Carlo estimate I_N of the integral in Eq.(9.11) by drawing $N \gg 1$ independent samples of $\mathbf{x}_i, i = 1, \dots, N$ from the prior distribution $p(\mathbf{x})$:

$$I_N = \sum_{i=1}^N \mathbf{x}_i w_i \quad (9.12)$$

where $w_i = \frac{p(\mathcal{Y}|\mathbf{x}_i)}{\sum_{i=1}^N p(\mathcal{Y}|\mathbf{x}_i)}$ [43]. Also note that it is straight-forward to use a parameter prior distribution instead of a prior distribution over \mathbf{x} , since there is a one-to-one mapping between the state variables and the parameters (in a deterministic model, θ_i unambiguously maps to \mathbf{x}_i).

Note that Eq.(9.10) is more commonly used in practical applications than Eq.(9.9); a finite set of distinct models is commonly the preferred choice in biochemical modeling. However, applications for which a continuous spectrum of models is at hand include the SDE models discussed in the previous section [33], and the nonlinear mixed effects models in [4].

9.4 Methods for Model Assessment and Model Selection

In this section we will review the most common computational approaches and methods for the assessment of, and selection between, (biochemical) models. We will in particular focus on the conceptual ideas and the connections between the approaches commonly used in the domain of dynamical (ODE) models. Most model selection methods that are not commonly used for dynamical models, such as minimum mes-

Table 9.1 Properties of some of the most common model selection methods

Method	Approach	Parameter points	Non-nested	Large N
Likelihood ratio test	HT	Single	No	Yes
AIC	IT	Single	Yes	Yes
DIC/BPIC	IT	Multiple	Yes	No
Bayes factors	Bayesian	Multiple	Yes	No
BIC	Bayesian	Single	Yes	Yes
Cross validation	Prediction	N/A	Yes	N/A

Abbreviations: *AIC* Akaike information criterion. *DIC* Deviance information criterion. *BIC* Bayesian information criterion. *BPIC* Bayesian predictive information criterion. *HT* Hypothesis testing. *IT* Information theoretic. *Approach* refers to the ideas underlying the method. *Parameter points* refers to if the conclusions are drawn from one (optimal) parameter point (see also [9]), or from an ensemble of parameter points. *Non-nested* refers to if the method works also for comparison of models that are not nested. *Large N* refers to if the method is only justified asymptotically in the limit of a large number of data points

sage length [13] and methods employing the Vapnik-Chervonenkis dimension [24], will therefore be omitted. For a more general discussion of issues commonly encountered in the modeling and inference process we refer the reader to [53].

The available approaches to model assessment can be assigned to two broad categories, depending on how they relate models and data. The prediction based approach evaluates the predictive power of the model after dividing the available data into two sets; a training set and a test set. Alternatively, in the likelihood based approach the data is not divided, but instead all available data is used in the assessment of each model. The corresponding methods are either based on a comparison between the assessed model and the data generating model (assumed to belong to the class of considered models), or between models within a pre-defined set of candidate models. Table 9.1 shows the key properties of the methods discussed here. For other reviews on the topic of model selection in computational biology see [6, 32, 61]. Note that the choice of model assessment approach is to some extent subjective. However, for a correct use of the approaches it is necessary to understand the underlying reasoning and assumptions. The steps of the typical model assessment process are summarized in Fig. 9.2.

Prediction Based Approaches

A straight-forward approach to investigate the predictive ability of a model is to divide the available data into two sets, namely one set of training data and one set of test data. The training data is used to calibrate the model parameters, and the test data is used to assess the predictive ability of the model. Note that this is different from the likelihood based approaches discussed later in this chapter, for which all data is used in the model assessment; this may be advantageous in situations for which few experimental observations are available.

Cross-validation can be used to estimate the expected prediction error, which then serves as a measure of the predictive ability of the model. The idea behind K -fold cross validation is to divide the N available data into $K \leq N$ equally sized parts, and to use all parts but one to train the model and the remaining part for model assessment. Let \mathcal{Y}^k denote all data in part k and \mathcal{Y}^{-k} denote all experimental data except for those in part $k = 1, \dots, K$. Let $\hat{\mathbf{y}}(\mathcal{Y}^{-k})$ be the prediction of \mathbf{y} for a model that has been calibrated on \mathcal{Y}^{-k} . An estimate of the expected prediction error, ε , is then:

$$\varepsilon = \frac{1}{N} \sum_{k=1}^K \sum_{i=1}^{N/K} L(Y_j, \hat{\mathbf{y}}_j(\mathcal{Y}^{-k})), j = (k - 1) \left(\frac{N}{K}\right) + i, \tag{9.13}$$

where Y_j is the j th experimental observation and $L(Y, \hat{\mathbf{y}}(\mathcal{Y}^{-k}))$ is a loss function for measuring errors (squared or absolute errors are commonly used) [24]. It is not immediately clear how large K should be. If we let $K = N$ (leave-one-out cross validation) the error estimation is as accurate as possible (for $K < N$ the error may be overestimated), but unfortunately this approach may be computationally expensive depending on the application [24]. It is therefore common to choose a $K < N$ in practice. For a given set of models, it is reasonable to select the model with the lowest expected prediction error. For a biochemical application of cross-validation, see the work by Kuepfer et al. [34] where 5-fold cross-validation was used to identify a predictive subset of models among the 18 models in the original hypothesis class.

The idea behind bootstrapping is to estimate the statistical accuracy of a quantity of interest by re-sampling B times from a set of data with replacement. In a fashion similar to that of cross-validation, the expected prediction error for a model can be computed as a leave-one-out bootstrap estimate (for each sample of size N):

$$\varepsilon = \frac{1}{N} \sum_{i=1}^N \frac{1}{|C_i|} \sum_{j \in C_i} L(Y_i, \hat{\mathbf{y}}_i^j), \tag{9.14}$$

where C_i are the indices of the bootstrap samples that do not contain observation i ($|C_i|$ is the number of such samples), and $\hat{\mathbf{y}}_i^j$ is the prediction of Y_i with the model trained on samples corresponding to bootstrap sample index j [18, 27]. Leave-one-out bootstrapping can be considered as a smoothed version of leave-one-out cross-validation, since for each i we average over multiple predictions made on the bootstrap samples in C_i [27]. At the same time, a bootstrap sample of size N contains only $\approx 0.632N$ distinct observations on average [27], resulting in a prediction error that tends to be overestimated. As for cross-validation, we select the model with the lowest expected prediction error from a given set of models, according to Eq. (9.14). However, in the commonly data-deprived domain of dynamical modeling in biochemistry it is difficult to find applications for which bootstrapping is used for model selection.

Another recent prediction based model validation strategy, which is rooted in information theory and particularly useful for model selection in clustering, is approximation set coding [5]. However, although biological applications have been reported [11], we are not aware of applications in the domain of dynamical models.

9.4.1 Hypothesis Testing: The Likelihood Ratio Test

Traditional hypothesis testing only evaluates evidence against a model, typically by integrating over the distribution of “possible experimental outcomes” and computing p-values (with the null hypothesis that the data was generated from the evaluated model). Correspondingly, p-values represent the probabilities of observing results at least as extreme as the experimental data, given that the model has the same data generating characteristics as the system. However, this may lead to situations in which all available models are rejected, which was expressed by Kass and Raftery [30] in the following statement: “There has not been as single date in the history of the law of gravitation when a modern significance test would not have rejected all laws and left us with no law.”

A commonly used approach for comparing two models is the likelihood ratio test. Assume that a model M_1 with d_1 free parameters is nested in a model M_2 with d_2 free parameters. By nested we mean that M_2 can be transformed into M_1 by imposing restrictions on a subset of the parameters in M_1 (e.g., by eliminating the parameters). Note that the fit of the more general model M_2 will always be at least as good as that of the nested model M_1 . We can then test whether the null hypothesis H_0 that M_1 best fits the data should be rejected in favor of the alternative hypothesis H_1 that M_2 best explains the data. Assuming Gaussian measurement noise, the test statistic takes the form:

$$R \equiv -2 \log \left(\frac{p(\mathbf{y}|\hat{\theta}_1, M_1)}{p(\mathbf{y}|\hat{\theta}_2, M_2)} \right) \sim \chi_{d_2-d_1}^2, \quad (9.15)$$

where $\hat{\theta}_1$ and $\hat{\theta}_2$ are maximum likelihood estimates, and $\chi_{d_2-d_1}^2$ is the chi-squared distribution with $d_2 - d_1$ degrees of freedom. Note that evidence against H_0 gives a small likelihood ratio (large R). An asymptotic (with the number of observations) rejection region for H_0 can therefore be computed from the $1 - \alpha$ quantile $\chi_{1-\alpha, d_2-d_1}^2$.

As a disadvantage of the likelihood ratio test, it requires that the models are nested, and it may also be clumsy for multiple nested models [32]. Hence it has been argued that the likelihood ratio test may not be “the optimal strategy for model selection” [40]. However, the likelihood ratio test, in combination with bootstrapping for non-nested models, was for example used by Muller et al. [39] in a two step process to analyze experimental data for the JAK-STAT signaling pathway. This analysis revealed that a cycling mechanism must be incorporated in the model to explain the data, as well as a delay mechanism for the nuclear entry of STAT5.

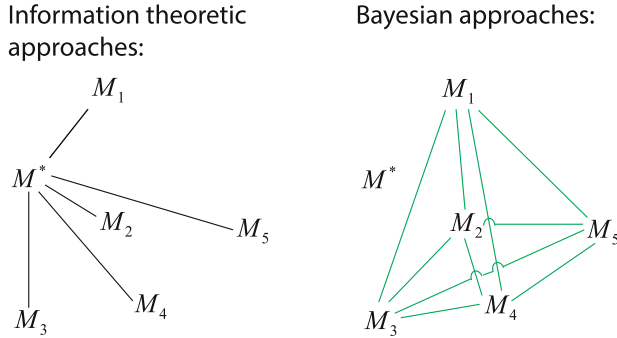


Fig. 9.3 The conceptual difference between the information theoretic and Bayesian approaches to model selection. *Left* The information theoretic approaches are based on a measure of the distance indicated by the *black lines* (from the marginal likelihoods) of each hypothetical model (M_{1-5}) to the best possible model (M^*). For example, model M_2 is most predictive since the distance to M^* is the shortest. *Right* Bayesian approaches are instead based on mutual comparisons between the hypothetical models (ratio of marginal likelihoods), indicated by the *green lines*, but M^* is not directly considered (color online)

9.4.2 Likelihood Based Approaches: The Information Theoretic versus the Bayesian Approach

There are also likelihood based methods for model assessment that can be applied to non-nested models. These methods simultaneously evaluate evidence both against and in favor of models [55]. The methods are also capable of weighing and ranking models, which enables model averaging for robust parameter estimations and predictions [28] (see also Sect. 9.3.2.2). There are in principle two main directions among the most popular approaches to model assessment; the information theoretic approach and the Bayesian approach. The conceptual difference between the two approaches is illustrated in Fig. 9.3.

The information theoretic approach assesses the information loss by using an approximate model M_i to analyze \mathcal{Y} instead of the data generating model (or process) M^* . The information in the data, given a model M and parameters θ , is given by $-\ln(p(\mathbf{y} = \mathcal{Y} | \theta, M))$, which results in the information difference:

$$\Delta(\theta_i, M_i) = \ln\left(\frac{1}{p(\mathcal{Y} | \theta_i, M_i)}\right) - \ln\left(\frac{1}{p(\mathcal{Y} | \theta^*, M^*)}\right), \quad (9.16)$$

where we assume that the correct model M^* belongs to the class of parameterized models with correct parameterization θ^* . Akaike [1] showed that the expected $\Delta(\theta_i, M_i)$, under certain assumptions, may be estimated from the parameter set that maximizes the likelihood function (given a set of data) for model M_i . This results in a model evaluation criterion that is commonly referred to as the Akaike information criterion (AIC). Other criteria that originate from Eq. (9.16) are the deviance

information criterion (DIC) and the related Bayesian predictive information criterion (BPIC), which require an integration over the parameter space of the evaluated model. We will discuss the information-theoretic criteria for model assessment in greater detail below. For a nice review about model selection using information theoretic approaches see also [6].

The Bayesian approach for assessment and selection of models (see also [25]) is to compute the posterior ratio in support for two models M_i and M_j as:

$$\frac{p(M_i|\mathcal{Y})}{p(M_j|\mathcal{Y})} = \frac{p(\mathcal{Y}|M_i) p(M_i)}{p(\mathcal{Y}|M_j) p(M_j)} \equiv B_{i,j} \frac{p(M_i)}{p(M_j)}, \quad (9.17)$$

where $B_{i,j}$ as defined in Eq. (9.17) is conventionally referred to as the Bayes factor. We also discuss an approximation of the Bayes factor, referred to as the Bayesian information criterion (BIC) below. At the moment there are high expectations on the Bayesian approach for applications in biochemical modeling, which are reflected in the statement by Wilkinson [61]: “Experience from closely related disciplines suggests that fully Bayesian approaches will turn out to provide the most satisfactory solutions to the complex statistical inference problems which lie at the heart of computational systems biology.”

9.4.3 Information Theoretic Approaches

9.4.3.1 Akaike Information Criterion

Assume that we want to assess the performance of a set of m models $\mathcal{M} = \{M_1, M_2, \dots, M_m\}$. The expected difference in information resulting from using a model $M_i \in \mathcal{M}$ and parameters θ_i , instead of the data generating model M^* and parameters θ^* , is:

$$I(\theta_i, M_i) \equiv D_{KL}[p(\mathbf{y}|\theta^*, M^*)||p(\mathbf{y}|\theta_i, M_i)] = E_*[\Delta(\theta_i, M_i)], \quad (9.18)$$

where E_* denotes the expectation under $p(\mathbf{y}|\theta^*, M^*)$, and D_{KL} denotes the Kullback-Leibler divergence. However, computing this quantity is not straight-forward in practice, mainly because the data generating model is unknown. We will now discuss the principles for how the expression of the Kullback-Leibler information may be approximated as suggested by Akaike [1] (see also [16]). Our discussion is similar to that by Cavanaugh [8].

We first multiply the expression in Eq. (9.18) with 2 (which is convenient for Gaussian likelihood functions):

$$2I(\theta_i, M_i) = 2E_* \left[\ln \frac{p(\mathbf{y}|\theta^*, M^*)}{p(\mathbf{y}|\theta_i, M_i)} \right] = 2E_*[\ln p(\mathbf{y}|\theta^*, M^*)] + E_*[-2 \ln p(\mathbf{y}|\theta_i, M_i)]. \quad (9.19)$$

We note that the first term in Eq. (9.19) is independent of the evaluated model and parameter (M_i and θ_i) and will be the same for all models in \mathcal{M} . For model assessment

we therefore only need to focus on the second term, which is a substitute for $I(\theta_i, M_i)$. For the substitute term we have that:

$$d(\theta_i, M_i) \equiv E_*[-2 \ln p(\mathbf{y}|\theta_i, M_i)] = \int_{-\infty}^{\infty} -2 (\ln p(\mathbf{y}|\theta_i, M_i)) p(\mathbf{y}|\theta^*, M^*) d\mathbf{y}, \tag{9.20}$$

which cannot be directly computed due to the unknown density $p(\mathbf{y}|\theta^*, M^*)$. However, Akaike's main achievement was to show that $-2 \ln p(\mathbf{y}|\hat{\theta}_i, M_i)$ can under certain conditions be employed as a biased estimator for $d(\theta_i, M_i)$, where $\hat{\theta}$ is the optimum of the likelihood function. To see this we can use the exact formulation (note that the terms on the right and left hand side cancel):

$$E_{\hat{\theta}}[d(\theta_i, M_i)] = E_{\hat{\theta}}[-2p(\mathbf{y}|\hat{\theta}_i, M_i)] + \underbrace{(E_{\hat{\theta}}[-2 \ln p(\mathbf{y}|M^*, \theta^*)] - E_{\hat{\theta}}[-2p(\mathbf{y}|\hat{\theta}_i, M_i)])}_{T_1} + \underbrace{(E_{\hat{\theta}}[d(\theta_i, M_i)] - E_{\hat{\theta}}[-2 \ln p(\mathbf{y}|M^*, \theta^*)])}_{T_2}. \tag{9.21}$$

where $E_{\hat{\theta}}$ denotes the expectation with respect to the distribution for $\hat{\theta}(\mathbf{y})$. Taylor expansions of T_1 (around $\hat{\theta}_i$) and T_2 (around θ^*) as defined in Eq. (9.21) to the second terms result in quadratic forms that both asymptotically converge to d (within the first order of approximation; for a detailed derivation see [1, 8, 16]). For a given set of data \mathcal{Y} , the Akaike information criterion (AIC) is then an unbiased estimator of Eq. (9.21):

$$AIC(\theta_i, M_i) \equiv -2p(\mathcal{Y}|\hat{\theta}_i, M_i) + 2d. \tag{9.22}$$

where $\hat{\theta}_i$ is the maximum likelihood estimate of θ_i given \mathcal{Y} . Note that we only need the optimum of the likelihood function ($\hat{\theta}_i$) to compute the AIC for a given model and a set of experimental data. This makes the criterion easy to handle and has contributed to its popularity.

In situations with small data sample sizes, a correction to the AIC was suggested, which is commonly referred to as the corrected Akaike information criterion (cAIC) [6, 26, 51]. For linear models the correction takes the form:

$$cAIC = AIC + \frac{2d(d+1)}{(N-d-1)} \tag{9.23}$$

where we note that $cAIC = AIC$ when $N \rightarrow \infty$. However, the AIC is more universally applicable than the cAIC, which is only justified for certain model classes (e.g., for linear regression) [8].

The (corrected) AIC can also be used to compute (Akaike) weights for the different models, which can be interpreted as probabilities for each model to be the best within the set of candidate models [60]. Let cAIC_{\min} denote the cAIC for the model that best fits the data (that is, it has the smallest cAIC among the evaluated models). Now let δ_i be the difference between the cAIC for model M_i and cAIC_{\min} . We can then compute Akaike weights, that is, the relative weight of model M_i among the models in \mathcal{M} for a uniform model prior, as in [6]:

$$w_i = p(M_i|\mathcal{Y}) = \frac{l_i}{\sum_{j=1}^m l_j}, \quad (9.24)$$

where $l_i = e^{-\frac{1}{2}\delta_i}$.

The AIC and the related cAIC appear to be some of the most commonly used criteria for model ranking in the literature. The cAIC was for example used by Turkheimer et al. [57] in combination with Akaike weights to compare three models of [11C]flumazenil kinetics in the brain for positron emission tomography data. In a recent example, Pezze et al. [15] used the AIC to compare three models of the mammalian target of rapamycin (TOR) pathway to study the unknown regulation of mTORC2 and its relation to the TSC1-TSC2 complex. These three models all turned out to be incompatible with the data, leading to the (manual) construction of a fourth model.

9.4.3.2 Deviance Information Criterion

Another model selection criterion based on information theory is the deviance information criterion (DIC) as proposed by Spiegelhalter et al. [48]. This criterion takes the form:

$$\text{DIC} = 2\bar{D} - D(\bar{\theta}) \quad (9.25)$$

with

$$\bar{D} \equiv E_{\theta|\mathcal{Y}}[D(\theta)] \quad (9.26)$$

and

$$D(\bar{\theta}) \equiv D(E_{\theta|\mathcal{Y}}[\theta]), \quad (9.27)$$

where $E_{\theta|\mathcal{Y}}$ denotes the expectation w.r.t. the posterior distribution $p(\theta|\mathcal{Y})$, and

$$D(\theta) = -2 \ln[p(y|\theta)] + C(y) \quad (9.28)$$

where $C(y)$ is a function of the data only (and not the parameters).

Note that this criterion is appropriate for sampling based methods, since it is straight-forward to compute $\bar{D}(\theta)$ and $D(\bar{\theta})$ from parameter samples. However, the DIC is much less frequently used in the systems biology literature than the AIC, which may be due to its recent introduction, but also likely results from the popularity of optimization based methods. The DIC has been criticized for the assumption that the true model belongs to the class of parameterized models, and that the observational data is used twice in the inference process [2]. Ando et al. [2] therefore proposed the Bayesian predictive information criterion to circumvent these issues in situations for which the sample size is large.

9.4.4 Bayesian Approach: Bayes Factors and BIC

The posterior ratio between models M_i and M_j , given experimental data, can be computed from the Bayes factor $B_{i,j}$ as defined in Eq. (9.17). However, for parameterized models the computation of $B_{i,j}$ requires integration over the parameter space of the two models according to:

$$B_{i,j} = \frac{p(\mathcal{Y}|M_i)}{p(\mathcal{Y}|M_j)} = \frac{\int_{\Theta_i} p(\mathcal{Y}|\theta, M_i)p(\theta|M_i)d\theta}{\int_{\Theta_j} p(\mathcal{Y}|\theta, M_j)p(\theta|M_j)d\theta}, \quad (9.29)$$

and we discuss strategies for the direct computation of such integrals in Sect. 9.6. However, all known sampling strategies share the unfortunate property that the computational cost is high and grows with the dimension of the sampled parameter space.

In some situations, Eq. (9.29) takes a significantly simpler form that can be derived from the Bayesian information criterion (BIC) (see also [30]). Assume that the posterior parameter distribution for model M_i is unimodal and approximately Gaussian around the maximum $\hat{\theta}$. This is often the case if the likelihood function is peaked due to a large samples size [30]. Let:

$$\tilde{l}(\theta) = \ln[p(\mathcal{Y}|\theta, M_i)p(\theta|M_i)], \quad (9.30)$$

and approximate \tilde{l} with a Taylor expansion around $\hat{\theta}$ up to the quadratic term:

$$\tilde{l}(\theta) \approx \tilde{l}(\hat{\theta}) + \frac{1}{2}(\theta - \hat{\theta})^T H_{\tilde{l}}(\hat{\theta})(\theta - \hat{\theta}), \quad (9.31)$$

where the first-order term of the expansion vanishes since $\hat{\theta}$ is an optimum, and $H_{\tilde{l}}(\hat{\theta})$ is the Hessian matrix of \tilde{l} . We can now approximate the terms in Eq. (9.29) as:

$$\begin{aligned}
p(\mathcal{Y}|M_i) &= \int_{-\infty}^{\infty} e^{\tilde{l}(\boldsymbol{\theta})} d\boldsymbol{\theta} \approx e^{\tilde{l}(\hat{\boldsymbol{\theta}})} \int_{-\infty}^{\infty} e^{\frac{1}{2}(\boldsymbol{\theta}-\hat{\boldsymbol{\theta}})^T H_i(\hat{\boldsymbol{\theta}})(\boldsymbol{\theta}-\hat{\boldsymbol{\theta}})} d\boldsymbol{\theta} \\
&= e^{\tilde{l}(\hat{\boldsymbol{\theta}})} \int_{-\infty}^{\infty} e^{-\frac{1}{2}(\boldsymbol{\theta}-\hat{\boldsymbol{\theta}})^T [N \bar{\mathcal{J}}(\hat{\boldsymbol{\theta}})](\boldsymbol{\theta}-\hat{\boldsymbol{\theta}})} d\boldsymbol{\theta} \\
&= p(\mathcal{Y}|\hat{\boldsymbol{\theta}}, M_i) p(\hat{\boldsymbol{\theta}}|M_i) (2\pi)^{d/2} |N \bar{\mathcal{J}}(\hat{\boldsymbol{\theta}})|^{-1/2} \\
&= p(\mathcal{Y}|\hat{\boldsymbol{\theta}}, M_i) p(\hat{\boldsymbol{\theta}}|M_i) (2\pi)^{d/2} N^{-d/2} |\bar{\mathcal{J}}(\hat{\boldsymbol{\theta}})|^{-1/2},
\end{aligned} \tag{9.32}$$

where $\bar{\mathcal{J}} = -\frac{1}{N} H_i(\hat{\boldsymbol{\theta}})$. This is called Laplace's method of approximation, and it is commonly used in population studies of pharmacokinetic and pharmacodynamic models (see for example [44]). Now if the parameter prior $p(\boldsymbol{\theta}|M_i)$ is noninformative compared to the likelihood $p(\mathcal{Y}|\boldsymbol{\theta}, M_i)$ in a region around $\hat{\boldsymbol{\theta}}$ (e.g., a uniform prior), it holds that [30]:

$$\bar{\mathcal{J}}(\hat{\boldsymbol{\theta}}) = -\frac{1}{N} H_i(\hat{\boldsymbol{\theta}}) \approx -\frac{1}{N} H_{l_p}(\hat{\boldsymbol{\theta}}) = \frac{1}{N} \mathcal{J}(\hat{\boldsymbol{\theta}}) = \bar{\mathcal{J}}(\hat{\boldsymbol{\theta}}), \tag{9.33}$$

where $l_p = \ln[p(\mathcal{Y}|\boldsymbol{\theta}, M_i)]$ and $\bar{\mathcal{J}}(\hat{\boldsymbol{\theta}})$ is the average Fisher information matrix (per datum). If $\bar{\mathcal{J}}(\hat{\boldsymbol{\theta}}) = \mathcal{J}(\hat{\boldsymbol{\theta}})$ we have that (-2 is convenient for Gaussians):

$$\begin{aligned}
-2 \ln(p(M_i|\mathcal{Y})) &\sim -2 \ln(p(\mathcal{Y}|M_i) p(M_i)) = \\
&-2 \left(\ln(p(\mathcal{Y}|\hat{\boldsymbol{\theta}}, M_i)) - \frac{d}{2} \ln(N) + \ln \left(\frac{(2\pi)^{d/2} p(\hat{\boldsymbol{\theta}}|M_i) p(M_i)}{|\bar{\mathcal{J}}|^{1/2}} \right) \right),
\end{aligned} \tag{9.34}$$

where Eq. (9.32) was used in the last step. Note that the last term of Eq. (9.34) does not grow with the sample size, so for $N \rightarrow \infty$:

$$\text{BIC} \equiv -2 \ln(p(M_i|\mathcal{Y})) \approx -2 \ln(p(\mathcal{Y}|\hat{\boldsymbol{\theta}}, M_i)) + d \ln(N). \tag{9.35}$$

This quantity is commonly referred to as the Bayesian information criterion (BIC) and used for model assessment. Note that the BIC is proportional to the AIC with the factor 2 in the second term replaced by $\ln(N)$ [24].

However, it has been reported that retention of the term $-\frac{d}{2} \ln 2\pi$ in Eq. (9.34) may improve the accuracy of the expression in practice [17]. This results in the alternative criterion BIC_{alt} :

$$\text{BIC}_{\text{alt}} = -2 \ln(p(M_i|\mathcal{Y})) \approx -2 \ln(p(\mathcal{Y}|\hat{\boldsymbol{\theta}}, M_i)) + d \ln\left(\frac{N}{2\pi}\right). \tag{9.36}$$

Choosing the model with the smallest BIC corresponds to choosing the model with the largest posterior probability [24]. However, it is also possible to quantify the ratio of plausibility of two models in light of the experimental data with the BIC:

$$\begin{aligned}
\ln \left(\frac{p(M_i|\mathcal{Y})}{p(M_j|\mathcal{Y})} \right) &= \ln(p(M_i|\mathcal{Y})) - \ln(p(M_j|\mathcal{Y})) \\
&\approx \ln(p(\mathcal{Y}|\hat{\theta}_i, M_i) - \frac{d_i}{2} \ln(N)) - \left(\ln(p(\mathcal{Y}|\hat{\theta}_j, M_j) - \frac{d_j}{2} \ln(N)) \right) \\
&= \ln \left(\frac{p(\mathcal{Y}|\hat{\theta}_i, M_i)}{p(\mathcal{Y}|\hat{\theta}_j, M_j)} \right) - \frac{(d_i - d_j)}{2} \ln(N). \tag{9.37}
\end{aligned}$$

Kass et al. [30] argued for the validity of this approximation as $N \rightarrow \infty$.

Equations (9.35) and (9.37) tell us that we only need to identify the parameter optima for the models to be assessed and compared. This simplifies the computations, but it does not solve the issues with large parameter spaces, despite contrary arguments in the literature.

Bayes factors have been employed for a number of biochemical applications, unlike the BIC that appears to be less frequently used in practice. Different methods to compute the marginal likelihood, which requires integration of a (high-dimensional) parameter space, have been investigated for synthetic data [38, 59]. Xu et al. [63] used Bayes factors to compute the posterior probabilities of four models of the extracellular signal-regulated kinase (ERK) pathway. This analysis suggested that B-Raf is necessary to fully activate the pathway, and indicated the absence of a mechanism for degradation of the bound epidermal growth factor receptor. In several studies, approximate Bayesian computation (ABC) methods have been employed to compute Bayes factors, see for example [19, 54, 55]. The ABC approach has the advantage that the likelihood function does not need to be defined; the posterior distributions are approximated from simulations instead. This may be an advantage (or even necessary) for certain model selection applications [53].

9.5 Minimum Description Length

Minimum description length (MDL) constitutes a fundamentally different approach to model selection; for an extensive introduction see the tutorial by Grünwald [23]. The main idea behind MDL is to select the model that can compress the experimental data the most, which requires the interpretation of probability distributions in terms of code lengths.

The connection between model probabilities and code lengths can be illustrated by letting model M_j be represented by $j = 1, \dots, m$ in binary representation. For the example of four models, M_1 is represented by 00, M_2 by 10, M_3 by 01, M_4 by 11, so that the number of bits needed to specify m models is $L = \log_2(m)$. Now assume that our model class takes the form $\mathcal{M} = \{M_1, M_2, \dots, M_m\}$, and that each model has the prior probability $p_i = 1/m, i = 1 \dots, m$. Then $L = \log_2(m) = -\log(p_i)$ [23]. However, this principle also extends to continuous non-uniform probability distributions and non-integer bits.

In the crude two parts version of MDL, the goal is to minimize the sum of the code lengths required to describe the model $L(M)$ and the data when encoded with the model $L(D|M)$. Although the data term can be represented by the negative log-likelihood of the data given the model, little advice is offered on how to encode $L(M)$. However, a refined version of MDL was developed, which states that we should select the model with minimal regret in the worst case sense (over the possible observations). The regret roughly represents the additional information needed to encode the data with M compared to an encoding with the optimal model in \mathcal{M} (after the data is observed). Shtarkov showed that this is uniquely achieved by a normalized maximum likelihood distribution p_{nml} [23, 46]. As a criterion for model selection we can therefore use:

$$-\log(p_{nml}) = -\log(p(y|\hat{\theta}, M)) + \text{COMP}(M|y). \quad (9.38)$$

where the term $\text{COMP}(M|y)$ can be regarded as a measure of the number of hypothetical data sequences that can be generated by the model. Interestingly, it was shown by Rissanen [42] that for a model M with d parameters, and with encoding probability distribution in the exponential family, the parametric complexity asymptotically (with sample size) takes the form:

$$\text{COMP}(M|y) = \frac{d}{2} \log\left(\frac{N}{2\pi}\right) + \log \int_{\theta \in \Theta} \sqrt{|I(\theta)|} d\theta, \quad (9.39)$$

where $I(\theta)$ is the Fisher information matrix. If we compare Eqs. (9.38)–(9.39) to (9.34) we see that $\text{BIC} = -2 \log(p_{nml})$ if we choose $p(\hat{\theta}|M_i) = \frac{\sqrt{|I(\hat{\theta})|}}{\int_{\theta \in \Theta} \sqrt{|I(\theta)|}}$ in BIC (apart from the term $p(M_i)$ in BIC). This is known as the (normalized) Jeffreys prior. Finally, note that the model prior term $p(M_i)$ in the BIC is insignificant for $N \rightarrow \infty$.

Perhaps because the resulting criterion for model evaluation is similar to that of the BIC, although the underlying approach is very different, MDL has so far received more attention for theoretical than for practical applications in the domain of dynamical biochemical modeling.

9.6 Sampling Strategies for Numerical Integration

Model selection methods such as the AIC and the BIC are based on a single optimal parameter point, which is convenient for practical applications. However, the reliability of inference procedures based on a single optimal parameter point has been questioned [62]. For example, ambiguity in parameter values translates into ambiguity in model predictions and consequently affects the model assessment. Also, most single parameter point methods are only asymptotically valid for a sufficiently large number of experimental observations (the cAIC is an exception). For these reasons, among others, methods based on parameter sampling are commonly considered as advantageous for the analysis of biochemical systems.

However, methods based on multiple parameter points are computationally considerably more demanding. For example, from Eq. (9.29) we see that the computation of Bayes factors requires the evaluation of integrals over the parameter space:

$$p(\mathcal{Y}|M) = \int_{\Theta} p(\mathcal{Y}|\theta, M)p(\theta|M)d\theta \equiv I \propto \int_{\Theta} p(\theta|\mathcal{Y}, M)d\theta, \quad (9.40)$$

where $p(\theta|M)$ is the prior, $p(\mathcal{Y}|\theta, M)$ is the likelihood, and $p(\theta|\mathcal{Y}, M)$ is the posterior distribution. In general, it is not possible to find an analytic closed form expression of the integral in Eq. (9.40), although exceptions exist [65]. It may also be possible to find closed form approximations of the integral as in the derivation of the BIC. However, to compute Bayes factors with limited available data, we typically need to resort to numerical approximations by simulation. The simplest way of approximating (9.40) would be to sample a large number of parameter points directly from the prior. The integral in Eq. (9.40) can then be asymptotically (w.r.t. the number of sampled parameter points N) approximated as [59]:

$$p(\mathcal{Y}|M) \simeq \frac{1}{N} \sum_{i=1}^N p(\mathcal{Y}|\theta^{(i)}, M) \equiv I_N, \theta^{(i)} \sim p(\theta|M) \quad (9.41)$$

It is also interesting to investigate how accurately I_N approximates I with the number of sampled points. From the central limit theorem we know that [43]:

$$\lim_{N \rightarrow \infty} \sqrt{N}(I_N - I) \sim \mathcal{N}(0, \sigma^2) \quad (9.42)$$

where:

$$\sigma^2 = \int (p(\mathcal{Y}|\theta, M) - I)^2 p(\theta|M)d\theta \quad (9.43)$$

However, the high likelihood region of a biochemical model typically only constitutes a small part of the parameter space, with a negligible likelihood for a vast majority of the drawn parameter points. An accurate estimate of the integral in Eq. (9.40) would therefore require a very large, or even intractable, number of sample points. For this reason, algorithms have been developed that can be used to approximate integrals in the form of Eq. (9.40) more efficiently.

One popular approach is to construct a Metropolis-Hastings Markov chain Monte Carlo (MH-MCMC) chain. A Markov chain in our setting corresponds to a sequence of random draws of parameter points, for which each new draw $\theta' \in \Theta$ is only conditioned on the current point $\theta \in \Theta$, which makes the process “memoryless”. Now define the probability of a transition from θ to θ' as $P(\theta \rightarrow \theta')$, and assume that the chain converges to a stationary target distribution $p^*(\theta)$ for which:

$$p^*(\theta)P(\theta \rightarrow \theta') = p^*(\theta')P(\theta' \rightarrow \theta). \quad (9.44)$$

We may then define that $P(\theta \rightarrow \theta') = q(\theta \rightarrow \theta')\alpha(\theta \rightarrow \theta')$, so that Eq. (9.44) takes the form:

$$p(\theta)q(\theta \rightarrow \theta')\alpha(\theta \rightarrow \theta') = p(\theta')q(\theta' \rightarrow \theta)\alpha(\theta' \rightarrow \theta) \quad (9.45)$$

where $q(\theta \rightarrow \theta')$ is a proposal distribution (used to draw the new point θ'), and $\alpha(\theta \rightarrow \theta') \in [0, 1]$ is the probability to accept θ' . If we then define that $\alpha(\theta' \rightarrow \theta) = 1$, we have that the acceptance probability for a transition from θ to θ' is:

$$\alpha(\theta \rightarrow \theta') = \min \left[\frac{p(\theta')q(\theta' \rightarrow \theta)}{p(\theta)q(\theta \rightarrow \theta')}, 1 \right] \quad (9.46)$$

To run the Metropolis Hastings algorithm we first need to choose a starting point (e.g., a random point in Θ), choose a proposal distribution (e.g., a Gaussian distribution) and a target distribution (e.g., $p(\mathcal{Z}|\theta, M)$), and then use Eq. (9.46) to update the chain. If the i th proposed parameter point $\hat{\theta}^{(i)}$ is accepted $\theta^{(i+1)} = \hat{\theta}^{(i)}$, and otherwise $\theta^{(i+1)} = \theta^{(i)}$. In practice, Eq. (9.45) can only be satisfied after sufficiently many points have been drawn and discarded in the initial “burn-in” phase of the algorithm, and a measure such as the autocorrelation of the chain is used to check if the target density is indeed stationary. For a more detailed introduction to the Metropolis-Hastings method see [12], and for a nice intuitive explanation see [61]. We generated artificial data for a model of degradation in Fig. 9.4a with the only parameter $k = 1$, and the Metropolis Hastings algorithm was then applied with a starting guess of $k = 0.5$. The Metropolis-Hastings algorithm was run for 10^4 (Fig. 9.4b), 10^5 (Fig. 9.4c), and 10^6 (Fig. 9.4d) steps, and the first 10% of the steps were discarded. As can be seen in Fig. 9.4b–d a large number of steps is required to compute a smooth and accurate posterior in this case. Potentially, the parameters of the method could be tuned so that fewer steps are required.

Gibbs sampling is another common strategy when it is not possible to sample directly from the joint parameter distribution (for an introduction see [7]). Probability densities for individual parameters, conditioned on the other parameters, are instead used to update one parameter at a time, making it suitable for Bayesian networks applications. However, Gibbs sampling constitutes a special case of the Metropolis-Hastings sampling [21]. Another popular approach to compute posterior distributions is sequential Monte Carlo methods (particle filters), for which one data point is incorporated at a time with a Bayesian updating equation (compare for example to Kalman filters) [43]. For comparative studies of how commonly used sampling methods perform for inference on artificial data, see the work by Miliias-Argeitis et al. [38] and Vyshemirsky et al. [59].

However, as the example in Fig. 9.4 demonstrates, accurate approximations may involve (prohibitively) high numbers of samples even for very small biochemical models. Two strategies have been proposed recently to deal with this problem. One strategy to compute Eq. (9.40) is to identify the high-likelihood regions of the parameter space, assuming that the contribution to the integral from the other regions is




Fig. 9.4 Posterior computation with the Metropolis-Hastings MCMC algorithm. We estimate the rate constant k for a linear degradation model $X(t) = X_0 e^{-kt}$, given the set of artificial data for X (assuming a std of 10% of the data values) shown in (a) with artificial data: stars, and model simulation with $k = 1$: solid line. **b–d** Posterior distributions for k for 10^4 steps (b), 10^5 steps (c), and 10^6 steps (d) of the MH-MCMC method with the starting guess $k = 0.5$. A Gaussian proposal distribution with a std of 0.1 was used to draw samples, and the first 10% of the samples were considered “burn-in” samples and discarded

negligible, and then to sample uniformly in the high-likelihood regions. A two step method based on this principle was proposed by Zamora-Sillero et al. [64]. Advanced numerical methods to speed up the evaluation of probability densities in parameter space are also being developed, such as those based on polynomial approximations with adaptive sparse grids [22]. If a valid polynomial approximation of the likelihood function is available it is trivial to compute the integral in Eq. (9.40).

9.7 Discussion

The construction and selection of a mathematical model as a compact and predictive representation of a biochemical system is an active area of research that is likely only in its infancy for biochemical applications. Actually, when browsing the literature there are surprisingly few examples of comparisons between dynamical models of biochemical systems. In the analysis of many biochemical systems a single candidate model is assessed, which was already argued against by Chamberlain in 1890 [10]. As pointed out by Chamberlain, the construction of several models typically requires

the researcher to think hard about alternative plausible explanations for the studied phenomena. However, recently a number of approaches has become available for the automatic construction of candidate models [20, 33, 52].

The most suitable method to assess a set of models depends on the structure of, and relation between, the candidate models as well as the experimental data. The likelihood ratio test is convenient for model discrimination, but it requires the candidate models to be nested, which is commonly not the case. Other criteria such as the AIC and BIC are also easy to handle, since only the optimal parameter point is required, but they are only asymptotically valid for large data sets. For small data sets, the researcher is left with the corrected AIC (and DIC) among the information theoretic approaches, and with the direct computation of posterior probabilities from Bayes factors. In the computation of posterior probabilities the researcher will commonly face issues such as the choice of the prior parameter distribution and the curse-of-dimensionality in the integration over high-dimensional parameter spaces. However, there are no available methods to avoid these issues, with a few exceptions (e.g., the BIC does not require prior distributions to be defined). Advanced computational methods for the integration of high-dimensional parameter spaces have been proposed [22], but it remains to be demonstrated that they are applicable for biochemical models.

We also note that (ad hoc) model selection criteria that do not fall under any of the main categories are sometimes useful. For example, a by now classical model for segmentation gene expression patterns in *Drosophila* was inferred by von Dassow et al. [58] using a goodness-of-fit function for pattern matching with the goal to resemble the judgment of an expert researcher. Lillacci et al. [36] used an extended Kalman filter to estimate statistical moments of states and parameters simultaneously, and used the overlap between model predictions and data as well as a χ^2 test as criteria for model selection.

To accurately apply model selection methods it is important to understand the concepts underlying the alternative methods, and to be able to judge when a certain method is applicable and suitable. It is also important to understand the inherent issues and limitations of the available methods. With this chapter we hope to provide the reader with an overview of the current situation in this domain as well as an easily accessible resource for the purpose of model selection in biochemical applications. However, many of the discussed principles apply equally well in other domains of science and engineering.

Acknowledgments We acknowledge funding from the Swiss Initiative for Systems Biology SystemsX.ch (project YeastX) evaluated by the Swiss National Science Foundation.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

1. Akaike, H.: Information theory and an extension of the maximum likelihood principle. In Petrov, B.N., Csaki, F. (Eds.) 2nd International Symposium on Information Theory, pp. 267–281 (1973)
2. Ando, T.: Bayesian predictive information criterion for the evaluation of hierarchical Bayesian and empirical Bayes models. *Biometrika* **94**, 443–458 (2007)
3. Augusto, D.A., Barbosa, H.J.C.: Symbolic regression via genetic programming. In: IEEE Proceedings of the Sixth Brazilian Symposium on Neural Networks, pp. 173–178 (2000)
4. Berglund, M., Sunnåker, M., Adiels, M., Jirstrand, M., Wennberg, B.: Investigations of a compartmental model for leucine kinetics using non-linear mixed effects models with ordinary and stochastic differential equations. *Math. Med. Biol.* (2011)
5. Buhmann, J.M.: Information theoretic model validation for clustering. In: International Symposium on Information Theory, pp. 1398–1402, Austin Texas, IEEE (2010)
6. Burnham, K.P., Anderson, D.R., Huyvaert, K.P.: Aic model selection and multimodel inference in behavioral ecology: some background, observations, and comparisons. *Behav. Ecol. Sociobiol.* **65**, 2335 (2011)
7. Casella, G., George, E.I.: Explaining the gibbs sampler. *Am. Stat.* **46**(3), 167–174 (1992)
8. Cavanaugh, J.E.: Unifying the derivations for the akaike and corrected akaike information criteria. *Stat. Probab. Lett.*, pp. 201–208 (1997)
9. Cedersund, G., Samuelsson, O., Ball, G., Tegnér, J., Gomez-Cabrero, D.: Optimization in biology parameter estimation and the associated optimization problem. In: Uncertainty in Biology, A Computational Modeling Approach. Springer, Chem (2016, this volume)
10. Chamberlin, T.C.: The method of multiple working hypotheses. *Science* **15**, 92–96 (1890)
11. Chehreghani, M.H., Busetto, A.G., Buhmann, J.M.: Information theoretic model validation for spectral clustering. In: Proceedings of the 15th International conference on artificial intelligence and statistics (AISTATS), pp. 495–503, 2012
12. Chib, S., Greenberg, E.: Understanding the metropolis-hastings algorithm. *Am. Stat.* **49**, 327–335 (1995)
13. Chris, S.: Wallace and David M Boulton. An information measure for classification. *Comput. J.* **11**(2), 185–194 (1968)
14. Csilléry, K., Blum, M.G.B., Gaggiotti, O.E., François, O.: Approximate bayesian computation (abc) in practice. *Trends Ecol. Evol.* **25**(7), 410–418 (2010)
15. Dalle Pezze, P., Sonntag, A.G., Thien, A., Prentzell, M.T., Goedel, M., Fischer, S., Neumann-Haefelin, E., Huber, T.B., Baumeister, R., Shanley, D.P., Thedieck, K.: A dynamic network model of mtor signaling reveals tsc-independent mtorc2 regulation. *Sci. Sig.* **5**(217), ra25 (2012)
16. DeLeeuw, J.: Introduction to Akaike (1973) information theory and an extension of the maximum likelihood principle. (1992)
17. Draper, D.: Assessment and propagation of model uncertainty. *J. R. Stat. Soc. Ser. B* **57**, 45–97 (1995)
18. Efron, B.: Estimating the error rate of a prediction rule: improvement on cross-validation. *J. Am. Stat. Assoc.* **78**, 316–331 (1983)
19. Fagundes, N.J.R., Ray, N., Beaumont, M., Neuenchwander, S., Salzano, F.M., Bonatto, S.L., Excoffier, L.: Statistical evaluation of alternative models of human evolution. *Proc. Natl. Acad. Sci. U.S.A.* **104**(45), 17614–17619 (2007). Nov
20. Floettmann, M., Schaber, J., Hoops, S., Klipp, E., Mendes, P.: Modelmage: A tool for automatic model generation, selection and management. *Genome. Inform.* **20**, 52–63 (2008)
21. Gerstner, T., Griebel, M.: Iterative and non-iterative simulation algorithms. *Computing Science and Statistics (Interface Proceedings)* **24**, 433–438 (1992)
22. Gerstner, T., Griebel, M.: Dimension-adaptive tensor-product quadrature. *Computing* **71**(1), 65–87 (2003)

23. Grünwald, P.: A tutorial introduction to the minimum description length principle. In: Grünwald, M.A.P.P., Myung, I.J. (Ed.) *Advances in Minimum Description Length: Theory and Applications*. MIT Press, US, 2005
24. Hastie, T., Tibshirani, R., Friedman, J.: *The Elements of Statistical Learning - Data Mining, Inference, and Prediction*, Second edition. Springer, 2008
25. Hug, S., Schmidl, D., Li, W.B., Greiter, M.B., Theis, F.J.: Bayesian model selection methods and their application to biological ODE systems. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
26. Hurvich, C.M., Tsai, C.-L.: Regression and time series model selection in small samples. *Biometrika* **76**, 297–307 (1989)
27. Jiang, W., Simon, R.: A comparison of bootstrap methods and an adjusted bootstrap approach for estimating the prediction error in microarray classification. *Stat. Med.* **26**(29), 5320–5334 (2007)
28. Johnson, J.B., Omland, K.S.: Model selection in ecology and evolution. *Trends. Ecol. Evol.* **19**(2), 101–108 (2004). Feb
29. Karlsson, M.O., Beal, S.L., Sheiner, L.B.: Three new residual error models for population pk/pd analyses. *J. Pharmacokinet. Biopharm.* **23**(6), 651–672 (1995)
30. Kass, R.E., Raftery, A.E.: Bayes factors. *J. Am. Stat. Assoc.* **90**(430), 773–795 (1995)
31. Kirk, P., Silk, D., Stumpf, M.P.H.: Reverse engineering under uncertainty. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
32. Kirk, P., Thorne, T., Stumpf, M.P.H.: Model selection in systems and synthetic biology. *Curr. Opin. Biotechnol.* (0):– (2013)
33. Kristensen, N.R., Madsen, H., Ingwersen, S.H.: Using stochastic differential equations for pk/pd model development. *J. Pharmacokinet. Pharmacodyn.* **32**(1), 109–141 (2005)
34. Kuepfer, L., Peter, M., Sauer, U., Stelling, J.: Ensemble modeling for analysis of cell signaling dynamics. *Nat. Biotechnol.* **25**(9), 1001–1006 (2007). Sep
35. Kuwahara, H., Myers, C.J., Samoilov, M.S., Barker, N.A., Arkin, A.P.: Automated abstraction methodology for genetic regulatory networks. In: *Transactions on computational systems biology VI*, pp. 150–175. Springer, 2006
36. Lillacci, G., Khammash, M.: Parameter estimation and model selection in computational biology. *PLoS Comput. Biol.* **6**(3), e1000696 (2010). Mar
37. Michael, D.: Schmidt, Ravishankar R Vallabhajosyula, Jerry W Jenkins, Jonathan E Hood, Abhishek S Soni, John P Wikswo, and Hod Lipson. Automated refinement and inference of analytical models for metabolic networks. *Phys. Biol.* **8**(5), 055011 (2011)
38. Miliás, A., Porreca, R., Summers, S., Lygeros, J.: Bayesian model selection for the yeast gata-factor network: a comparison of computational approaches. In: *IEEE Conference on Decision and Control*, Atlanta, Georgia, USA, 2010
39. Müller, T.G., Faller, D., Timmer, J., Swameye, I., Sandra, O., Klingmüller, U.: Tests for cycling in a signalling pathway. *J. Royal Stat. Soc.: Ser. C (Appl. Stat.)* **53**(4), 557–568 (2004)
40. Posada, D., Buckley, T.R.: Model selection and model averaging in phylogenetics: advantages of akaike information criterion and bayesian approaches over likelihood ratio tests. *Syst. Biol.* **53**(5), 793–808 (2004). Oct
41. Raue, A., Kreutz, C., Maiwald, T., Bachmann, J., Schilling, M., Klingmüller, U., Timmer, J.: Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics* **25**(15), 1923–1929 (2009). Aug
42. Rissanen, J.: Fisher information and stochastic complexity. *IEEE Transact. Inf. Theor.* **42**(1), 40–47 (1996)
43. Ristic, B., Arulampalam, S., Gordon, N.: *Beyond the Kalman filter*. Arctec House, 2004
44. Rune, V.: Overgaard, Niclas Jonsson, Christoffer W. Tornøe, and Henrik Madsen. Non-linear mixed-effects models with stochastic differential equations: implementation of an estimation algorithm. *J. Pharmacokinet. Pharmacodyn.* **32**(1), 85–107 (2005). Feb
45. Schliemann-Bullinger, M., Fey, D., Bastogne, T., Findeisen, R., Scheurich, P., Bullinger, E.: The experimental side of parameter estimation. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)

46. Shtarkov, Y.M.: Universal sequential coding of single messages. (Translated From) *Probl. Inf. Transm.* **23**(3), 3–17 (1987)
47. Sleep, D.J.H.: Statistical versus biological hypothesis testing: response to Steidl. *J. Wildl. Manag.* **71**, 21202121 (2007)
48. Spiegelhalter, D.J., Best, N.G., Carlin, B.P., van der Linde, A.: Bayesian Measures of Model Complexity and Fit. *J. Royal Stat. Soc.* 1–34 (2002)
49. Steidl, R.J.: Limits of Data Analysis in Scientific Inference: Reply to Sleep, et al. *J. Wildl. Manag.* **71**, 2122–2124 (2007)
50. Steidl, R.J.: Model selection, hypothesis testing, and risks of condemning analytical tools. *J. Wildl. Manag.* **70**, 14971498 (2006)
51. Sugiura, N.: Further analysts of the data by akaike's information criterion and the finite corrections. *Commun. Stat.* **7**(1), 13–26 (1978)
52. Sunnåker, M., Zamora-Sillero, E., Dechant, R., Ludwig, C., Busetto, A.G., Wagner, A., Stelling, J.: A method for automatic generation of predictive dynamic models reveals nuclear phosphorylation as the key msn2 control mechanism. *Sci. Signal.* **6**, ra41 (2013)
53. Sunnåker, M., Busetto, A.G., Numminen, E., Corander, J., Foll, M., Dessimoz, C.: Approximate bayesian computation. *PLoS Comput. Biol.* **9**(1), e1002803 (2013). Jan
54. Toni, T., Welch, D., Strelkowa, N., Ipsen, A., Stumpf, M.P.H.: Approximate bayesian computation scheme for parameter inference and model selection in dynamical systems. *J. R. Soc. Interf.* **6**(31), 187–202 (2009). Feb
55. Toni, T., Stumpf, M.P.H.: Simulation-based model selection for dynamical systems in systems and population biology. *Bioinformatics* **26**(1), 104–110 (2010). Jan
56. Transtrum, Mark K., Machta, Benjamin B., Sethna, James P.: Geometry of nonlinear least squares with applications to sloppy models and optimization. *Phys. Rev. E Stat. Nonlin. Soft. Matter. Phys.* **83**(3 Pt 2), 036701 (2011). Mar
57. Turkheimer, F.E., Hinz, R., Cunningham, V.J.: On the undecidability among kinetic models: from model selection to model averaging. *J. Cereb. Blood Flow Metab.* **23**(4), 490–498 (2003). Apr
58. von Dassow, G., Meir, E., Munro, E.M., Odell, G.M.: The segment polarity network is a robust developmental module. *Nature* **406**(6792), 188–192 (2000). Jul
59. Vyshemirsky, V., Girolami, M.A.: Bayesian ranking of biochemical system models. *Bioinformatics* **24**(6), 833–839 (2008). Mar
60. Wagenmakers, Eric-Jan, Farrell, Simon: AIC model selection using akaike weights. *Psychonomic. Bull. Rev.* **11**(1), 192–196 (2004)
61. Wilkinson, D.J.: Bayesian methods in bioinformatics and computational systems biology. *Brief Bioinform.* **8**(2), 109–116 (2007). Mar
62. Wilkinson, D.J.: Stochastic modelling for quantitative description of heterogeneous biological systems. *Nat. Rev. Genet.* **10**(2), 122–133 (2009). Feb
63. Xu, T.-R., Vyshemirsky, V., Gormand, A., von Kriegsheim, A., Girolami, M., Baillie, G.S., Ketley, D., Dunlop, A.J., Milligan, G., Houslay, M.D., Kolch, W.: Inferring signaling pathway topologies from multiple perturbation measurements of specific biochemical species. *Sci. Signal* **3**(113), ra20 (2010)
64. Zamora-Sillero, E., Hafner, M., Ibig, A., Stelling, J., Wagner, A.: Efficient characterization of high-dimensional parameter spaces for systems biology. *BMC Syst. Biol.* **5**, 142 (2011)
65. Zellner, A., Chung-Ki, M.: *Bayesian Analysis, Model Selection and Prediction*, 1st edn. Cambridge University Press, Cambridge Books Online, Cambridge (1993)

Chapter 10

Bayesian Model Selection Methods and Their Application to Biological ODE Systems

Sabine Hug, Daniel Schmidl, Wei Bo Li, Matthias B. Greiter
and Fabian J. Theis

Abstract In this chapter, we focus on Bayesian model selection for biological dynamical systems. We do not present an overview over existing methods, but showcase their comparison and the application to ordinary differential equation (ODE) systems, as well as the inference of the parameters in the ODE system. For this, our method of choice is the Bayes factor, computed by Thermodynamic Integration. We first present several model selection methods, both alternatives to the Bayes factor as well as several methods for calculating the Bayes factor, foremost among them said Thermodynamic Integration. As a simple example for the selection problem, we resort to a choice between normal distributions, which is analytically tractable. We apply our chosen method to a medium sized ODE model selection problem from radiation science and demonstrate how predictions can be drawn from the model selection results.

S. Hug · F.J. Theis (✉)

Institute of Computational Biology, Helmholtz Zentrum München, Neuherberg, Germany
e-mail: fabian.theis@helmholtz-muenchen.de

S. Hug

e-mail: sabine.hug@helmholtz-muenchen.de

S. Hug · F.J. Theis

Department of Mathematics, Technische Universität München, Garching, Germany

D. Schmidl

Institute of Bioinformatics and Systems Biology, Helmholtz Zentrum München,
Neuherberg, Germany
e-mail: daniel.schmidl@helmholtz-muenchen.de

W.B. Li · M.B. Greiter

Research Unit Medical Radiation Physics and Diagnostics, Helmholtz Zentrum München,
Neuherberg, Germany
e-mail: wli@helmholtz-muenchen.de

M.B. Greiter

e-mail: matthias.greiter@helmholtz-muenchen.de

© Springer International Publishing Switzerland 2016

L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_10

10.1 Introduction

The last few years have seen a steady increase in the use of dynamical models as powerful modeling tools for the modeling of biochemical systems such as signaling pathways [3, 4, 15, 44]. However, tuning these dynamical models to fit and explain the experimental data is a task that is far from trivial [23, 27, 38, 43]. The situation is even more complicated if the structure of the interaction mechanisms in the dynamical models itself is uncertain. Then we are faced with the problem of model extension and selection. In this chapter, we focus on the situation where a few (often hand curated by the expert) models already exist and the question at hand is which of these models best fits the data [24]. As there usually is no perfect model for a biological system, this is then a choice which model performs best according to the applied quality criteria. It is important to keep in mind that “Essentially, all models are wrong, but some are useful.” [6]. The question should always be what we expect our model to explain, and what knowledge can be gained from a fitted model. For this reason, we focus in our contribution on giving two practical examples on how to apply model selection methods and what insight might be gained from them.

Bayesian approaches are increasingly popular for model inference. They often require Markov chain Monte Carlo (MCMC) sampling procedures, which only now become feasible in larger systems due to the increase in computational power. Bayesian approaches then approximate the full distribution of the system in otherwise intractable systems. While this is computationally expensive, the gain is a natural access to all the information contained in full parameter distributions like confidence/credible regions or parameter correlations. For this reason we focus on likelihood or sampling based methods. In general, methods based on single point estimates are rather easy and computationally non-expensive to compute, however the single points might be unrepresentative for the whole distribution. Our method of choice for model selection is thus the **Bayes factor**, the quotient of the marginal likelihoods of the model. Since the marginal likelihood is a higher-dimensional integral, it is typically intractable and we resort to approximative methods. We have found the best performing method to be **thermodynamic integration**, a numerically stable if computationally expensive sampling scheme for the Bayes factor.

In order to test the reliability of such a sampling scheme, we did model selection on two simple models. We computed the Bayes Factor and the expected value in the thermodynamic integration both analytically and numerically and compared our results, finding a good agreement between the both values.

Next we present an application from radiation science, where two compartmental models corresponding to linear ordinary differential equations (ODEs) are compared based on experimental data. We demonstrate how to apply thermodynamic integration in these systems of twelve and fifteen parameters, respectively. Furthermore we showcase how predictions of divergent behavior can be drawn from the model.

10.2 Likelihood Based Model Selection

Usually, a model fit depends on parameters, especially in ODE models, this is often the case, see also [23]. Here the parameters are for example rate constants in the ODE or initial conditions. The shape of the ODE solution changes depending on these parameters. Quite a few rather easily accessible ways of doing model selection are then based on the maximum likelihood estimates (MLE) of the parameters for each model. The likelihood $p(\mathbf{Y}|\boldsymbol{\theta}_i, \mathbf{M}_i)$ is a measure for the agreement between data \mathbf{Y} and model \mathbf{M}_i , $i = 1, \dots, \mathcal{I}$ parametrized with parameters $\boldsymbol{\theta}_i$ and was already introduced in [43]. The parameter vector which maximizes the likelihood, i.e. the agreement between model and data, is called the **maximum likelihood estimate**, often written as $\hat{\boldsymbol{\theta}}_i$. It can be seen as the single best point estimate.

As seen previously, an ODE model can usually be written as

$$\frac{dx(t)}{dt} = f(\mathbf{x}(t), \boldsymbol{\theta}, \mathbf{u}(t)), \quad (10.1)$$

with a (usually nonlinear) function f , the state variables $\mathbf{x}(t)$, the parameters $\boldsymbol{\theta}$ and the external inputs $\mathbf{u}(t)$. Often, not all state variables can be observed, so we have to define a response function h that relates the state variables to the experimental observables $\mathbf{y} = h(\mathbf{x}) \in \mathbb{R}^m$. Considering that we only have noisy observations at discrete time points, cf. also Chapter [38], we arrive at:

$$\tilde{\mathbf{y}}_n = h(\mathbf{x}_n) + \boldsymbol{\varepsilon}_n, \quad \boldsymbol{\varepsilon}_n \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}), \quad (10.2)$$

in which the subscript $n = 1, \dots, N$ enumerates the time points t_n at which the measurement $\tilde{\mathbf{y}}_n$ was taken. The measurement noise $\boldsymbol{\varepsilon}_n$ is in most cases assumed to be normally distributed with covariance matrix $\boldsymbol{\Sigma}$. We denote the collection of all available experimental data by \mathbf{Y} . We can then write down the general likelihood explicitly:

$$p(\mathbf{Y}|\boldsymbol{\theta}) = \prod_{n=1}^N \frac{1}{(2\pi)^{m/2} \sqrt{|\boldsymbol{\Sigma}|}} \exp\left(-\frac{1}{2} (\tilde{\mathbf{y}}_n - \mathbf{y}(t_n))^\top \boldsymbol{\Sigma}^{-1} (\tilde{\mathbf{y}}_n - \mathbf{y}(t_n))\right) \quad (10.3)$$

We now briefly repeat some definitions of likelihood based model selection methods from the previous chapter [43]. Based on the MLE, several model selection criteria or tests have been proposed. Best known among them might be the Akaike Information Criterion (AIC) [1]. It is defined as

$$AIC(\hat{\boldsymbol{\theta}}_i) = -2 \log p(\mathbf{Y}|\hat{\boldsymbol{\theta}}_i, \mathbf{M}_i) + 2d_i, \quad (10.4)$$

where d_i is the number of independently adjusted parameters of model \mathbf{M}_i . The preferred model is the one with the minimal value for the AIC. The AIC weighs the goodness of fit, given by the loglikelihood value, with the associated number of

parameters, preferring smaller models over large models. The AIC is asymptotically efficient, but not consistent. Some care has to be taken concerning the conditions under which the AIC is applicable, see also the preceding chapter [43].

Somehow closely related is the Bayesian Information Criterion (BIC) [42]. In contrast to the AIC, it also takes into account the number of data points on which the choice is based:

$$BIC(\hat{\theta}_i) = -2 \log p(\mathbf{Y}|\hat{\theta}_i, \mathbf{M}_i) + d_i \log(N). \quad (10.5)$$

Here, N is the number of data points in \mathbf{Y} . Again, the model with the lowest BIC value should be chosen. As with the AIC, the BIC is only a good approximation if the models are identifiable and the number of data points is large, cf. [43].

Neither of the two criteria gives an absolute measure of how much one model is “better” than another model.

While these two criteria are rather closely related, a very different, yet also MLE based model choice method between two models is the likelihood ratio test (LRT) [24]. This method requires the models to be nested, meaning that the smaller of the models needs to be a special case of the larger model. The LRT is a hypothesis test with the null hypothesis that the smaller model (without loss of generality from now on model \mathbf{M}_1) is the true model that generated the data versus the alternative hypothesis that the larger model \mathbf{M}_2 generated the data. As the models are nested, the ratio of the logarithms of the maximum likelihood values is approximately χ^2 -distributed, with degrees of freedom d_1 and d_2 corresponding to the numbers of parameters in the two models:

$$-2 \log \left(\frac{p(\mathbf{Y}|\hat{\theta}_1, \mathbf{M}_1)}{p(\mathbf{Y}|\hat{\theta}_2, \mathbf{M}_2)} \right) \sim \chi_{d_2-d_1}^2. \quad (10.6)$$

For two nested models, the larger model always explains the data at least as well as the smaller model, thus $\frac{p(\mathbf{Y}|\hat{\theta}_1, \mathbf{M}_1)}{p(\mathbf{Y}|\hat{\theta}_2, \mathbf{M}_2)} < 1$. With the LRT, it is possible to determine if the improvement is significant by deriving a p-value under the appropriate χ^2 -distribution. Classical hypothesis testing then reveals if the null model can be rejected at the desired significance level.

10.3 Bayesian Model Selection Methods

10.3.1 The Bayes Factor

The BIC already points towards Bayesian inference. Bayesian inference is widely applied in systems biology in different forms, cf. [23, 27, 43]. Here, the likelihood $p(\mathbf{Y}|\theta)$ is complemented with prior information $p(\theta)$ available for the parameters to yield the general posterior distribution $p(\theta|\mathbf{Y})$ of the parameters given the data:

$$p(\boldsymbol{\theta}|\mathbf{Y}) = \frac{p(\mathbf{Y}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\mathbf{Y})} \quad (10.7)$$

An important quantity for the purpose of model selection is actually the marginal likelihood $p(\mathbf{Y})$ in the denominator of the posterior distribution.

With Bayes' theorem once again, we get:

$$p(\mathbf{M}_i|\mathbf{Y}) = \frac{p(\mathbf{Y}|\mathbf{M}_i)p(\mathbf{M}_i)}{\sum_j p(\mathbf{Y}|\mathbf{M}_j)p(\mathbf{M}_j)}, \quad (10.8)$$

which is to compute the marginal likelihood $p(\mathbf{Y}|\mathbf{M}_i)$ for the desired model \mathbf{M}_i , $i = 1, \dots, \mathcal{J}$. It is important to notice that the marginal likelihood is not straightforward to compute, since it is a usually high-dimensional and analytically intractable integral:

$$p(\mathbf{Y}|\mathbf{M}_i) = \int_{\mathbb{R}^{d_i}} p(\mathbf{Y}|\boldsymbol{\theta}_i, \mathbf{M}_i)p(\boldsymbol{\theta}_i|\mathbf{M}_i) d\boldsymbol{\theta}_i. \quad (10.9)$$

This integral has to be approximated, usually with sampling based approaches. Nevertheless, if we then want to compare two models \mathbf{M}_1 and \mathbf{M}_2 , we can do so by computing the ratio of the two marginal likelihoods, the so-called **Bayes Factor**

$$B_{12} = \frac{p(\mathbf{Y}|\mathbf{M}_1)}{p(\mathbf{Y}|\mathbf{M}_2)}, \quad (10.10)$$

in which a value of B_{12} greater than one indicates a preference for model \mathbf{M}_1 and a value less than one one for model \mathbf{M}_2 .

Jeffreys established a widely used interpretation of the Bayes factor in [22]. It is based on a classification of the evidence in favor of model \mathbf{M}_1 in \log_{10} -half-scale units as:

$\log_{10}(B_{12})$	B_{12}	Evidence in favor of model \mathbf{M}_1
0–0.5	1–3.2	Not worth more than a bare mention
0.5–1	3.2–10	Substantial
1.0–1.5	10–32.6	Strong
1.5–2.0	32.6–100	Very strong
2.0– ∞	100– ∞	Decisive

This scale has become known as Jeffreys' scale of evidence. While it certainly can be challenged, it nevertheless is well established and widely used in the Bayesian community. The Bayes factor offers certain advantages over the presented point-based model selection methods, see also [43] for additional arguments. First, in contrast to the likelihood ratio test, it provides evidence for either of the models, since the Bayes factor in favor of model \mathbf{M}_2 can easily be interpreted by the same Jeffreys' scale by taking $B_{21} = 1/B_{12}$. Secondly, it works for non-nested models. Thirdly, point-based methods might not be appropriate in cases where the MLE is

not representative for the whole distribution, e.g. for multimodal likelihoods, see also Chap. 6 of [31] for an example. Furthermore, by taking into consideration the whole parameter space, the Bayes factor is more efficient in preventing overfitting [32] than the other introduced methods.

10.3.2 Sampling Based Methods for Calculating Bayes Factors

As already mentioned, the crux is that the marginal likelihood is computed by integrating over the whole parameter space, which is computationally costly and also often not straightforward. Because of this, standard methods for computing Bayes factors are mostly sampling based. This can be seen from the following relationship:

$$p(\mathbf{Y}|\mathbf{M}_i) = \int_{\mathbb{R}^{d_i}} p(\mathbf{Y}|\boldsymbol{\theta}_i, \mathbf{M}_i) p(\boldsymbol{\theta}_i|\mathbf{M}_i) d\boldsymbol{\theta}_i = \mathbb{E}_{p(\boldsymbol{\theta}_i|\mathbf{M}_i)} [p(\mathbf{Y}|\boldsymbol{\theta}_i, \mathbf{M}_i)] \quad (10.11)$$

In this chapter we present the following methods for estimating marginal likelihoods and thus Bayes factors:

Method	Sample from	Remarks
Prior arithmetic mean	Prior	Can be very efficient
Posterior harmonic mean	Posterior	Known variance issues
Chib's method	Posterior	Basically point estimate
Thermodynamic integration	Power posterior	Numerically stable

The easiest approach for sampling any of the marginal likelihoods, here now simply denoted $p(\mathbf{Y}|\mathbf{M})$, is the prior arithmetic mean. For this approach, a total of T samples $\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}, \dots, \boldsymbol{\theta}^{(T)}$ are drawn from the prior distribution $p(\boldsymbol{\theta})$. From Eq. (10.11) it can then be inferred that

$$p(\mathbf{Y}|\mathbf{M}) = \mathbb{E}_{p(\boldsymbol{\theta})} [p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})] \approx \frac{1}{T} \sum_{j=1}^T p(\mathbf{Y}|\boldsymbol{\theta}^{(j)}, \mathbf{M}) \quad (10.12)$$

The right hand side of this equation is known as the prior arithmetic mean estimate. The strong law of large numbers (almost surely) guarantees convergence as the sample number tends to infinity. However, in many practical applications, the prior does not contain too much information about the actual shape of the posterior. Then many samples might have very low likelihood values, thus a large number of samples might be needed for accurate results, which will be demonstrated later in this chapter.

Slightly more involved is the approach by Newton and Raftery [35] called the posterior harmonic mean. As the name already implies, for this approach samples are not drawn from the prior, but from the posterior distribution directly. Usually, this sampling is not directly possible anymore, but Markov chain Monte Carlo (MCMC)

methods [37] have to be applied. Similarly to the prior arithmetic mean, we draw a total of T samples $\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(T)}$ from the posterior distribution $p(\theta|\mathbf{Y}, \mathbf{M})$. This then yields the following marginal likelihood approximation:

$$p(\mathbf{Y}|\mathbf{M}) \approx \left(\frac{1}{T} \sum_{j=1}^T \frac{1}{p(\mathbf{Y}|\theta^{(j)}, \mathbf{M})} \right)^{-1} \quad (10.13)$$

The derivation can for example be found in [41]. However, already Neal [34] showed that this estimate suffers from severe issues. Newton and Raftery [35] also proposed a weighted combination of the prior arithmetic mean estimator and posterior harmonic mean estimator called the stabilized harmonic mean estimator. This helps to reduce the issues of the individual estimators.

Also often mentioned is Chib's method, which is originally also a point-based estimate. In [8], Chib and Jeliazkov show how to apply the method to the output of a Metropolis-Hastings sampling algorithm. The basic idea is to rearrange Bayes's theorem:

$$\log p(\mathbf{Y}|\mathbf{M}) = \log p(\mathbf{Y}|\theta^*, \mathbf{M}) + \log p(\theta^*|\mathbf{M}) - \log p(\theta^*|\mathbf{Y}, \mathbf{M}) \quad (10.14)$$

with a suitable θ^* , for example the Maximum Likelihood estimate. While this might yield an easily computable result, it might suffer from the same issues as other point-based estimates. Furthermore, the posterior probability $p(\theta^*|\mathbf{Y}, \mathbf{M})$ is often not readily available, since sampling and optimization are mostly based on "likelihood times prior" instead of the posterior, ignoring the proportionality constant that is actually the marginal likelihood at question here. For estimating the posterior value at the chosen point estimate, Chib and Jeliazkov propose to use the output of a Metropolis-Hastings sampler $\theta^{(1)}, \dots, \theta^{(T)}$.

If $q(\theta, \theta'|\mathbf{Y})$ denotes the proposal density of the Metropolis-Hastings algorithm for the transition from θ to θ' , where the proposal density is allowed to depend on the data \mathbf{Y} , and $\alpha(\theta, \theta'|\mathbf{Y})$ denotes the standard Metropolis-Hastings acceptance probability of a move, then it can be shown that

$$\hat{p}(\theta^*|\mathbf{Y}, \mathbf{M}) = \frac{T^{-1} \sum_{s=1}^S \alpha(\theta^{(s)}, \theta^*|\mathbf{Y}) q(\theta^{(s)}, \theta^*|\mathbf{Y})}{J^{-1} \sum_{j=1}^J \alpha(\theta^*, \theta^{(j)}|\mathbf{Y})} \quad (10.15)$$

is a simulation-consistent estimate of the posterior value. Here, the $\theta^{(s)}$ are samples drawn from the posterior, while the $\theta^{(j)}$ are drawn from $q(\theta^*, \cdot|\mathbf{Y})$ with θ^* fixed. This can then be plugged into Eq. (10.14) to yield an estimate for the marginal likelihood.

10.3.3 Thermodynamic Integration

Thermodynamic integration is a method for computing the Bayes factor also based on MCMC sampling. While it is computationally costlier than the other methods, it yields more robust and numerically stable results. Thermodynamic integration is based on path sampling ideas [11], and was first discussed for marginal likelihoods in the papers by [10, 28]. It has recently found increasing application in systems biology [7, 9, 44].

Central to the method is the **power posterior**, a variant of the usual posterior of the Bayesian setting,

$$p_\tau(\boldsymbol{\theta}|\mathbf{Y}, \mathbf{M}) = \frac{1}{p_\tau(\mathbf{Y})} p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})^\tau p(\boldsymbol{\theta}) \tag{10.16}$$

where $\tau \in [0, 1]$ is a so-called temperature parameter and $p_\tau(\mathbf{Y}) = \int_{\mathbb{R}^d} p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})^\tau p(\boldsymbol{\theta}) d\boldsymbol{\theta}$ is a normalization term necessary for making the power posterior a probability density. For $\tau = 0$, we get $p_{\tau=0}(\mathbf{Y}) = 1$, since this is the prior integrated over $\boldsymbol{\theta}$ and thus simply 1. The power posterior is then equal to the prior $p(\boldsymbol{\theta})$. For $\tau = 1$, we get $p_{\tau=1}(\mathbf{Y}) = p(\mathbf{Y}|\mathbf{M})$, the marginal likelihood and thus the power posterior is the regular posterior. Intuitively, a power posterior with a low value of τ thus corresponds to a smoother distribution closer to the prior and allows for more movement of the Markov chains through the parameter space, while a higher value of τ corresponds to a distribution closer to the posterior, which might be e.g. spiky due to the influence of the likelihood. The power posterior in total thus corresponds to a smooth transition from the prior to the posterior, which can also be seen in Fig. 10.1 with the example we will present later. We now derive an expression for the log marginal likelihood with respect to the power posterior which can be evaluated using MCMC methods. First we note that




Fig. 10.1 Visualization of the smooth transition from prior to posterior through the power posterior. Shown is the power posterior of the two parameter model \mathbf{M}_2 introduced in a later section, for five different temperatures

$$\frac{d}{d\tau} \log p_\tau(\mathbf{Y}) = \frac{d}{d\tau} \log \int_{\mathbb{R}^d} p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})^\tau p(\boldsymbol{\theta}) d\boldsymbol{\theta} \quad (10.17)$$

$$= \frac{1}{p_\tau(\mathbf{Y})} \int_{\mathbb{R}^d} \frac{dp(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})^\tau}{d\tau} p(\boldsymbol{\theta}) d\boldsymbol{\theta} \quad (10.18)$$

$$= \frac{1}{p_\tau(\mathbf{Y})} \int_{\mathbb{R}^d} \frac{dp(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})^\tau}{dt} \frac{p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})^\tau p(\boldsymbol{\theta})}{p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})^\tau} d\boldsymbol{\theta} \quad (10.19)$$

$$= \int_{\mathbb{R}^d} \frac{d \log p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})^\tau}{d\tau} \frac{p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})^\tau p(\boldsymbol{\theta})}{p_\tau(\mathbf{Y})} d\boldsymbol{\theta} \quad (10.20)$$

$$= \int_{\mathbb{R}^d} \log p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M}) \frac{p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})^\tau p(\boldsymbol{\theta})}{p_\tau(\mathbf{Y})} d\boldsymbol{\theta} \quad (10.21)$$

$$= \mathbb{E}_{p_\tau} \{ \log p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M}) \} \quad (10.22)$$

Integrating both sides with respect to τ yields the **thermodynamic integral**,

$$\int_0^1 \mathbb{E}_{p_\tau} \{ \log p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M}) \} d\tau = \int_0^1 \frac{d}{d\tau} \log p_\tau(\mathbf{Y}) \quad (10.23)$$

$$= \log p_{\tau=1}(\mathbf{Y}) - \log p_{\tau=0}(\mathbf{Y}) \quad (10.24)$$

$$= \log p(\mathbf{Y}|\mathbf{M}) \quad (10.25)$$

The integral in Eq. (10.25) can be solved numerically by choosing a discretization $0 = \tau_0 < \tau_1 < \dots < \tau_{K-1} < \tau_K = 1$, then the numerical approximation is

$$\log p(\mathbf{Y}|\mathbf{M}) \approx \frac{1}{2} \sum_{k=0}^{K-1} (\tau_{k+1} - \tau_k) (\mathbb{E}_{p_{\tau_{k+1}}} \{ \log p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M}) \} + \mathbb{E}_{p_{\tau_k}} \{ \log p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M}) \}) \quad (10.26)$$

The expectation for a specific τ can be obtained by Monte Carlo estimates,

$$\mathbb{E}_{p_{\tau_k}} \{ \log p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M}) \} \approx \frac{1}{S} \sum_{s=1}^S \log p(\mathbf{Y}|\boldsymbol{\theta}^{(s)}, \mathbf{M}) \quad (10.27)$$

where $\boldsymbol{\theta}^{(s)}$ denotes a sample drawn from $p_{\tau_k}(\boldsymbol{\theta}|\mathbf{Y}, \mathbf{M})$.

Thus we are able to compute the log marginal likelihood for our model. Doing so for two models provides us with the possibility to calculate the Bayes Factor and thus do model selection.

We provide an algorithm in pseudo code for thermodynamic integration:

Algorithm 1: Thermodynamic integration for calculating marginal likelihoods

input : number of temperature steps K , power posterior $p_\tau(\boldsymbol{\theta}|\mathbf{Y}, \mathbf{M})$
 depending on the temperature τ , number of samples S , loglikelihood
 $\log p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})$

output: log marginal likelihood $\log p(\mathbf{Y}|\mathbf{M})$

for $k \leftarrow 1$ **to** K **do**

$\tau_k \leftarrow (\frac{k}{K})^5$;

Draw S samples from $p_{\tau_k}(\cdot|\mathbf{Y}, \mathbf{M})$ with a Metropolis-Hastings algorithm;

Approximate $\mathbb{E}_{p_{\tau_k}}\{\log p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})\}$ by $\frac{1}{S} \sum_{s=1}^S \log p(\mathbf{Y}|\boldsymbol{\theta}^{(s)}, \mathbf{M})$

end

Set the log marginal likelihood to

$\log p(\mathbf{Y}|\mathbf{M}) =$

$\frac{1}{2} \sum_{k=0}^{K-1} (\tau_{k+1} - \tau_k) (\mathbb{E}_{p_{\tau_{k+1}}}\{\log p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})\} + \mathbb{E}_{p_{\tau_k}}\{\log p(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{M})\})$

10.4 Comparison of Model Selection Methods on a Simple Example

10.4.1 A Tractable Example for Model Selection

We now want to present an example where the Bayes factor could be computed analytically and thus the error made by the presented approximations was accessible. We have found this possible with a very simple model selection where we chose between the following two models:

- Model \mathbf{M}_1 : a normal distribution with expected value μ and standard deviation σ , with N data points drawn.
- Model \mathbf{M}_2 : two normal distributions with expected values μ_1 and $\mu_2 = -\mu_1$ and standard deviation σ , with N_1 data points drawn from the first normal distribution and N_2 drawn from the second.

Our data came from model \mathbf{M}_2 , so that the analytical Bayes Factor and the results of thermodynamic integration should both point towards model \mathbf{M}_2 significantly, since this model selection problem was rather simple.


We drew $N = 10$ samples from model \mathbf{M}_2 in the following way to obtain the data $\mathbf{Y} = (Y_1, \dots, Y_N)$

- $Y_1, \dots, Y_{N_1} \sim \mathcal{N}(\mu_1, \sigma^2)$ with $N_1 = 3$ and $\mu_1 = -2$.
- $Y_{N_1+1}, \dots, Y_N \sim \mathcal{N}(\mu_2, \sigma^2)$ with $N_2 = N - N_1 = 7$ and $\mu_2 = 2$.

and compared that to

- $Y_1, \dots, Y_N \sim \mathcal{N}(\mu, \sigma^2)$ with $N = 10$.

Fig. 10.2 The two models \mathbf{M}_1 (green) and \mathbf{M}_2 (purple, correct model) and the data (magenta dots) drawn from model \mathbf{M}_2



A visualization of the data can be seen in Fig. 10.2. To make the computation tractable, we chose a fixed σ , to which we assigned the value $\sigma = 2$ in our implementation. This left us with one free parameter, μ , for model \mathbf{M}_1 and two parameters μ_1 and μ_2 for model \mathbf{M}_2 .

In the following, we considered the following scenario: strong prior information for the parameters is available in the form of Gaussian priors. We thus chose:

- $\mu \sim \mathcal{N}(0, \sigma^2)$
- $\mu_1 \sim \mathcal{N}(-2, \sigma^2)$
- $\mu_2 \sim \mathcal{N}(+2, \sigma^2)$

10.4.2 Analytically Computing the Bayes Factor

The likelihoods we obtained for both models are given by

$$p(\mathbf{Y}|\mathbf{M}_1, \mu) = \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^N \exp\left(-\frac{1}{2\sigma^2}\left(\sum_{n=1}^N (Y_n - \mu)^2\right)\right) \tag{10.28}$$

and

$$p(\mathbf{Y}|\mathbf{M}_2, \mu_1, \mu_2) = \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^N \exp\left(-\frac{1}{2\sigma^2}\left(\sum_{n=1}^{N_1} (Y_n - \mu_1)^2 + \sum_{n=N_1+1}^N (Y_n - \mu_2)^2\right)\right) \tag{10.29}$$

After some straightforward calculations (see Schmidl [41] for a detailed derivation) we found that the posterior distributions in our scenario (Gaussian priors) for our two models were

- $\mathcal{N}\left(\frac{1}{N+1} \sum_{n=1}^N Y_n, \frac{\sigma^2}{N+1}\right)$ for model \mathbf{M}_1

- $\mathcal{N}\left(\left(\frac{1}{N_1+1}(-2 + \sum_{n=1}^{N_1} Y_n), \left(\frac{\sigma^2}{N_1+1} \quad 0\right)\right), \left(\frac{1}{N_2+1}(+2 + \sum_{n=1}^{N_2} Y_n), \left(0 \quad \frac{\sigma^2}{N_2+1}\right)\right)\right)$ for model \mathbf{M}_2

In order to compute the Bayes Factor, we needed to compute the marginals $p(\mathbf{Y}|\sigma, \mathbf{M}_1)$ and $p(\mathbf{Y}|\sigma, \mathbf{M}_2)$. We began with model \mathbf{M}_1 and thus $p(\mathbf{Y}|\sigma, \mathbf{M}_1)$:

$$\begin{aligned} p(\mathbf{Y}|\sigma, \mathbf{M}_1) &= \int_{\mathbb{R}} p(\mathbf{Y}|\mu)p(\mu) \, d\mu \\ &= \int_{\mathbb{R}} \left(\prod_{n=1}^N \mathcal{N}(Y_n|\mu, \sigma^2)\right) \cdot \mathcal{N}(\mu|0, \sigma^2) \, d\mu \\ &= \int_{\mathbb{R}} \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^{N+1} \exp\left(-\frac{1}{2\sigma^2} \sum_{n=1}^N (Y_n - \mu)^2 - \frac{1}{2\sigma^2} \mu^2\right) \, d\mu \\ &= \int_{\mathbb{R}} \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^{N+1} \exp\left(-\frac{1}{2\sigma^2} \left(\sum_{n=1}^N Y_n^2 - 2\mu \sum_{n=1}^N Y_n + (N+1)\mu^2\right)\right) \, d\mu \\ &= \int_{\mathbb{R}} \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^{N+1} \exp\left(-\frac{N+1}{2\sigma^2} \left(\frac{1}{N+1} \sum_{n=1}^N Y_n^2 - 2\mu \frac{N}{N+1} \bar{Y} + \mu^2\right)\right) \, d\mu \\ &= \int_{\mathbb{R}} \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^{N+1} \exp\left(-\frac{N+1}{2\sigma^2} \left(\frac{1}{N+1} \sum_{n=1}^N Y_n^2 - \left(\frac{N}{N+1}\right)^2 \bar{Y}^2\right)\right) \\ &\quad \times \exp\left(-\frac{N+1}{2\sigma^2} \left(\frac{N}{N+1} \bar{Y} - \mu\right)^2\right) \, d\mu \\ &= \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^N \frac{1}{\sqrt{N+1}} \exp\left(-\frac{N+1}{2\sigma^2} \left(\frac{1}{N+1} \sum_{n=1}^N Y_n^2 - \left(\frac{N}{N+1} \bar{Y}\right)^2\right)\right) \end{aligned}$$

where $\bar{Y} = \frac{1}{N} \sum_{n=1}^N Y_n$ is the sample mean. In a very similar fashion, we could also calculate $p(\mathbf{Y}|\sigma, \mathbf{M}_2)$ for Gaussian priors. For that, we introduced the notation $\bar{Y}_1 = \frac{1}{N_1} \sum_{n=1}^{N_1} Y_n$ and $\bar{Y}_2 = \frac{1}{N_2} \sum_{n=N_1+1}^N Y_n$:

$$\begin{aligned} p(\mathbf{Y}|\sigma, \mathbf{M}_2) &= \int_{\mathbb{R}} \int_{\mathbb{R}} p(\mathbf{Y}|\mu_1)p(\mu_1)p(\mathbf{Y}|\mu_2)p(\mu_2) \, d\mu_1 \, d\mu_2 \\ &= \int_{\mathbb{R}} \int_{\mathbb{R}} \left(\prod_{n=1}^{N_1} \mathcal{N}(Y_n|\mu_1, \sigma^2)\right) \cdot \mathcal{N}(\mu_1|-2, \sigma^2) \left(\prod_{n=N_1+1}^N \mathcal{N}(Y_n|\mu_2, \sigma^2)\right) \cdot \mathcal{N}(\mu_2|+2, \sigma^2) \, d\mu_1 \, d\mu_2 \\ &= \int_{\mathbb{R}} \int_{\mathbb{R}} \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^{N+2} \exp\left(-\frac{1}{2\sigma^2} \left(\sum_{n=1}^N Y_n^2 + 8 - 2(N_1\bar{Y}_1 - 2)\mu_1 - 2(N_2\bar{Y}_2 + 2)\mu_2\right.\right. \\ &\quad \left.\left.+ (N_1+1)\mu_1^2 + (N_2+1)\mu_2^2\right)\right) \, d\mu_1 \, d\mu_2 \\ &= \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^N \frac{1}{\sqrt{N_1+1}\sqrt{N_2+1}} \exp\left(-\frac{1}{2\sigma^2} \left(\sum_{n=1}^N Y_n^2 + 8 - \frac{(N_1\bar{Y}_1 - 2)^2}{N_1+1} - \frac{(N_2\bar{Y}_2 + 2)^2}{N_2+1}\right)\right) \end{aligned}$$

Having obtained the marginals, we could now compute the Bayes Factor B_{21} in favor of model \mathbf{M}_2 :

$$B_{21} = \frac{p(\mathbf{Y}|\sigma, \mathbf{M}_2)}{p(\mathbf{Y}|\sigma, \mathbf{M}_1)} \quad (10.30)$$

$$= \frac{\sqrt{N+1}}{\sqrt{N_1+1}\sqrt{N_2+1}} \exp\left(-\frac{1}{2\sigma^2} \left(\frac{(N\bar{\mathbf{Y}})^2}{N+1} - \frac{(N_1\bar{\mathbf{Y}}_1 - 2)^2}{N_1+1} - \frac{(N_2\bar{\mathbf{Y}}_2 + 2)^2}{N_2+1} + 8 \right)\right) \quad (10.31)$$

Since B_{21} only depends on the data \mathbf{Y} and the standard deviation σ , which we fixed, we could easily evaluate the Bayes Factor in our implementation.

10.4.3 Computational Results

All computations were performed in MATLAB 2013a. For the optimization based criteria AIC, BIC and LRT, we used 10,000 runs of a local optimization routine in MATLAB. Starting values were drawn uniformly random from the intervals $[-5, 5]$ for μ in \mathbf{M}_1 and from $[-5, 0]$ and $[0, 5]$ for μ_1 and μ_2 in \mathbf{M}_2 , respectively, to find the maximum likelihood estimates.

For the sampling based approaches (prior arithmetic mean, posterior harmonic mean and Chib's method), we drew 100,000 samples each from the required densities. For the prior, sampling was directly available. For the posterior, we sampled with Haario's Adaptive Metropolis Sampler [14]. Since Chib's method is in our opinion not tailored to accommodate adaptive sampling, we there chose a regular Metropolis-Hastings algorithm.

For thermodynamic integration, we followed the recommendations of Calderhead and Girolami [7] and chose a power law temperature schedule $\tau_k = (k/20)^5$ with $K = 20$ temperature steps. For each temperature, we drew 5000 samples, yielding also a total of 100,000 samples.

All sampling algorithms were initialized at the maximum likelihood estimates found for the AIC/BIC/LRT approaches. For the sampling based approaches, we ran the sampling 30 times on the same data to correct for randomness. This took less than an hour on a standard desktop computer. These 30 runs yielded mean results for the Bayes factor and the standard error.

The AIC for model \mathbf{M}_1 was 49.58, while the AIC for the correct model \mathbf{M}_2 was 40.74, thus the AIC made the correct choice in this simple example. For the BIC we found that the value for model \mathbf{M}_1 is 49.88, while the value for model \mathbf{M}_2 was 41.35, also indicating a preference for the correct model. The likelihood ratio test rejected the smaller model \mathbf{M}_1 with a p-value of 0.000956. For the sampling based approaches, we found the following:

Method	Mean Bayes factor and standard error
TRUE	139.23
Thermodynamic integration	140.66 \pm 0.13
Posterior harmonic mean	149.37 \pm 1.85
Prior arithmetic mean	138.76 \pm 0.03
Chib's method	55.76 \pm 0.01

For the Bayes factor, the analytical computation based on our drawn data showed a true value of 139.23 for model \mathbf{M}_2 over \mathbf{M}_1 , i.e. decisive preference for model \mathbf{M}_2 . All sampling based approaches also found a preference for this model.

The posterior harmonic mean estimate overestimated the Bayes factor rather significantly, also the standard error is very large compared to the other estimation methods. Combined, this indicated a rather bad approximation and reliability, which is in agreement with the general issues of this sampling method. Chib's method in our case performed worst. While the Bayes factor was still very strong in favor of model \mathbf{M}_2 , it underestimated the true value by a factor of 2. This was mostly due to a systematic underestimation of both log marginal likelihoods, e.g. for model \mathbf{M}_2 , all sampling results were < -21.4 while the true value was $\log(p(\mathbf{Y}|\sigma, \mathbf{M}_2)) = -20.12$. This seems to be a systematic issue, since the sampling passed Geweke's convergence criterion with all p-values larger than 0.98. Furthermore, the mean of the samples for model \mathbf{M}_2 for example could be compared to the analytical posterior distribution. We found that the sample means of -2.3007 and 2.1318 agree very well with the analytical values -2.2994 and 2.1260 . Also the sample covariance matrix $[0.9915, -0.0053; -0.0053, 0.4974]$ agreed very well with the analytical one $[1, 0; 0, 0.5]$. We conclude that Chib's method seems to suffer from severe numerical issues and should thus only be used very carefully.

The prior arithmetic mean and thermodynamic integration performed best with mean Bayes factors of 138.76 and 140.66, respectively. The good performance of the prior arithmetic mean was certainly due to the simplicity of the model selection problem, as well as the goodness of the prior. To further illustrate this, consider the following: let the prior information for the parameters have not the same σ as the data, but a different σ_p . Then the marginal likelihoods and the Bayes factor change accordingly. But more importantly, we can showcase the problems which might occur with the prior arithmetic mean. For $\sigma_p = \sigma$, all samples drawn from the prior have a non-zero likelihood value. However, if we increased σ_p to $\sigma_p = 20$ (i.e. if we weakened the prior information), around 22 % of the prior samples within both models had a likelihood value of zero within the limits of computational accuracy. Further increasing the hyper parameter for the prior to $\sigma_p = 50$ leads to an even worse situation: now around 63 % of the prior samples for model \mathbf{M}_1 and 77 % of prior samples for model \mathbf{M}_2 had a likelihood value of zero within the limits of computational accuracy. We expect this situation to worsen in larger models with more parameters due to Bellman's curse of dimensionality [5]. Thus in larger models a large number of samples might have to be drawn from the prior to ensure the

statistical power of the estimator. Thermodynamic integration performed very well in our scenario and also in the following larger model and we expect it to also perform well in other applications.

10.5 Model Selection for a Whole-Body Multi-compartmental Model

We now want to present a larger example for model selection in ODE systems, previously published in Schmidl et al. [40]. The example compares a model with 12 parameters to one with 15 parameters on the basis of 16 data sets. It comes from radiation protection where biokinetic ODE models are of crucial importance in dose estimation and further risk analysis for humans exposed to radioactive substances. More correctly, we examined the processing of zirconium in the human body after intake by ingestion. The models in question provided limiting values of detrimental effects and built the basis for applications in internal dosimetry, the prediction for radioactive zirconium retention in various organs as well as retrospective dosimetry.

Mathematically, the models for zirconium processing are multi-compartmental models, corresponding to linear ODE systems. In a compartmental model, all major human organs are represented as separate compartments representing a kinetically homogeneous amount of radionuclides [16, 21]. Transfer between these compartments is governed by the law of mass balance and described by time-constant transfer rates, which are the parameters that have to be inferred to fit the model to the data. Since this linear structure is rather straightforward, it can easily be interpreted, however, determining the exact interaction mechanisms is a challenging task. In the present case, there exist two competing models as suggestions for these interaction mechanisms. For the first time, *in vivo* experimental data with measurements in humans from blood plasma and urine were now available. Applying thermodynamic integration for the computation of Bayes factors, we could establish dominance of one model over the other. Furthermore, the availability of samples from the posterior distribution of the models allowed for the prediction of accretion in compartments where no direct measurements were technically possible in humans.

The first model is well established in the community and was put forward by the International Commission on Radiological Protection (ICRP) [16]. The transfer rates for this model were mostly derived from animal data and yielded extensive prior information for our inference. The Helmholtz Zentrum München (HMGU) recently published another, physiologically more plausible biokinetic model [12]. It is the first model based on measurement data in humans, taken in 16 investigations from 12 healthy human subjects. *In vivo* measurements were taken from blood plasma and urine of up to 100 days after ingestion by application of the double tracer technique. More details on the measurement process as well as a global statistical uncertainty and sensitivity analysis of this HMGU model can be found in the respective publications [29, 30].

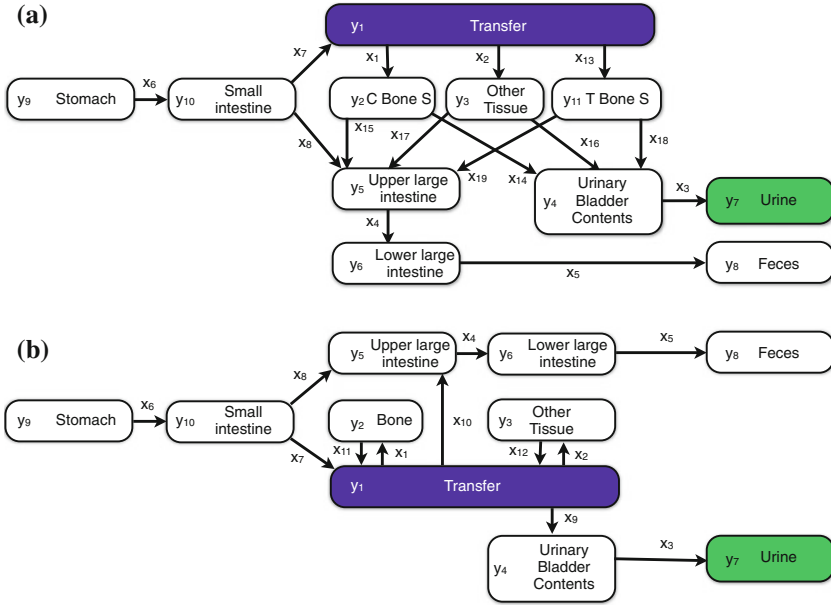


Fig. 10.3 Models for zirconium processing. **a** ICRP model. This model contains eleven compartments y_1, \dots, y_{11} and 15 time-constant transfer rates $x_1, \dots, x_8, x_{13}, \dots, x_{19}$. **b** HMGU model. This model contains ten compartments y_1, \dots, y_{10} and twelve transfer rates x_1, \dots, x_{12} . In both models zirconium enters the body in the stomach compartment y_9 and is processed through the system until it reaches either one of the two final compartments urine, y_7 , or feces, y_8 . The colored compartments y_1 and y_7 are corresponding to the ones where measurements are taken

10.5.1 Inference Setup

The first of the two compartmental models under examination here was recommended by the ICRP in [16, 17, 20] (Fig. 10.3a). The model contains eleven compartments, which are linked through 15 transfer rates. Since zirconium is ingested, it enters the body through the stomach compartment y_9 and is processed until it reaches one of the two final compartments urine, y_7 , or feces, y_8 . The HMGU model [12] differs from this model: it contains only ten compartments and twelve transfer rates, since the physiologically questionable distinction between the two bone compartments of the ICRP model, trabecular bone surface and cortical bone surface, was abolished in this model. Furthermore, most mass transfers are now mediated by the transfer compartment representing the blood plasma instead of direct transfers, which is physiologically more plausible. To better represent that some rates are shared between the models, we use the notation x . Both models share eight transfer rates, which are thus denoted x_1, \dots, x_8 . The additional rates present in only one of the models then have a unique index x_9, \dots, x_{19} , beginning with the HMGU model specific rates.

Both compartmental models correspond to systems of coupled linear first-order ordinary differential equations. For the time-dependent concentration $y_j(t)$ in compartment y_j , the rate of change is given by

$$\frac{d}{dt}y_j(t) = \sum_{\alpha \in \mathcal{A}_{y_j}^+} x_\alpha y_{[x_\alpha]}(t) - \sum_{\beta \in \mathcal{A}_{y_j}^-} x_\beta y_j(t) \quad (10.32)$$

For this, we introduce the set $\mathcal{A}_{y_j}^+$, containing all indices to transfer rates x_α flowing INTO compartment y_j and the set $\mathcal{A}_{y_j}^-$ containing all indices to transfer rates flowing OUT of compartment y_j . The compartment $y_{[x_\alpha]}$ is then the one connected to y_j by the transfer rate x_α . We illustrate this by an example: consider the HMGU model. Here we have $\mathcal{A}_{y_5}^+ = \{8, 10\}$, $y_{[x_8]} = y_{10}$ and $y_{[x_{10}]} = y_1$. Initial concentrations are needed for a unique solution of the ODEs, thus we choose $y_9(0) = 100\%$ and $y_{j \neq 9}(0) = 0\%$, since we assume that all zirconium is in the stomach compartment at the start of the investigation. The detailed ODEs can be found in the supplementary material of the original publication [40].

Zirconium was measured in plasma and urine through the double tracer technique in 16 investigations [12, 13]. The raw data tracer concentrations were then normalized to the respective investigation-specific tracer amount to yield 100% at $t = 0$ in the stomach compartment y_9 . For the development of the model, the transfer compartment was taken to be identical with blood plasma, the measured concentrations were then expressed as % per kg plasma. Absolute concentrations were obtained by scaling with the total amount of plasma in the body [2]. The measurements in urine correspond to an excretion rate in % per day.

10.5.2 Model Likelihoods

For each investigation $z = 1, \dots, Z = 16$, we found the likelihood by assuming Gaussian noise on the ODE solution $\mathbf{c}_{x_i}(t)$ of the differential equation for any of the two models \mathbf{M}_1 or \mathbf{M}_2 and some corresponding parameter vector x^i , where the model index $i \in \{1, 2\}$. Here, \mathbf{M}_1 is the HMGU model and \mathbf{M}_2 the ICRP model. Corresponding to the notation in Fig. 10.3a, b, $x^1 = (x_1, \dots, x_{12})$ and $x^2 = (x_1, \dots, x_8, x_{13}, \dots, x_{19})$. Then the data for investigation z was given by the measurements in plasma and urine

$$\mathbf{Y}_z = (y_1^{z,1}, y_1^{z,2}, \dots, y_1^{z,N_z^b}, \dot{y}_7^{z,1}, \dot{y}_7^{z,2}, \dots, \dot{y}_7^{z,N_z^u}) \quad (10.33)$$

While $y_1^{z,i}$ indicated measurements in plasma, i.e. in the transfer compartment y_1 , $\dot{y}_7^{z,i}$ designated measurements of the excretion rate in the urine compartment y_7 . There are N_z^b measurements in plasma and N_z^u measurements in urine for investigation z . Assuming normally distributed noise, the likelihood was then given for each investigation z and for model i by

$$p(\mathbf{Y}_z|x^i, M_i) = \underbrace{\prod_{\alpha=1}^{N_z^b} \mathcal{N}\left(y_1^{z,\alpha} | c_{x^i}^b(t_\alpha), \sigma_z^b\right)}_{p^{(b)}(\mathbf{Y}_z|x^i, M_i)} \underbrace{\prod_{\beta=1}^{N_z^u} \mathcal{N}\left(\dot{y}_7^{z,\beta} | \frac{d}{dt} c_{x^i}^u(t_\beta), \sigma_z^u\right)}_{p^{(u)}(\mathbf{Y}_z|x^i, M_i)}.$$

Again $c_{x^i}^b(t_\alpha)$ was the ODE solution for the parameter vector x^i in the transfer compartment y_1 at time point t_α , corresponding to the measurement at $y_1^{z,\alpha}$. Accordingly $\frac{d}{dt} c_{x^i}^u(t_\beta)$ was the derivative of the solution for the urine compartment y_7 at time point t_β , corresponding to the measurement $\dot{y}_7^{z,\beta}$. The standard deviations of the normal distributions for plasma, σ_z^b , and for urine, σ_z^u , were fitted for each investigation separately by applying the simulated annealing algorithm [25] before starting the MCMC sampling process. This error model corresponds to the combined strength of all deviations from the “true” ODE solution, which include (possibly amongst others) measurement error as well as natural internal fluctuations not considered by an ODE approach. With these assumptions, both models were able to fit the data in principle, justifying our ODE approach with additive normally distributed noise, see also Fig. 10.5. It is however still important to account for biological variability between the individual investigations, for which we accounted by fitting different σ_z^b and σ_z^u for each investigation z and thus get investigation-specific likelihoods. This lead to individual credible intervals for each parameter in each investigation in the MCMC sampling procedure later on. In contrast to this individual treatment, it also makes sense to consider the complete (i.e. concatenated) data. Then the likelihood was given by $p_{ALL}(\mathbf{Y}|x^i, M_i) = \prod_{z=1}^{16} p(\mathbf{Y}_z|x^i, M_i)$, with $\mathbf{Y} = \{\mathbf{Y}_1, \dots, \mathbf{Y}_{16}\}$ and fitting investigation independent σ^b and σ^u .

For the calculation of the likelihood, the ODE had to be solved depending on x^i . Since the ODE is of first order, it can be reformulated as

$$\frac{d\mathbf{y}_{x^i}(t)}{dt} = A(x^i) \cdot \mathbf{y}_{x^i}(t), \quad (10.34)$$

with $\mathbf{y}_{x^i}(t)$ the vector of all the compartments of model i and the time independent matrix $A(x^i)$ representing all the interactions between these compartments, depending on the transfer rate values x^i . First order ODEs can then be solved analytically:

$$\mathbf{y}_{x^i}(t) = e^{A(x^i)t} \cdot \mathbf{y}_{x^i}(t=0). \quad (10.35)$$

The matrix exponential $e^{A(x^i)t}$ can be computed by eigenvalue decomposition using e.g. MATLAB's `eig` function.

This now enabled us to compute a Bayes factor for each investigation z

$$B_{1,2}^z = \frac{p(\mathbf{Y}_z|\mathbf{M}_1)}{p(\mathbf{Y}_z|\mathbf{M}_2)}. \quad (10.36)$$

as well as an overall Bayes factor

$$B_{1,2} = \frac{p_{ALL}(\mathbf{Y}|\mathbf{M}_1)}{p_{ALL}(\mathbf{Y}|\mathbf{M}_2)}. \quad (10.37)$$

10.5.3 Prior Information

Models for zirconium processing have been used for a few decades already, and quite a large number of animal studies has been held. From these, comprehensive prior information for both models could be curated. The priors for each single transfer rate were given as triangular, normal or lognormal distributions with known hyperparameters. Of the eight transfers present in both models, only x_8 had a different distributions in the ICRP and HMGU model. Each univariate prior distribution was truncated at zero. For details, we again refer to the original publication [40]. In contrast to the investigation specific Bayes factors, the prior information was naturally the same for each investigation, since it represents information from a large number of preceding examinations and was not specific to the present investigations.

10.5.4 Copula Based Monte Carlo Sampling

In order to be able to do model selection via thermodynamic integration, we needed to be able to sample from the model, investigation and temperature specific distribution. For this we used a sophisticated MCMC method named copula-based independence/random walk Metropolis-Hastings approach (CIMH) [39] with the setup specified in [40]. As already mentioned, we used simulated annealing to find the maximum a posteriori estimate and used this as starting point for the sampling, enabling us to skip the burn-in phase. For this application, we chose to apply thinning by the autocorrelation based Effective Sample Size (ESS) [33]. Though generally not necessary, we used this as an additional quality insurance. The ESS corresponds to the number of samples that are left if a Markov chain is thinned such that two consecutive samples can be considered approximately independent. Our sampling algorithm CIMH was able to provide a high ESS at simultaneously high acceptance rates. From all required distributions we generated 30,000 proposals.

10.5.5 Model Selection Results and Discussion

Since the experimental data as basis for the model selection came from 16 investigations, one can ask if the models should be compared based on the complete data, yielding one overall Bayes factor, or on each dataset separately, yielding 16




Fig. 10.4 The experimental data: Plasma concentrations and excretion rate in urine for all 16 investigations on log–log-timescale

Bayes factors, which cannot be directly compared. When taking a closer look at the data (Fig. 10.4), we saw that all investigations exhibit a pulse-like time course in the plasma measurements, while the excretion rates in urine pointed more to an exponential decay behavior. Despite these shared characteristics, the actual zirconium tracer concentrations showed up to a 50-fold difference between maximum plasma concentrations, i.e. for investigation 10 (1.616 %) and 6 (0.033 %).

This already suggested that the investigations should be treated separately, since the differences in the concentrations propagate to differences in the transfer rates. To verify this, we did a pairwise comparison of the posterior samples marginal (corresponding to the temperature $\tau = 1$) by the Kolmogorov-Smirnov test. Since this test is univariate, we picked parameter x_7 in the ICRP model as example, as it directly affects the observed concentrations in plasma [30]. Except for one pair, all obtained p-values were $< 6 \times 10^{-8}$. We took this as a strong indication that all investigations should be treated separately.

However, for many applications of the models, reference values for an average subject are needed. This is why we also included the Bayes factor for the complete data in our analysis. The differences between the overall Bayes factor and the investigation specific ones can also be the basis for the study of influence factors like gender or weight.

As already mentioned, the analysis was based on 30,000 proposals for every one of the 30 temperature level in all 17 examined cases (one for each investigation and one for \mathbf{Y}). From the 510 sampling runs of the HMGU model, we achieved an average ESS of 5832 ± 405 (including one standard error), for the ICRP model 5808 ± 252 . This corresponded to high acceptance rates and a good capture of the power posterior by the sampling procedure.

The obtained posterior samples yielded credible intervals for the parameters as hand as well as an maximum a posteriori estimate based on the complete data, which can be used if single parameter values for an average subject are required. When propagating these posterior samples to the ODE solution, we saw (Fig. 10.5)

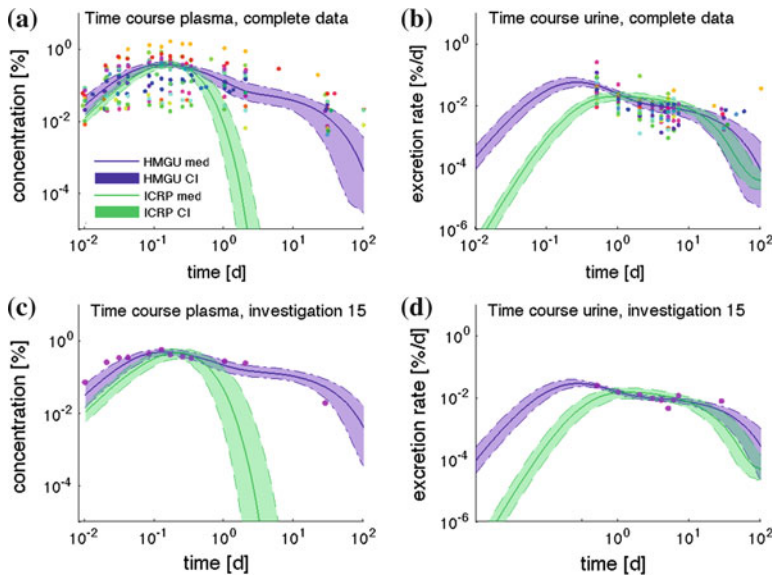


Fig. 10.5 Posterior time courses: Shown are the sample median (*solid line*) and 90 % credible interval (CI, *shaded area*) drawn from the time courses based on the $\tau = 1$ HMGU (*purple*) and ICRP (*green*) MCMC samples for **a** the complete plasma data, **b** the excretion rate in urine over time of the complete data, **c** as a single example the plasma data of investigation 15, and **d** the corresponding urinary excretion rate over time of investigation 15, all plotted on a log–log scale. The plasma plots were truncated at $1 \times 10^{-5}(\%)$ and urine plots at $1 \times 10^{-6}(\%/d)$. Note that the median and CI represent only the uncertainty in the parameters, in contrast to measurement uncertainty (not shown). *Colored markers* are the raw experimental data points. At each time point the median and the 90 % credible interval were computed pointwisely over all MCMC based solutions

that both models were in principle able to fit the measurement data, justifying our approach. While no rigorous model selection was possible merely from these fits, especially the plasma data already hinted at a better suitability of the HMGU model. In the urine data the difference between the two models was not as pronounced. We want to point out that the credible intervals in Fig. 10.5 represent only the uncertainty coming from the parameters, in contrast to the measurement uncertainties accounted for by the noise parameters σ_z^b and σ_z^u , which are not shown.

For the actual model selection, we now compared the HMGU and ICRP models based on both the complete data as well as the individual investigations, yielding 17 Bayes factors. We found that all Bayes factors favor the HMGU model, 14 out of 17 even decisively (Table 10.1).

Since the time courses already indicated that the HMGU model might be more suitable since it fitted the plasma data better, we computed additional Bayes factors based on either only the plasma or urine data. This corresponds to considering either only "the plasma likelihood" $p^{(b)}(\mathbf{Y}_z|x^i, M_i)$ or "the urine likelihood" $p^{(u)}(\mathbf{Y}_z|x^i, M_i)$ from Eq. 10.33, where $i = 1, 2$ and $z = 1, \dots, 16$ and accordingly for the complete

Table 10.1 Bayes factors for the HMGU versus the ICRP model ($B_{1,2}^A$) for the individual investigations as well as for the complete data (ALL) and the according Bayes factors for the plasma ($B_{1,2}^P$) and urine ($B_{1,2}^U$) data

Inv.	$B_{1,2}^A$	$B_{1,2}^P$	$B_{1,2}^U$
1	<i>7.17×10^1</i>	<i>7.12×10^1</i>	<i>1.05</i>
2	<i>1.15×10^2</i>	<i>2.93×10^2</i>	<i>3.94×10^3</i>
3	<i>5.95×10^4</i>	<i>5.23×10^4</i>	<i>1.34</i>
4	<i>1.07×10^3</i>	<i>2.64×10^3</i>	<i>3.47×10^1</i>
5	<i>2.19×10^2</i>	<i>4.73×10^2</i>	<i>1.34×10^2</i>
6	<i>4.64×10^3</i>	<i>3.93×10^3</i>	<i>2.38×10^3</i>
7	<i>2.18×10^2</i>	<i>2.30×10^2</i>	<i>1.34×10^3</i>
8	<i>3.75×10^1</i>	<i>1.28×10^2</i>	0.22
9	<i>4.62×10^2</i>	<i>2.32×10^2</i>	0.18
10	<i>8.62×10^2</i>	<i>1.16×10^2</i>	0.20
11	<i>1.17×10^5</i>	<i>1.81×10^1</i>	<i>2.94×10^3</i>
12	<i>1.78×10^2</i>	<i>5.48</i>	<i>1.14×10^1</i>
13	<i>7.19×10^2</i>	<i>1.41×10^1</i>	<i>4.41</i>
14	<i>3.58×10^1</i>	<i>7.43</i>	<i>9.77</i>
15	<i>6.29×10^3</i>	<i>2.17×10^1</i>	<i>1.60×10^2</i>
16	<i>6.22×10^2</i>	<i>1.34×10^1</i>	<i>1.20×10^4</i>
ALL	<i>1.20×10^{11}</i>	<i>3.43×10^4</i>	<i>4.73×10^7</i>

Italic indicates a Bayes factor in favor of the HMGU model and bold a Bayes factor in favor of the ICRP model. The HMGU model is favored *substantially*, when $B_{1,2} > 3$ and *decisively*, when $B_{1,2} > 100$. Also, $1/B_{1,2} = B_{2,1}$

data. The Bayes factors support our theory that the plasma is fitted better by the HMGU model: all 17 Bayes factors based on the plasma data favored the HMGU model, in ten cases again decisively (Table 10.1, 3rd column). For the urine data, the situation was slightly more ambivalent, as here three investigations favor the ICRP model (Table 10.1, 4th column), but not decisively, while still eight Bayes factors favor the HMGU model decisively. All in all, we asserted that all decisive Bayes factors are in favor of the HMGU model, meaning that the ICRP model was decisively rejected in the majority of cases. Thus we conclude that the HMGU model is superior to the ICRP model for representing zirconium processing in the human body, both on an individual level as well as for an average subject represented by the complete data.

With this knowledge in mind, we now want to determine if this difference in the quality of the models can be used to make useful and distinguishable predictions of system behavior. In internal exposure monitoring, biokinetic models like the two under scrutiny here are often used to predict the organ retention or daily excretion of incorporated radionuclides [18], often in organs that are technically hard to measure. The solution of the model in each compartment gives the organ retention function. The organ doses are then directly computable from the integral of the radioactivity




Fig. 10.6 Retention of zirconium in bones: Shown are the median (*solid lines*) as well as 90% credible intervals (*shaded areas*) for the extrapolated retention of ^{95}Zr in the bone compartment(s) as predicted by the HMGU (*purple*) and ICRP (*green*) models, correcting for radioactive decay

of ^{95}Zr in source organs over e.g. 50 years with the inference results from our data from the non-radioactive ^{98}Zr .

We now want to compare the retention of ^{95}Zr in the bone compartment as predicted by the two models. For this we computed the 90% credible intervals for the time courses in the bone compartments, based on the posterior samples of the complete data on a longer timescale. Figure 10.6 immediately shows a significant difference between the two models, with much lower accretion in the HMGU model, where we added the two bone compartments of the ICRP model together. While the inference was based on a stable isotope of zirconium, we now assume that the radioactive ^{95}Zr behaves the same, however now radioactive decay with half-life $64.032d$ [19] has to be taken into account. The reduction of retention in the HMGU model could have significant effects on monitoring, for which at the moment the ICRP model is used.

10.6 Conclusions

In this chapter, we have tried to advocate the use of Bayes factors for model selection. For this purpose, we have first demonstrated several different model selection methods on a simple, analytically tractable example. We have found in this example

that while likelihood based methods like the AIC also give correct results, the Bayes factor is easier interpretable. For the calculation of the Bayes factor as such, we have found that both our preferred method thermodynamic integration as well as the prior arithmetic mean perform very well in our example. We have also shortly shown that the prior arithmetic mean strongly depends on the quality of prior information. We went on to demonstrate the computation of the Bayes factor with thermodynamic integration on a real world example from radiation science.

In summary, we could show there that the newer HMGU model was unequivocally superior with 14 of 17 Bayes factors being decisive when compared to the well-established ICRP model. Also, when restricting the data on plasma and urine measurements only, we found that the HMGU model was clearly favored. The HMGU model thus best covers human data. We found a significant difference between the two models regarding the predicted accumulation of zirconium in bones, which might be experimentally tested in animal studies in the future.

More generally, we believe that the presented methodology is suitable for any ODE-based model selection task and that Bayes factors should be the method of choice for doing model selection, a view that is also increasingly shared in the current literature [9, 15, 24]. Especially tasks like the modeling of protein signaling, gene regulation, or drug processing [26], nowadays frequently put forward in systems biology [4, 36] or pharmacogenetics [45] could benefit from these approaches.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

1. Akaike, H.: A new look at the statistical model identification. *Automatic Control, IEEE Transactions on* (1974)
2. Alberts, B., Johnson, A., Lewis, J., Raff, M., Roberts, K., Walter, P.: *Molecular biology of the cell*, 4th edn. Garland Science, New York (2002)
3. Bachmann, J., Raue, A., Schilling, M., Böhm, M.E., Kreutz, C., Kaschek, D., Busch, H., Gretz, N., Lehmann, W.D., Timmer, J., Klingmüller, U.: Division of labor by dual feedback regulators controls jak2/stat5 signaling over broad ligand range. *Mol. Syst. Biol.* **7**(1), 2011
4. Becker, V., Schilling, M., Bachmann, J., Baumann, U., Raue, A., Maiwald, T., Timmer, J., Klingmüller, U.: Covering a broad dynamic range: information processing at the erythropoietin receptor. *Science* **328**(5984), 1404–1408 (2010)
5. Bellman, R.: Dynamic programming and lagrange multipliers. In: Bellman, E. (Ed.) *The Bellman Continuum: A Collection of the Works of Richard*, pp. 49 (1986)
6. Box, G.E.P., Draper, N.R.: *Empirical model-building and response surfaces*. Wiley (1987)
7. Calderhead, Ben, Girolami, Mark: Estimating Bayes factors via thermodynamic integration and population MCMC. *Comput. Stat. Data Anal.* **53**(12), 4028–4045 (2009)
8. Chib, S., Jeliazkov, I.: Marginal likelihood from the metropolis-hastings output. *J. Am. Stat. Assoc.* **96**(453), 270–281 (2001)

9. Eydgahi, H., Chen, W.W., Muhlich, J.L., Vitkup, D., Tsitsiklis, J.N., Sorger, P.K.: Properties of cell death models calibrated and compared using bayesian approaches. *Mol. Syst. Biol.* **9**(1), (2013)
10. Friel, N., Pettitt, A.N.: Marginal likelihood estimation via power posteriors. *J. Royal Stat. Soc.: Ser. B (Stat. Methodol.)* **70**(3), 589–607 (2008)
11. Gelman, A., Meng, X.L.: Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Stat. Sci.*, 163–185 (1998)
12. Greiter, M.B., Giussani, A., Höllriegel, V., Li, W.B., Oeh, U.: Human biokinetic data and a new compartmental model of zirconium - a tracer study with enriched stable isotopes. *Sci. Total Environ.* **409**, 3701–3710 (2011)
13. Greiter, M.B., Höllriegel, V., Oeh, U.: Method development for thermal ionization mass spectrometry in the frame of a biokinetic tracer study with enriched stable isotopes of zirconium. *Int. J. Mass. Spectrom.* **304**(1), 1–8 (2011)
14. Haario, H., Saksman, E., Tamminen, J.: An adaptive Metropolis algorithm. *Bernoulli* **7**(2), 223–242 (2001)
15. Hug, S., Raue, A., Hasenauer, J., Bachmann, J., Klingmüller, U., Timmer, J., Theis, F.J.: High-dimensional Bayesian parameter estimation: Case study for a model of JAK2/STAT5 signaling. *Math. Biosci.* (2013)
16. ICRP: Age-dependent Doses to Members of the Public from Intake of Radionuclides (Part 1). ICRP Publication 56. Number Ann. ICRP, Vol. 20, No. 2. Pergamon Press, Oxford, 1989
17. ICRP: Age-dependent Doses to Members of the Public from Intake of Radionuclides (Part 2: Ingestion dose coefficients). ICRP Publication 67. Number Ann. ICRP, Vol. 23, No. 3–4. Pergamon Press, Oxford, 1993
18. ICRP: Individual Monitoring for Internal Exposure of Workers. ICRP Publication 78. Number Ann. ICRP, Vol. 27, No. 3–4. Pergamon Press, Oxford, 1997
19. ICRP: Nuclear Decay Data for Dosimetric Calculations. ICRP Publication 107. Number Ann. ICRP, Vol. 38, No. 3. Pergamon Press, Oxford, 2008
20. ICRP: Report on the Task Group on Reference Man. ICRP Publication 23. Pergamon Press, Oxford, 1975
21. Jacquez, J.A.: Compartmental analysis in biology and medicine, 3rd edn. BioMedware, Ann Arbor, MI (1996)
22. Jeffreys, H.: The theory of probability. Oxford University Press, 1998
23. Kirk, P., Silk, D., Stumpf, M.P.H.: Reverse engineering under uncertainty. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
24. Kirk, P., Thorne, T., Stumpf, M.P.: Model selection in systems and synthetic biology. *Curr. Opin. Biotechnol.* 1–8 (2013)
25. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by simulated annealing. *Science* **220**, 671–680 (1983)
26. Krumsiek, J., Pölsterl, S., Wittmann, D.M., Theis, F.J.: Odefy-from discrete to continuous models. *BMC Bioinform.* **11**(1), 233 (2010)
27. Lagani, V., Triantafyllou, S., Ball, G., Tegnér, J., Tsamardinos, I.: Probabilistic computational causal discovery for systems biology. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
28. Lartillot, N., Philippe, H.: Computing Bayes factors using thermodynamic integration. *Syst. Biol.* **55**(2), 195–207 (2006)
29. Li, W.B., Greiter, M., Oeh, U., Hoeschen, C.: Reliability of a new biokinetic model of zirconium in internal dosimetry part i, parameter uncertainty analysis. *Health Phys.* **101**(6), 660–676 (2011)
30. Li, W.B., Greiter, M., Oeh, U., Hoeschen, C.: Reliability of a new biokinetic model of zirconium in internal dosimetry part ii, parameter sensitivity analysis. *Health Phys.* **101**(6), 676–692 (2011)
31. Marin, J.M., Robert, C.P.: Bayesian core: a practical approach to computational Bayesian statistics. Springer, 2007

32. Myung, I.J., Pitt, M.A.: Applying occam's razor in modeling cognition: A bayesian approach. *Psychonomic Bull. Rev.* **4**(1), 79–95 (1997)
33. Neal, R.M.: Probabilistic inference using markov chain monte carlo methods. Technical report, University of Toronto, 1993
34. Neal, R.: The Harmonic Mean of the Likelihood: Worst Monte Carlo Method Ever. <http://radfordneal.wordpress.com/2008/08/17/the-harmonic-mean-of-the-likelihood-worst-monte-carlo-method-ever>, (2008)
35. Newton, M.A., Raftery, A.E.: Approximate Bayesian inference with weighted likelihood bootstrap. *J. Royal Stat. Soc. Ser. B* (1994)
36. Raia, V., Schilling, M., Böhm, M., Hahn, B., Kowarsch, A., Raue, A., Sticht, C., Bohl, S., Saile, M., Möller, P., Gretz, N., Timmer, J., Theis, F., Lehmann, W.-D., Lichter, P., Klingmüller, U.: Dynamic mathematical modeling of il13-induced signaling in hodgkin and primary mediastinal b-cell lymphoma allows prediction of therapeutic targets. *Cancer Res.* **71**(3), 693–704 (2011)
37. Robert, C.P., Casella, G.: *Monte Carlo Statistical Methods*. Springer, New York, 2004
38. Schliemann-Bullinger, M., Fey, D., Bastogne, T., Findeisen, R., Scheurich, P., Bullinger, E.: The experimental side of parameter estimation. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
39. Schmidl, D., Czado, C., Hug, S., Theis, F.J.: A vine-copula based adaptive mcmc sampler for efficient inference of dynamical systems. *Bayesian Anal.* **8**(1), 1–22 (2013)
40. Schmidl, D., Hug, S., Li, W., Greiter, M.B., Theis, F.J.: Bayesian model selection validates a biokinetic model for zirconium processing in humans. *BMC Syst. Biol.* **6**(95), (2012)
41. Schmidl, D.: Bayesian model inference in dynamic biological systems using Markov Chain Monte Carlo methods. Ph.D. thesis, München, Technische Universität München, 2012
42. Schwarz, G.: Estimating the dimension of a model. *Annals Stat.* (1978)
43. Sunnåker, M., Stelling, J.: Model extension and model selection. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
44. Xu, T.-R., Vyshemirsky, V., Gormand, A., von Kriegsheim, A., Girolami, M., Baillie, G.S., Ketley, D., Dunlop, A.J., Milligan, G., Houslay, M.D., et al.: Inferring signaling pathway topologies from multiple perturbation measurements of specific biochemical species. *Sci. Signal.* **3**(113), ra20 (2010)
45. Zhao, W., Elie, V., Roussey, G., Brochard, K., Niaudet, P., Leroy, V., Loirat, C., Cochat, P., Cloarec, S., André, J.L., Garaix, F., Bensman, A., Fakhoury, M., Jacqz-Aigrain, E.: Population pharmacokinetics and pharmacogenetics of tacrolimus in de novo pediatric kidney transplant recipients. *Clin. Pharmacol. Ther.* **86**(6), 609–618 (2009)

Part IV
Sensitivity Analysis and Model Adaptation

Chapter 11

Sloppiness and the Geometry of Parameter Space

**Brian K. Mannakee, Aaron P. Ragsdale, Mark K. Transtrum
and Ryan N. Gutenkunst**

Abstract When modeling complex biological systems, exploring parameter space is critical, because parameter values are typically poorly known a priori. This exploration can be challenging, because parameter space often has high dimension and complex structure. Recent work, however, has revealed universal structure in parameter space of models for nonlinear systems. In particular, models are often *sloppy*, with strong parameter correlations and an exponential range of parameter sensitivities. Here we review the evidence for universal sloppiness and its implications for parameter fitting, model prediction, and experimental design. In principle, one can transform parameters to alleviate sloppiness, but a parameterization-independent information geometry perspective reveals deeper universal structure. We thus also review the recent insights offered by information geometry, particularly in regard to sloppiness and numerical methods.

Keywords Sloppiness · Hessian · Experimental design · Bayesian ensembles · Cost functions · Information geometry

B.K. Mannakee
Graduate Interdisciplinary Program in Statistics, University of Arizona,
Tucson, AZ, USA
e-mail: mannakee@email.arizona.edu

A.P. Ragsdale
Graduate Interdisciplinary Program in Applied Mathematics, University of Arizona,
Tucson, AZ, USA
e-mail: aragsdale@math.arizona.edu

M.K. Transtrum
Department of Physics and Astronomy, Brigham Young University, Provo, UT, USA
e-mail: mktranstrum@byu.edu

R.N. Gutenkunst (✉)
Department of Molecular and Cellular Biology, Tucson, AZ, USA
e-mail: rgutenk@email.arizona.edu

11.1 Introduction

Mathematical models of cell-signaling, metabolic, and gene networks play a critical role in developing mechanistic understanding of these networks [45]. Building models can be difficult, however, because such networks often have complex nonlinear dynamics and not all components may be known. In fact, important uses of network models are to infer network structure [19] or choose between hypotheses regarding network function [43]. (For more on the challenges in reverse engineering biological networks, see Chap. 2 in this volume [42].) Even when the network is well-known, however, modeling can still be difficult, because mechanistic models typically depend upon a large number of kinetic parameters [29, 52]. Such parameters are often unknown and are difficult to measure experimentally.

In this chapter, we review methods for exploring the spaces of model parameters and data, and we review recent work on *sloppiness*, a general property of complex nonlinear models. Sloppy models have highly anisotropic parameter and data spaces, with complex relationships between parameter values and model output. Sloppiness results in several difficulties for modelers, including that numerical tools used to estimate parameters can be slow, confidence intervals for parameter values and model output can be large, and experiments to improve the model can be difficult to design. We review recent work on how these challenges arise and how they can be overcome.

11.1.1 Parameter Space and Data Space

In a model with N parameters $\theta_1, \dots, \theta_N$, the parameter space encompasses the set of all possible values for each parameter. Most commonly, parameters are real numbers, in which case the N -dimensional parameter space is a subspace of \mathbb{R}^N . A particular realization of parameter values $\theta = [\theta_1, \theta_2, \dots, \theta_N]$ is a vector representing a single point in parameter space. For biological models, in which parameters may have different units and scales that differ by orders of magnitude, it is convenient to consider logarithmic rather than absolute parameter values, so in this chapter when we speak of parameters and parameter space we are always referring to log-parameters. One can think of the model as a function mapping points or regions in parameter space to output values, or points in data space [18], and a general problem in systems biology is to understand this mapping.

A complete description of this map is a useful mathematical tool. For instance, once the mapping is understood, it is easy to enumerate the possible outcomes a model can generate. Thus, given a model that matches experimental data, one can generate hypotheses about other states the system might enter and perform experiments to look for those states [61]. Additionally, estimating parameter values by fitting them to data [40] employs the reverse map, in that the modeler seeks the point or region in parameter space that maps to the point in data space closest to the measured data. These maps are commonly constructed and analyzed using a cost function.

11.1.2 Cost Functions

In the context of fitting models to data, the cost function measures the difference between the model output for a given set of parameters and the data that is being fit. As such, the cost function acts as a map between parameter space and data space, structuring the parameter space in such a way that moves in parameter space correspond to changes in model fit. This structure is often called the cost landscape, and we will use these concepts, of maps and landscapes, interchangeably here. The most common cost function used in data fitting, and the one we will focus on in this chapter, is the least-squares cost function [62]. (For more on cost functions, see Chap. 7 in this volume [60]). Given a model $y(\theta, t)$ with parameter vector θ we define the least-squares cost function as

$$\begin{aligned} C(\theta) &\equiv \frac{1}{2} \sum_s \sum_c \sum_{T_c} \left[\frac{y_{s,c}(\theta, t) - Y_{s,c}(t)}{\sigma_{s,c,t}} \right]^2 \\ &= \frac{1}{2} \sum_s \sum_c \sum_{T_c} r_{s,c,t}^2(\theta) = \frac{1}{2} \sum_{m=1}^M r_m^2(\theta), \end{aligned} \quad (11.1)$$

which is half the squared difference over M data points collected for species s under experimental conditions c at time points T_c , or the sum of squares of the M normalized residuals r between model and data. Measurement uncertainty for each data point is denoted $\sigma_{s,c,t}$. The sum over T_c can be replaced with an integral when fitting continuous data. The best-fit parameter vector θ^* occurs at the global minimum of C .

Most commonly, the model residuals are assumed to be independent and normally distributed. The probability density that the model will produce the data given parameters θ is then

$$P(\text{Data}|\theta) = \prod_s \prod_c \prod_{T_c} \frac{1}{\sqrt{2\pi}\sigma_{s,c,t}} \exp \left[-\frac{1}{2} \left(\frac{y_{s,c}(\theta, t) - Y_{s,c}(t)}{\sigma_{s,c,t}} \right)^2 \right]. \quad (11.2)$$

In statistics, this probability density is called the likelihood of θ [12]. Taking the negative logarithm of the likelihood function yields the least-squares cost function $C(\theta)$ (Eq. 11.1). Thus, minimizing the cost function to find the best-fit parameters θ^* is equivalent to maximizing the likelihood, so the best-fit parameter vector inherits the statistical properties of the maximum likelihood estimator [26]. This statistical connection, arising from the assumption of normally-distributed residuals, makes the sum-of-squares cost function $C(\theta)$ particularly useful in describing the structure of parameter space. Other assumptions about the distribution of residuals are, however, possible and imply different cost functions. Note that much of what we discuss in this chapter has only been shown for sums-of-squares cost functions.

11.2 Multivariate Sensitivity Analysis

We use sensitivity analysis to explore parameter space, observing how model output changes as model parameters vary. In systems biology, sensitivity analysis is commonly used to quantify the uncertainty associated with both best-fit parameter values, and new model predictions generated using those parameter values. However, methods used in sensitivity analysis also provide additional useful information about parameter space. In this chapter, we are particularly interested in the correlation structure of parameter space, i.e. the relationships among combinations of parameters. Ultimately, we will define sloppiness in terms of these correlations. In this section, we describe one local and three global methods of multivariate sensitivity analysis. We then use an example model to illustrate how sensitivity analysis reveals parameter correlations and how those correlations affect the precision of parameter inference.

11.2.1 Local (The Hessian)

The sensitivity of the fit of the model $y(\theta, t)$ to a given data set is determined by how quickly the cost function $C(\theta)$ increases away from the best-fit parameters θ^* . A useful local perspective on the cost landscape is given by a quadratic expansion of the cost function:

$$C(\theta) \approx C(\theta^*) + \nabla C(\theta^*)(\log \theta^* - \log \theta) + \frac{1}{2}(\log \theta^* - \log \theta)^\top H(\log \theta^* - \log \theta) . \quad (11.3)$$

The gradient $\nabla C(\theta^*)$ of the cost function is, by definition, zero at the best-fit parameters θ^* . The $N \times N$ Hessian matrix H is defined as

$$H_{i,j} \equiv \left. \frac{\partial^2 C}{\partial \log \theta_i \partial \log \theta_j} \right|_{\theta=\theta^*} . \quad (11.4)$$

Because it involves second derivatives, the Hessian can be difficult to calculate. If the residuals r are small at the best-fit (i.e., the fit to the data is good) then the Hessian

can be well-approximated by $H \approx J^T J$. The Jacobian matrix J is the $M \times N$ matrix of partial derivatives of the residuals with respect to the parameters:

$$J_{m,n} = \frac{\partial r_m}{\partial \log \theta_n} . \tag{11.5}$$

The first-derivatives in the Jacobian can be evaluated by finite differences or, for ordinary-differential equation (ODE) models, by integrating sensitivity equations.

The Hessian describes the quadratic behavior of the cost function C near the point θ^* , so analyzing the Hessian corresponds to approximating the level curves of C as N -dimensional ellipsoids in parameter space (Fig. 11.1). The Hessian matrix is positive




Fig. 11.1 Local and global sensitivity analysis. **a** Local analysis around the best fit point (*black dot*). Ellipses are curves of constant cost calculated from the Hessian matrix. Sensitivity in each direction is proportional to the widths of the curves. **b** Scanning parameters along each axis around the best fit point. **c** Latin Hypersquare scan with uniform priors. Every equal probability bin of each parameter distribution is sampled exactly once. **d** Bayesian parameter vector ensemble for this quadratic model.

definite and symmetric, so it has real eigenvalues λ and eigenvectors v corresponding to the principle axes of those ellipsoids. Additionally, the relationship between the negative log-likelihood $-\log P(\text{Data}|\theta)$ and $C(\theta)$ (Sect. 11.1.2) implies that inverting the Hessian in Eq. 11.4 gives an asymptotic approximation of the covariance matrix of the parameters [22]. Thus, the covariance matrix has eigenvectors v and eigenvalues λ^{-1} , and the widths of the principle axes of the ellipsoids are proportional to $\lambda^{-\frac{1}{2}}$ (Fig. 11.1a).

11.2.2 Global

Local sensitivity analysis accurately measures model sensitivity when parameter space is linear and smooth, so the cost minimum is well defined, but these conditions are not guaranteed to hold in systems biology models. In particular, the relationship between parameters may be nonlinear near the best fit, so the quadratic map imposed by the Hessian may be a poor approximation to the actual cost surface. Figure 11.2b shows such a cost surface, in which strong nonlinearities cause the quadratic approximation to overestimate the variability of the parameters. Moreover, some models may have rough parameter spaces with multiple minima of similar model behavior separated by ridges [25]. In such a landscape, local sensitivity analysis can be misleading, because steep curvatures near a local minimum may obscure the true shape of the parameter space. Global methods of sensitivity analysis address these problems by sampling parameter space in a finite neighborhood around the best fit. Broadly, such methods fall into two categories, scanning methods and Bayesian methods.

Scanning methods sample parameters without regard to the data and look for correlations between locations in parameter space and the model behavior or value of the cost function at those locations. Bayesian methods sample from the posterior distribution of the parameters given the data and use those samples to make inferences about the sensitivity of the model. The challenge in both cases is to generate a sufficiently dense sample of the parameter space that valid inferences can be made about the sensitivity of the model to parameter changes. Here we describe two parameter scanning and two Bayesian methods that are frequently used for multivariate sensitivity analysis in biological modeling.

11.2.2.1 Parameter Scanning

One way to generate parameter sets is to simply scan the parameter space, varying one parameter at a time in small intervals over a range of values (Fig. 11.1b). Only in the unusual case where the principal axes of the model sensitivity line up with the parameter axes will this method provide an accurate measure of the uncertainty in parameter estimates, because it ignores correlations between parameters. Scanning combinations of parameters to capture those correlations is, however, often

prohibitive for large models, because the number of points needed to define a grid in N dimensions grows exponentially with N .

Latin Hypercube Sampling (LHS), a generalization of the Latin Square experimental design, is a method for sampling parameter space in such a way that correlations between parameters can be uncovered. Marino et al. [50] describe an application of LHS in the context of sensitivity analysis, in which each parameter is assigned a probability distribution. These probability distributions incorporate prior information about the range of values a parameter can take, are often normal or uniform, and need not be the same for every parameter. Each of the probability distributions is divided into B equal probability bins, and B parameter vectors are generated by randomly sampling one bin from each parameter distribution without replacement, keeping track of which bin each value came from. The result is a group of parameter vectors such that each value for a given parameter was drawn from a different part of its distribution (Fig. 11.1c). The cost function is evaluated for each of these parameter sets and the correlation between costs and bins describes the sensitivity of the model. Computing partial correlations among parameter combinations reveals the correlation structure of the parameter space. While LHS is computationally efficient due to the Latin Square randomization, its use in analyzing biological models requires special care because nonlinearities in parameter space can render correlation analysis inaccurate [50]. For other approaches to parameter scanning, see Chap. 13 in this volume [77].

11.2.2.2 Bayesian Ensembles

Parameter scanning methods sample parameter vectors without regard to the data. Thus, the resulting sample may contain many vectors that poorly fit the data and add little to our understanding of the relevant distribution of parameters. Bayesian approaches maximize information about the distribution of parameters around the best fit by sampling densely in areas corresponding to good fits and sparsely elsewhere. Bayesian Markov-chain Monte Carlo (MCMC) walks through parameter space have been widely used in systems biology to construct ensembles of parameter sets [9, 13, 14, 24, 27, 79], and flexible approximate Bayesian methods have recently been developed. (For more on Bayesian approaches to sampling parameters and choosing among models, see Chaps. 9 and 10 in this volume [38, 66].)

The goal of Bayesian MCMC is to sample from the posterior distribution $P(\theta|\text{Data})$ of parameter sets given the observed data. From Bayes' rule:

$$P(\theta|\text{Data}) = \frac{P(\text{Data}|\theta) P(\theta)}{P(\text{Data})}, \quad (11.6)$$

where $P(\text{Data}|\theta)$ is the likelihood defined in Eq. 11.2, $P(\theta)$ is the prior probability of the parameters, and $P(\text{Data})$ is the evidence for the data. $P(\text{Data})$ is often

difficult to calculate but is in many cases an unimportant normalization, leading to the proportionality:

$$P(\theta|\text{Data}) \propto P(\theta) P(\text{Data}|\theta) . \quad (11.7)$$

This proportionality allows a relative posterior probability to be calculated for any parameter set in terms of the likelihood and the prior. As we saw in Sect. 11.1.2, the likelihood can itself often be calculated in terms of the least squares cost function. The prior distribution reflects pre-existing knowledge of the distribution of parameter values, often from other experiments or analogy with similar molecular parameters. Early work focused on uniform priors [9, 13, 14], while more recent work employs log-normal [24, 27, 32] or gamma [79] priors.

The Markov chain is usually started at the best-fit parameter set and allowed to walk through parameter space sampling the posterior distribution of θ . At each step of the walk the Metropolis-Hastings criterion [17] is applied, such that at the j th step a new random vector θ_{test} is generated and

$$\theta_{j+1} = \begin{cases} \theta_{\text{test}} , & \text{with probability } \alpha \\ \theta_j , & \text{with probability } 1 - \alpha , \end{cases} \quad (11.8)$$

where

$$\alpha = \min \left\{ 1 , \frac{P(\theta_{\text{test}}|\text{Data})}{P(\theta_j|\text{Data})} \right\} . \quad (11.9)$$

The walk thus always accepts moves to parameter sets with higher posterior probability and sometimes accepts moves to parameter sets with lower posterior probability. This random walk generates an ensemble of parameter vectors that converges to the posterior distribution [17]. The marginal distributions for each parameter measure the sensitivity of the fit to changes in that parameter, integrating over changes in the other parameters, and provide confidence intervals for the best-fit value. The covariance matrix of the ensemble describes the correlation structure of the cost landscape.

11.2.2.3 Approximate Bayesian Computation

Approximate Bayesian Computation (ABC) allows sampling of approximate posterior parameter distributions when the likelihood function is analytically or computationally intractable. Both Markov-chain [11, 51] and Sequential Monte Carlo (SMC) [64, 67] methods exist. The SMC method uses sequential importance sampling [20] to shorten chain length by preventing the algorithm from getting stuck in areas of low probability [67]. ABC has the advantage that it can be used to sample the parameter space of stochastic models, in addition to deterministic models [67].

The main difference between ABC and likelihood-based MCMC is that at each step, rather than evaluating the likelihood of θ_{test} the algorithm instead generates a new simulated data set $Y'_{s,c}(t)$ from $y_{s,c}(\theta_{\text{test}}, t)$ and computes a distance metric

$\rho(Y_{s,c}(t), Y'_{s,c}(t))$. Possible distance metrics include the euclidian distance, squared distance, and total absolute deviation, among many others. At the j th step

$$\theta_{j+1} = \begin{cases} \theta_{\text{test}} , & \text{with probability } \alpha , \\ \theta_j , & \text{with probability } 1 - \alpha , \end{cases} \quad (11.10)$$


where

$$\alpha = \begin{cases} 0 , & \rho(Y_{s,c}(t), Y'_{s,c}(t)) > \varepsilon \\ \min \left\{ 1 , \frac{P(\theta_{\text{test}})}{P(\theta_j)} \right\} , & \rho(Y_{s,c}(t), Y'_{s,c}(t)) \leq \varepsilon , \end{cases} \quad (11.11)$$

and ε is some small number chosen to bound the acceptable distance between simulated and real data. This chain generates a collection of parameter vectors drawn from the joint distribution $P(\theta | \rho(Y_{s,c}(t), Y'_{s,c}(t)) \leq \varepsilon)$, which can be used in the same way as the joint distribution generated from Bayesian MCMC. In the case of a deterministic ODE model, choosing ρ to be the squared distance and performing an ABC analysis with decreasing ε is equivalent to the maximum-likelihood method used in Bayesian MCMC with a least-squares cost function [67].

11.2.3 Example: Robertson Model

To illustrate the concepts discussed in this section and throughout the chapter, we follow Eydgahi et al. [24] and consider a set of mass-action reactions among chemical species A, B, and C originally formulated by Robertson [58]:



These reactions yield the nonlinear system of ODEs:

$$\frac{d[A]}{dt} = k_3 \cdot [B] \cdot [C] - k_1 \cdot [A] , \quad (11.13)$$

$$\frac{d[B]}{dt} = k_1 \cdot [A] - k_2 \cdot [B]^2 - k_3 \cdot [B] \cdot [C] , \quad (11.14)$$

$$\frac{d[C]}{dt} = k_2 \cdot [B]^2 . \quad (11.15)$$

As initial conditions, we took

$$[A_0] = 1 , \quad [B_0] = [C_0] = 0 . \quad (11.16)$$

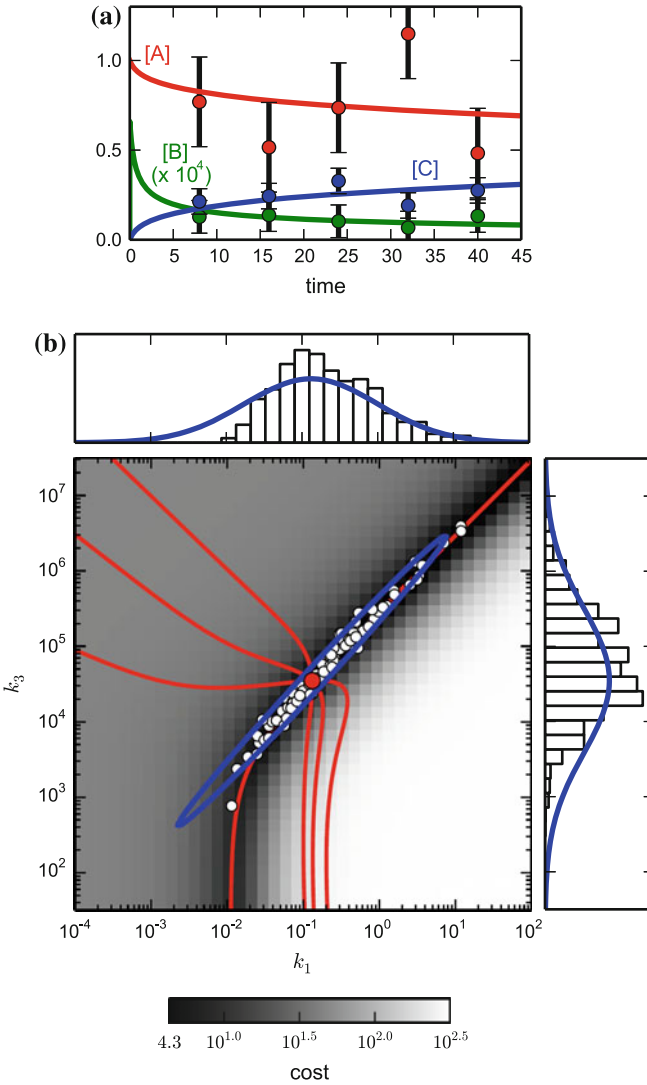


Fig. 11.2 Data fitting for the Robertson model. **a** Simulated data and best-fit trajectories. *Error bars* correspond to one standard deviation. **b** Corresponding cost landscape showing best-fit parameters (*red point*), and the confidence interval from the $J^T J$ approximation to the Hessian (*blue ellipse*). One hundred samples from a Bayesian MCMC ensemble (*white dots*), and geodesic curves starting at the best-fit (*red lines*) are also shown. *Top and right panels* show marginal distributions of k_1 and k_3 , respectively, inferred from the Hessian approximation (*blue curve*) and the Bayesian ensemble (*white histogram*)

We generated synthetic data points for [A], [B], and [C] by sampling every 8 time units from the model with initial parameters

$$k_1 = 0.04, \quad k_2 = 3 \times 10^7, \quad k_3 = 10,000, \quad (11.17)$$

and adding normally-distributed noise to each data point with standard deviation equal to 25% of the maximum value of the corresponding variable. We then fixed k_2 and used least-squares optimization to fit the synthetic data, estimating k_1 and k_3 (Fig. 11.2a).

We conducted sensitivity analysis using both the Hessian matrix (local) and an ensemble of parameter sets sampled by Bayesian MCMC (global), using Sloppy-Cell [55]. In both cases, we added log-normal priors that restricted k_1 and k_3 to remain within three orders of magnitude of the initial values (Eq. 11.17), with 95% confidence. The quadratic approximation (Fig. 11.2b, blue ellipse) mimics the shape of the cost landscape quite well in the vicinity of the best fit, but it overestimates the variability in these two parameters, due to the strong nonlinearity in their relationship. The ensemble (Fig. 11.2b, white dots), on the other hand, captures the true posterior distribution of parameters.

This example illustrates some of the difficulties encountered when using parameter scanning methods for sensitivity analysis. A simple scan along each parameter axis at the best fit value will dramatically underestimate the variability in the parameter estimates, and the nonlinearities in the landscape will render the correlation analysis used in LHS inaccurate.

11.3 Sloppiness

The topography of the cost landscape plays a critical role in modeling. For example, the cost landscape of the Robertson model (Fig. 11.2b) is highly anisotropic, as indicated by the eigenvalue spectrum in Fig. 11.3a(i). Near the best-fit, the parameter combination k_1/k_3 is tightly constrained (corresponding to the large eigenvalue). By contrast, the parameter combination $k_1 \times k_3$ is loosely constrained (corresponding to the small eigenvalue), so inferred values of k_1 and k_3 have large uncertainty. In 2003, Brown and Sethna noted similar behavior in a much more complex signaling model [13, 14], leading to the discovery of sloppiness.

Brown and Sethna used a system of 15 nonlinear differential equations, involving 48 rate constant parameters, to model the activation of ERK1/2 by epidermal growth factor (EGF) and neuronal growth factor (NGF) in PC12 cells [14]. They fit their model to 68 data points tracking the time-course of activation of several proteins in the system. When they analyzed the corresponding Hessian matrix, they found a surprising regularity in the eigenvalue spectrum (Fig. 11.3a(iv)). The eigenvalues spanned many orders of magnitude roughly evenly, a phenomenon they deemed *sloppiness*. The large eigenvalues and corresponding eigenvectors defined *stiff* combinations of parameters that were tightly constrained, whereas the small eigenvalues

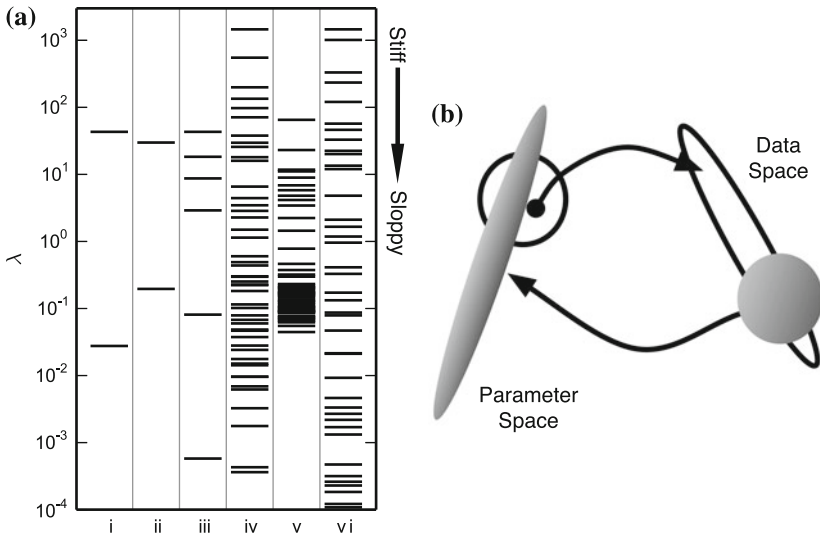


Fig. 11.3 Sloppy parameter spaces, eigenvalues, and data space. **a** Eigenvalue spectra for several sloppy models, illustrating different approaches to parameter and data space. *i* $J^T J$ eigenvalues of the Robertson model and data described in Sect. 11.2.3. *ii* Principle components analysis (PCA) eigenvalues of a Bayesian parameter ensemble for the Robertson model (inverted for comparison with column *i*). *iii* PCA eigenvalues of the model manifold for the Robertson model (inverted and rescaled so the largest eigenvalue matches the largest eigenvalue in column *i*). *iv* $J^T J$ eigenvalues of the Brown and Sethna model for differentiation in PC12 cells [14] that was fit to 68 data points. *v* PCA eigenvalues of a Bayesian parameter ensemble for the PC12 model [32], generated with log-normal priors similar to Robertson model (inverted for comparison with column *iv*). *vi* $J^T J$ eigenvalues for the PC12 model fit to continuous data on all molecular species in the model [32] (rescaled so the largest eigenvalue matches the largest eigenvalue in column *iv*). **b** The sloppy mapping between parameter and data space implies that spherical regions of parameter space map to distorted sloppy regions in model space, and vice-versa [18]

and corresponding eigenvectors defined *sloppy* combinations that were loosely constrained.¹ Moreover, a similar pattern of eigenvalues was found even when considering large amounts of perfectly-fit synthetic data on every species in the model (Fig. 11.3a(vi)), suggesting that sloppiness was a property of the model itself, not the particular data set. As illustrated in Fig. 11.3b, this sloppiness implies that large volumes of parameter space can map to a small volumes in data space, and vice-versa.

The importance of sloppiness to systems biology became more apparent in 2007, when Gutenkunst et al. found sloppiness in a diverse set of sixteen other systems biology models [32]. In a systematic survey of the BioModels database [48], Erguler and Stumpf later found sloppiness in 180 systems biology models [23].

¹Concurrent with Brown and Sethna's work, Rand et al. independently noted an exponential spacing of eigenvalues for several circadian clock models, although Rand et al. focused their analysis on the stiffest few eigenvalues [57].

Although similar in spirit, sloppiness differs from conventional conceptions of robustness [46]. Typically, when a biological system is deemed *robust* it means that a particular qualitative behavior is insensitive to a particular perturbation. That perturbation may be a change in parameter values [78], temperature [56], or structure of the system. Sloppiness, on the other hand, focuses on the quantitative behavior of the model and its sensitivity to changes in combinations of parameters. A system may be sloppy, but not be robust to changes in individual parameters. For example, at the best-fit set of parameters, the Robertson model is robust to changes in k_1 or k_3 that leave the stiff parameter combination k_1/k_3 unchanged, but it is fragile to changes in either parameter individually.

Brown and Sethna's discovery of sloppiness spawned a large body of literature exploring its theoretical basis. Early work on the origins of sloppiness focused on symmetries between parameter effects [76], but recent connections with information geometry and interpolation theory have revealed a more general origin (Sect. 11.4.1). Although it has been best-studied in the context of systems biology, sloppiness also appears in non-biological models [30, 76], including classic statistical problems such as fitting a sum of exponentials or polynomials to data [76]. In classic physics models for magnetism and diffusion, sloppiness emerges when observations are restricted to large length scales, so microscopic details of the system cease to matter [49]. A similar phenomenon may be occurring in systems biology, where most experiments probe the collective behavior of many interacting reactions. The ubiquity of sloppiness also suggests that it may have implications for biological evolution [18].

In the remainder of this chapter, we focus on the practical implications of sloppiness for modeling biological systems, through building predictive models, designing experiments, and developing numerical methods.

11.3.1 Local and Global Perspectives

We have defined sloppiness in terms of the distribution of the eigenvalues of the Hessian matrix. For nonlinear models, however, the Hessian depends on where in parameter space it is evaluated, as exemplified by the curved basin of the Robertson model (Fig. 11.2b). In the Robertson model and in the Brown and Sethna PC12 model [14], Hessian matrices calculated using multiple parameter sets from the MCMC posterior distribution are all sloppy, with similar eigenvalue spectra but differing eigenvectors [30]. This suggests that the curved basins are everywhere locally sloppy, but a more global perspective can be obtained from Principal Component Analysis (PCA) of the MCMC parameter set ensemble.

PCA is the eigen-decomposition of the covariance matrix of a set of points in space (here we focus on points in parameter space), and it has a broad range of applications in statistics [1, 37]. (For more on PCA and other statistical models in systems biology, see Chap. 6 in this volume [63].) PCA is defined such that the first principal component is the eigenvector with the largest eigenvalue, and it points in the direction that accounts for the largest amount of variance in the positions of the

points. The eigenvector with the second-largest eigenvalue points in the direction that accounts for the second-largest amount of variance in positions of the points, orthogonal to the previous direction, and so on. The Hessian matrix and the covariance matrix share eigenvectors (Sect. 11.2.1), so we can think of performing PCA on an ensemble as the global analog to the local analysis of the Hessian matrix. The eigenvalues of the Hessian are inversely related to those of the principal components, and in Fig. 11.3a(ii, v) we take the inverse of the PCA eigenvalues for comparison.

Because the ensemble captures nonlinearities in the parameter space if they exist, we might wonder whether models with sloppy Hessian eigenspectra also have sloppy PCA eigenspectra. Figure 11.3a(i, ii) shows the eigenvalues for the Robertson example computed using the Hessian and by PCA, respectively. The similar spacing of eigenvalues shows that the aspect ratio of the level curves of the cost manifold are preserved, even as nonlinearities cause them to curve. Figure 11.3a(iv, v) show corresponding Hessian and PCA spectra for the PC12 model [13], where the ensemble in Fig. 11.3a(v) was generated with priors similar to those we used for our Robertson model. The spectrum is truncated from below by the prior,² and the largest eigenvalues are reduced due to nonlinearities in the parameter space that are better captured by the ensemble. Although quantitative differences are evident in the eigenspectra generated by the two methods, qualitatively they are both sloppy, spanning several orders of magnitude with eigenvalues that are evenly spaced in the logarithm. In addition to these empirical comparisons of Hessian and PCA eigenspectra, recent work in information geometry (Sect. 11.4.1) also suggests that sloppiness is a global property.


11.3.2 Predictive Modeling from Sloppy Systems

Because sloppiness appears universal in systems biology models, attempting to fit individual parameters in such models is difficult and often uninformative. Even with extensive time-series data, inferred values for individual parameters are often imprecise, because the model is insensitive to changes to most parameter combinations. The common practice of reporting only the means and confidence intervals on individual parameters should thus be avoided. On the other hand, because the model is quite sensitive to changes in a few parameter combinations, with precise measurement of time-series data it is often possible to tightly constrain model predictions, despite large individual parameter uncertainties [32]. Moreover, in most cases precisely modeling and predicting system behavior is more compelling than precisely inferring individual parameters.

²A log-normal prior that bounds a parameter θ to be, with $\approx 95\%$ confidence, between θ_0/F and $\theta_0 \times F$ corresponds to an additional residual in the cost function (Eq. 11.1) of $r = (\log \theta - \log \theta_0) / \log \sqrt{F}$. Such a residual adds $1/(\log \sqrt{F})^2$ to the diagonal elements of the Hessian matrix, bounding the eigenvalues from below. In our case, $F = 10^3$, so the eigenvalues must be greater than ≈ 0.084 .

For the Robertson model (Sect. 11.2.3 and Fig. 11.2), the individual parameters k_1 and k_3 are only loosely constrained by time-course data, so our inferred values for these parameters, whether using the Hessian approximation or Bayesian MCMC sampling, span many orders of magnitude, as shown by the top and right panels in Fig. 11.2b. The cost landscape does show a distinct nonlinear canyon of parameter sets that fit the data well (Fig. 11.2b), but this canyon does not align with any individual parameter, so inferring individual parameters with high precision is difficult. On the other hand, combinations of parameters perpendicular to the canyon are tightly constrained.

Because predictions are often more important than individual parameter values, we tested the ability of our synthetic time-course data in the Robertson model to constrain a novel prediction. In particular, we added a new reaction to the model:



We then predicted the time course of $[C]$ in this four-reaction model by generating trajectories using the results of our data fit of the original three-reaction model (Eq. 11.12, Fig. 11.2). When we generated a set of predictions assuming that we knew k_1 and k_3 to high precision (95 % confidence interval of $\pm 50\%$), the prediction for $[C]$ was tightly constrained (Fig. 11.4a, b). If we instead knew k_1 precisely, but k_3 imprecisely, the prediction of $[C]$ was uninformative (Fig. 11.4c, d), because the corresponding parameter ensemble includes parameter sets with high cost, rather than exploring only the canyon. When we approximated the stiff and sloppy directions using the Hessian, as in (Fig. 11.4e, f) and generated predictions from this set of parameter combinations, we recovered some constraint on the prediction uncertainties.

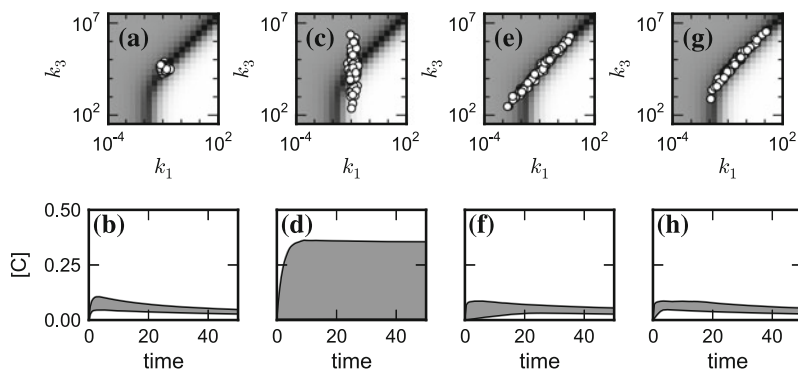


Fig. 11.4 Parameter inferences and prediction 95 % confidence intervals for the Robertson model (Sect. 11.2.3). **a, b** Assuming both parameters are measured to $\pm 50\%$ precision. **c, d** Assuming k_1 is measured to high precision, but k_3 must be guessed to low precision (95 % confidence interval spanning three orders of magnitude). **e, f** Evaluating the prediction using samples from the $J^T J$ approximation to the Hessian matrix. **g, h** Evaluating the prediction using samples from the Bayesian MCMC ensemble illustrated in Fig. 11.2b

Taking points from the MCMC-generated ensemble constrained uncertainty even further (Fig. 11.4g, h).

In the Robertson model, both the Hessian and Bayesian-ensemble approaches constrained prediction uncertainty much better than a mixture of well- and poorly-determined rate constants. However, the Hessian did not perform as well as Bayesian MCMC, because the cost manifold is nonlinear; in other words, the canyon of well-fitting parameter sets is curved. The Hessian approximates the stiff and sloppy directions at the best-fit parameters, but away from the best-fit parameters the stiff and sloppy directions inferred from Hessian deviate from the true shape of the cost manifold, so points outside the canyon are sampled. Bayesian MCMC avoids this problem, as there is no assumption of linearity of the cost manifold, so the sampling follows the curve of parameter sets that both fit the data well and yield accurate predictions. The Hessian approximation works well in the Robertson model, but it may fail in more complex models with stronger nonlinearities in the cost landscape [24, 31].

11.3.3 Experimental Design

Motivated by the previous example of precise predictions from a sloppy model without precise parameters, we turn to the design of informative experiments. Experimental design is a large sub-field of systems biology, and many methods have been developed for designing experiments and models to extract optimal information about a quantity of interest [3, 15, 34, 44, 47]. In this section, we discuss several studies that directly address sloppiness. One identifies additional time-series data points that improve system behavior prediction [15], and the others identify experimental conditions that improve parameter inference [3, 34, 68].

11.3.3.1 Optimal Design for Prediction

As illustrated by the Robertson model in Fig. 11.4c, d, precisely measuring individual parameters in a complex model may not improve the predictive power of the model. To overcome this difficulty, Casey et al. developed an approach for designing experiments to improve the prediction of unmeasurable quantities and applied it to a model of epidermal growth factor receptor (EGFR) activation [15].

In the EGFR network, Casey et al. were interested in predicting the dynamics of the triple complex of Cool-1, Cdc42 and Cbl, each of which potentially disrupts receptor down-regulation. The triple complex was not directly measurable, so they relied on a complex systems biology model to predict its dynamics. Casey et al. fit their model to existing experimental data to obtain best-fit parameter values θ^* and an ensemble of parameter sets that fit the data well, but they found that the predicted trajectory for the triple complex had large uncertainty.

Given the large prediction uncertainty from the existing data, Casey et al. set out to design a new experiment to minimize the variance of the prediction. Doing so

required searching over the space of possible experiments and evaluating prediction variance many times. Bayesian sampling (Sect. 11.2.2.2) is the preferred way to estimate prediction variance, but it is computationally very expensive, so Casey et al. used an approximation to the variance of their prediction p :

$$\text{Var}(p) \approx \left. \frac{\partial p}{\partial \theta} \right|_{\theta^*} (J^T J)^{-1} \left. \frac{\partial p}{\partial \theta} \right|_{\theta^*}. \quad (11.19)$$

Here $(J^T J)^{-1}$, the inverse Fisher Information Matrix (FIM), asymptotically approximates the covariance of the parameters, and $\partial p / \partial \theta$ is a linear approximation of the model response to changes in the parameters [15].

Casey et al. employed a sequential experimental design to minimize the prediction variance calculated via Eq. 11.19. They first searched over experimental conditions, measurable molecular species, and timepoints to find the single data point whose addition most greatly reduced the prediction variance. This was computationally feasible because adding a single data point to the collection of measurements is a rank-one update of the Fisher Information Matrix [15]. Assuming that single data point represented the optimal condition and species to measure, Casey et al. then optimized over possible combinations of measured timepoints to design a complete experiment.

Applying their computational analysis, Casey et al. carried out the experiment they had designed. Adding the new data points to their model, they built a new ensemble of parameter sets from which to make predictions. As desired, the new ensemble dramatically reduced uncertainty in the predicted dynamics of the triple complex. Uncertainties on individual parameter values, however, were not substantially smaller. The sloppiness of the EGFR model allowed Casey et al.'s experimental design to improve prediction precision without improving parameter precision, but experimental design can also improve parameter precision.

11.3.3.2 Optimal Design for Parameter Inference

Fitting time-course data typically poorly constrains individual parameter values in sloppy models, but careful experimental design can yield well-constrained parameters. In a recent manuscript, Tönsing et al. argue that sloppiness can be generically caused by autocorrelation and sparseness in the Jacobian matrix (Eq. 11.5) of parameter sensitivities for residuals between model and data [68]. Autocorrelation naturally arises in time-course measurements, and sparseness arises because different predictions may be sensitive to different parameters. Tönsing et al. further show using a model and *in silico* experiments from the DREAM 6 challenge [52] that careful experimental design can avoid autocorrelation and sparseness, minimizing sloppiness in the resulting parameter inferences.

In a more targeted study, Apgar et al. have shown [3] that carefully designed complementary experiments can in principle tightly constrain all parameter values

in the original Brown and Sethna sloppy model of EGF/NGF signaling in PC12 cells [13, 14]. Apgar et al. sought to design a set of experiments that would together yield uncertainties on all 48 model parameter values of less than 10%, based on the Hessian approximation. To do so, they considered 164,500 potential experimental conditions, encompassing various levels of EGF and NGF stimulation and protein overexpression or knockdown. To avoid computationally challenging re-optimization of parameter values, they assumed that each experiment would yield data that exactly matched the model prediction. Their design processes employed a greedy algorithm that, at each step, chose the experiment that constrained the most parameters to within 10% that were not constrained by any earlier experiment. Remarkably, they found that five carefully chosen experiments were enough to tightly constrain all parameters [3]. It was essential that the experiments be chosen in a complementary way; choosing random experiments or even the best individual experiments gave much poorer results. The computational experiments that Apgar et al. considered used continuously sampled species time courses, yielding effectively many more data points than typical experiments, which may account for much of the improvement in parameter constraint [16]. Such dense measurements are, however, becoming increasingly feasible, and even with fewer collected data points, tight constraint on all parameters in the model are still possible [33].

Recently, Hagen et al. have relaxed many of the simplifying assumptions made in Apgar et al.'s work [34]. Most importantly, they considered data at discrete points along the trajectory with some assumed experimental error instead of continuous measurements with zero error. As a result, they had to re-optimize parameters at each stage of the experimental design, so each experiment was chosen on the basis of a model with imperfect parameters. Nevertheless, they found that just six experiments were needed to constrain all parameters to within 10% as assessed by the Hessian approximation, confirming the previous results.

The experiments designed by Tönsing et al. [68], Apgar et al. [3], and Hagen et al. [34] are complex, and to date they have not been carried out in the lab. This work, however, demonstrates the power of experimental design and offers hope that parameter values can indeed be precisely inferred even for sloppy models.

11.4 Information Geometry Perspective

To this point, we have reviewed work whose focus was analyzing the properties of parameter space. However, we have seen it is often beneficial to focus on the model predictions rather than the parameter values. Recent results focusing on data space rather than parameter space have proven beneficial for understanding the properties of models and for advancing numerical techniques for exploring them. This approach, usually known as information geometry since it combines information theory with differential geometry, is a natural mathematical language for exploring parameterized models. As we have seen, in essence a model is a mathematical mapping from parameters to predictions. This recognition leads to the interpretation of a model as

a manifold embedded in the space of data. The approach is very general, applicable to any parameterized statistical model (although we focus on least squares models in this review) and has many deep connections to statistics [2, 4, 7, 41, 54]. Because differential geometry is foreign to most biologists, much of the technical aspects and insights of information geometry are not immediately accessible to much of the systems biology community. In this section, we give a summary of recent results without assuming a prior understanding of differential geometry, illustrating the types of analyses that can be performed and providing references for further study.

To illustrate the approach for least-squares models, we return to the Robertson model introduced in Sect. 11.2.3. This model has two parameters and was fit to $M = 15$ data points. Any experimental realization of the data can be interpreted as a single point in \mathbb{R}^M . Likewise, for any value of the parameters, the model predictions are similarly a point in \mathbb{R}^M . As the two model parameters k_1 and k_3 are varied over their allowed ranges, the model sweeps out within the 15-dimensional data space a two-dimensional surface known as the model manifold and denoted by \mathcal{M} . In general, for a model of N parameters fit to M data points, the model manifold is the N -dimensional surface embedded in \mathbb{R}^M constructed by varying the model parameters over their physically allowed values.

Figure 11.5d shows a three-dimensional projection of the high-dimensional data space for the Robertson model. The manifold was calculated using a grid of parameter values over the ranges shown in Fig. 11.2b. Model trajectories were sampled at the equally-spaced timepoints for which data was simulated in Fig. 11.2a. The axes of the visualization come from a principal component analysis (Fig. 11.5a–c) that was performed for the resulting set of model trajectories.

The Robertson model manifold illustrates several features of the information geometry perspective that make it a powerful tool for studying models that complements the approach of considering the cost surface in parameter space. First, there is no information loss in the model manifold, i.e. the manifold is mathematically equivalent to the model itself. In contrast the cost surface in parameter space condenses the M numbers making up the prediction and data vectors into a single number.

Second and relatedly, information geometry separates the model, i.e. the manifold embedded in data space, from the data to which it is being fit, i.e. a point in the data space (blue star in Fig. 11.5). This is a useful abstraction which allows one to study the properties of the model itself irrespective of what is experimentally observed. The cost surface in parameter space will vary with the observed data. The best-fit parameters correspond to the point on \mathcal{M} nearest to the data (red point in Fig. 11.5).

Third, the set of points that constitute the model manifold are the same regardless of how the model is parameterized. That is, the geometric properties of the model are the same if a systems biology model is expressed in terms of reaction rates or time constants, in bare or log-parameters. In fact, it will be unchanged if the model is reparameterized in a complicated, highly nonlinear way. Because of this, the geometric perspective places the emphasis on model predictions rather than the parameters. The parameters are not ignored completely, but act as coordinates on the manifold, i.e. labels for specific predictions. The grid lines in Fig. 11.5d–g correspond to a square grid in log-parameter space. In general, differential-geometric objects are

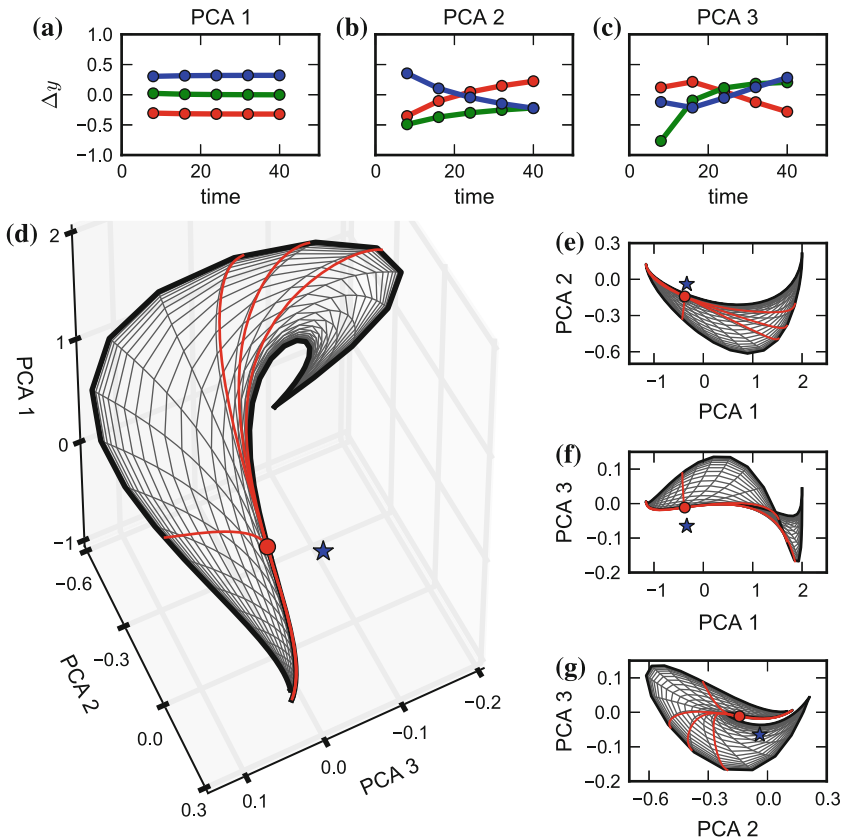


Fig. 11.5 Model manifold evaluated over the parameter space shown in Fig. 11.2b for the Robertson model (Sect. 11.2.3). **a–c** First three principle components of model prediction variation over the manifold. *Colors* identify species as in Fig. 11.2a. The first principle component, for example, represents an increase in [C] and a decrease in [A] that is roughly constant over the sampled timepoints. **d** Projection of the model manifold onto the first three principle components. *Blue star* shows the data, and *red dot* shows the best-fit trajectory. *Red lines* correspond to the geodesics in Fig. 11.2b. **e–g** Projections of the model manifold onto pairs of the first three principle components, as in (d)

constructed in terms of derivatives of the predictions with respect to the parameters in such a way that their relevant properties are the same for all possible parameterizations. Indeed, differential geometry is generally concerned with the properties of the manifold that are invariant under such reparameterizations.

A fourth and final point is that the language of differential geometry naturally accommodates the potentially large dimensionality of both the data space and the model manifold. Visualizations of both the cost surface and the model manifold are limited to only a few dimensions, but the geometric properties of high-dimensional spaces can be very different from those of the three-dimensional world in which

our visualizations live. For example, the specific properties of sloppiness are closely tied to the properties of high-dimensional manifolds. The mathematical formalism of differential geometry, however, has no such limitation and provides a framework within which the space can be systematically studied.

A particularly useful differential-geometric tool is a geodesic. A geodesic can be understood qualitatively as the generalization of straight lines to curved surfaces, i.e. the path connecting two points such that its image in data space is as close to a straight line as possible. A perfectly straight line is generally impossible, since the surface is typically curved for nonlinear models. The geodesic is constructed numerically as the solution of a nonlinear differential equation involving first and second derivatives of the model predictions with respect to the parameters. We refer the reader to any introductory text in differential geometry for more details [39, 65]. Several geodesic paths on the model manifold are shown in Fig. 11.5d–g, and the corresponding paths in parameter space are shown in Fig. 11.2b.

11.4.1 Models as Interpolation: Geometric Sloppiness

Geometry helps us understand the phenomenon of sloppiness. The observed universality of sloppiness across a wide range of models is perplexing; its ubiquity suggests some deep connecting principle [32, 76]. However, two other observations suggest otherwise. First, the hierarchy of Hessian eigenvalues can be transformed into any set of positive values by reparameterizing the model. Although such parameterizations might be unnatural from a human perspective, they are mathematically acceptable. Perhaps sloppiness is a reflection of how we humans choose to parameterize models. Is the human-preferred parameterization somehow perverse from a mathematical perspective? Second, sloppiness can be reduced by an appropriate choice of experiments [3, 34, 73]. Perhaps sloppiness is furthermore a reflection on what we choose to measure and not intrinsic to the system itself. Geometric arguments reconcile these apparently contradictory observations.

The key observation is that model manifolds are typically bounded, as is our example manifold in Fig. 11.5d. Considering the cost landscape in Fig. 11.2b, notice that away from the best fit, the cost surface plateaus, apparently approaching a limiting value. In fact, the parameters can be taken to zero and infinity without the cost becoming infinite, implying that the model manifold must be bounded. For any specific model this can be checked numerically using differential geometry. By numerically constructing geodesics, the manifold can be systematically explored to identify boundaries in any direction. Furthermore, by calculating the length of these geodesic paths in data space, one can measure the extent of the manifold in any given direction and calculate its aspect ratio. In this way it was found that typical sloppy model manifolds are not only bounded, but exhibit a hierarchy of widths analogous to the hierarchy of Hessian eigenvalues [71, 72]. For example, PCA analysis of the data space points used to construct the model manifold for our Robertson

model (Fig. 11.5d) reveals an exponential hierarchy of eigenvalues (Fig. 11.3b(iii)).³ This result is similar to the previously noted observation that sloppiness is generally reflected as a global property as measured by PCA analysis of a Bayesian ensemble (see Fig. 11.3b(ii, v)). However, in the current context the anisotropy is a reflection of an intrinsic property of the entire range of model predictions, rather than the ensemble of parameter values consistent with an instance of data.

This empirical observation of a hierarchy of widths in data space can be explained by applying interpolation theory [71]. Orthogonal geodesic paths identify cross sections of the model manifold. Geometrically, the cross section is formed by fixing the model output along a few axes (the directions orthogonal to the geodesic) and varying the output along others. Now consider a time series of model predictions for which a handful of time points have been fixed and the intermediate time points are allowed to vary. Although not fixed themselves, the values of intermediate time points can often be approximated by interpolating the values of those that are fixed. Therefore, the corresponding cross sections of the model manifold must be bounded by the accuracy of the interpolations.

For one-dimensional time series, the above argument can be made formal using theorems from interpolation theory [71]. In this case, cross sections become more narrow by roughly a constant factor for each additional fixed output, provided that the number of effective degrees of freedom probed by the model predictions is much less than the number of parameters. Qualitatively, this is understood to mean that the “complexity” of the data to be explained is much less than that of the model. In a sense, the model is over parameterized. However, it is often unclear how to remove the unnecessary parameters, because the stiff and sloppy parameter directions are almost always combinations of the bare parameters [32].

This argument suggests that models can be understood as generalized interpolation schemes and explains a number of observations. First, it explains why accurate predictions can be made by sloppy models when parameters are largely unconstrained; the predictions are interpolating from the existing data. It also explains why sloppiness disappears when complementary experiments are chosen; the number of effective degrees of freedom probed by the model becomes comparable to the number of parameters. In this case, the model needs all of the parameters to explain the data, resulting in tight bounds on their estimates.

The connection between manifold widths and the Hessian eigenvalues can be understood by dimensional analysis. The square roots of Hessian eigenvalues have units of data space distance per parameter space distance. If the parameters are expressed in the natural units of the problem, for example by using log-parameters, then we expect the eigenvalues to reflect the natural length scales of the manifold, i.e. the manifold widths as observed. What is interesting about this argument is the implication that the “natural” parameterization preferred by the human modeler is

³Note that this PCA was done on points sampled uniformly in parameter space, not data space, and this non-uniformity may bias the resulting eigenvalue summary of the model manifold. We expect, however, that this approach provides a good first approximation to the hierarchy of manifold widths that would be found by geodesics.

actually not perverse after all. Indeed, encoded in this natural parameterization is useful information: the length of the model manifold along several principal axes.

11.4.2 Applications to Numerical Methods

One of the most useful applications of the differential geometry approach involves the development and improvement of numerical methods for exploring parameter space. We now discuss two such improvements: the geodesic acceleration correction to the Levenberg-Marquardt algorithm for least squares data fitting [71, 72] and the Riemannian manifold sampling methods for Markov-Chain Monte Carlo (MCMC) [28].

11.4.2.1 Data Fitting

Fitting multi-parameter models to data via least squares can be notoriously difficult. One reason for this is that as algorithms approach the best fit they become agonizingly slow. This is because the cost surface in the vicinity of the best fit consists of a long, narrow canyon, as illustrated for the Robertson model in Fig. 11.2b. The algorithm must navigate this canyon en route to the best fit. The greater the aspect ratio of the canyon, the smaller the steps the algorithm must take. For many sloppy problems, as we have seen, it is not unusual for the canyon to have aspect ratios of 1000:1 or more, leading to very slow convergence rates.

A second reason that data fitting is difficult is that it is hard for the algorithm to even find the canyon to begin with. Observe in Fig. 11.2b how the cost surface plateaus away from the canyon. Because the cost surface is so flat, it is difficult for the algorithm to know in which direction to move. One typically finds that the results of a fitting algorithm are inconsistent, “converging” to wildly different parameter values depending on the starting point. This is typically attributed to multi-modality, or a rough cost surface with many local minima [25, 53, 59]. Closer inspection and the understanding of bounded model manifolds refines this picture in a very useful way. Specifically, the parameter values that result from failed runs of search algorithms typically contain parameters approaching their physical limits, e.g. zero or infinity. Geometrically, these points correspond to boundaries of the model manifold. The failure of search algorithms is due to them getting stuck in the boundaries of the model manifold en route to the best fit, i.e. being lost on the plateau. It was found that by adding weak, regularizing “prior” terms to the cost function that kept the algorithm away from the limiting parameter value, algorithms were much more successful at finding best fits [72]. These terms should be chosen in a way to force the algorithm to search in the region of parameter space to which the model behavior remains sensitive to changes in the parameters.

A second geometrically-inspired improvement to data fitting is an improvement to the common Levenberg-Marquardt algorithm known as the geodesic acceleration

correction [71, 72, 75]. The motivation for this algorithm is the observation that model manifolds typically have surprisingly small curvatures. (This observation had been noted by statisticians for several decades [5–8, 10, 36] and was finally explained by the same interpolation arguments that explain why the model manifolds are often bounded [71, 72].) Since the manifolds are relatively flat, the ideal path for an algorithm to follow is a geodesic, i.e. a straight line through data space. Notice in Fig. 11.2b how the geodesic path naturally follows the curvature of the canyon in parameter space. In the limit of small curvature, the second-order correction to the Levenberg-Marquardt algorithm reduces to the second-order term in the geodesic equation, which can be easily approximated with little computational cost compared to other aspects of the algorithm. The result is an algorithm that is dramatically faster at finding best fits. An open source FORTRAN implementation of this algorithm is available for download [69].

11.4.2.2 Bayesian Posterior Sampling

As discussed in Sect. 11.2.2.2, MCMC is a powerful technique for exploring parameter space and sampling the Bayesian posterior distribution. One of the challenges to effectively implementing this approach is the need to run the algorithm long enough to gather independent samples of the posterior. For a cost landscape with long narrow canyons, the Markov chain needs to effectively diffuse along the length of the long axes of the canyon for each sample. For the same reason that data fitting algorithms become sluggish in the canyon, the MCMC method also becomes very slow, requiring a very long chain before independent samples can be identified.

In order to alleviate this problem, it was suggested by Girolami and Calderhead that convergence could be improved by taking steps uniform in data space rather than parameter space [28]. Effectively, at each step of the chain, random parameter steps are proposed as a multivariate normal distribution with covariance chosen so that the corresponding steps in data space have covariance given by the identity. In this way, steps are preferentially aligned with the axis of the canyon, reducing the number of iterations necessary to generate independent samples. For extremely sloppy models with large aspect ratios in the canyons around their best fits, the improvement in convergence rate can be dramatic.

11.4.2.3 Curvature and Beyond

There are many other instances where differential geometry has provided insights and advancements in modeling and numerical methods. One of the most important concepts in differential geometry, and one that is beyond the scope of this review, is curvature. Measures of curvature have been used to quantify nonlinearity in models [5, 10, 36], measure kurtosis [35], and identify the global minimum in least-squares data fitting problems [21].

Information geometry has also led to a new approach to model reduction known as the manifold boundary approximation method [74]. By numerically constructing geodesics to the edge of the model manifold, limiting approximations are identified in the model that can be used to remove sloppy parameter combinations. The net result is a sequence of effective models of decreasing complexity. These reduced models remain expressed in terms of the microscopic parameters, i.e., there are no black boxes, and dramatically highlight the emergent control mechanisms that govern the system's behavior. Differential geometry also provides insights into questions of parameter identifiability, which combined with model reduction techniques, can be powerful tools for constructing appropriate mathematical representations of biological systems [70].

Although a relatively undeveloped approach, information geometry has provided a wealth of insight into modeling and the numerical methods for exploring model behavior. Much of the strength of the approach lies in its generality. Indeed, very little of what is summarized in this section is specific to systems biology. In this respect, applying information geometry to systems biology, with its wide array of models, is a compelling synthesis for the development of new theoretical and computational methods that are likely to not only advance biological understanding, but also find application in other complex systems.

11.5 Conclusion

Mechanistic models in systems biology typically possess a profusion of parameters, and this poses great challenge for modelers. In particular, understanding the multivariate sensitivity of the model to changes in parameter values is critical. Local and global analyses of sensitivity complement each other, and Bayesian methods are particularly powerful for assessing statistical confidence in parameter inferences and model predictions.

Analysis of many models in systems biology and other fields has revealed that nonlinear least-squares models are typically sloppy. Sloppy models have parameter sensitivity eigenvalues that span many decades roughly evenly and thus have highly anisotropic mappings between parameter and data spaces. Consequently, it is difficult to infer precise parameter values from data fits, but some predictions can nevertheless be tightly constrained. Careful experimental design can improve the precision of parameter inferences or model predictions, depending on the goals of the modeler. Information geometry offers a useful parameterization-independent perspective on modeling, and combining it with interpolation theory suggests that sloppiness arises because even complex models are often acting as interpolating functions between available data points. The information geometry perspective also suggests improved algorithms for optimization and Markov-chain Monte Carlo that account for the anisotropic and curved model and parameter spaces common in sloppy models.

Modelers have tackled a huge number of complex nonlinear systems in biology and other fields, and each model is unique. The study of sloppiness has shown,

however, that models of very different systems are nevertheless governed by shared deep statistical properties. Study of sloppy models thus offers insight and tools for not only systems biology, but also many other fields of science.

For readers who want hands-on experience with the methods and ideas discussed here, code implementing our analyses of the Robertson model is bundled with the SloppyCell software [55], available at: <http://sloppycell.sourceforge.net>.

Acknowledgments B.M. was supported by an ARCS Foundation Fellowship. A.R. was supported by NSF IGERT grant DGE-0654435. R.G. was supported by NSF grant DEB-1146074. We thank Alec Coffman for helpful discussions. R.G. and M.T. particularly thank Jim Sethna for his outstanding support and mentorship.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

1. Abdi, H., Williams, L.J.: Principal component analysis. *Wiley Interdiscip Rev: Comput. Statist.* **2**(4), 433–459 (2010)
2. Amari, S.I., Nagaoka, H.: *Methods of Information Geometry*, Translations of Mathematical Monographs, vol. 191. American Mathematical Society, New York (2000)
3. Apgar, J.F., Witmer, D.K., White, F.M., Tidor, B.: Sloppy models, parameter uncertainty, and the role of experimental design. *Mol. Biosyst.* **6**(10), 1890–1900 (2010)
4. Barndorff-Nielsen, O., Cox, D., Reid, N.: The role of differential geometry in statistical theory. *Int. Stat. Rev.* **54**(1), 83–96 (1986). doi:[10.2307/1403260](https://doi.org/10.2307/1403260)
5. Bates, D.M., Watts, D.G.: Relative curvature measures of nonlinearity. *J. Roy. Stat. Soc. B* **42**, 1–25 (1980)
6. Bates, D.M., Watts, D.G.: Parameter transformations for improved approximate confidence regions in nonlinear least squares. *Ann. Stat.* **9**(6), 1152–1167 (1981)
7. Bates, D.M., Watts, D.G.: *Nonlinear Regression Analysis and Its Applications*, Wiley Series in Probability and Statistics, vol. 32. Wiley, New York (1988)
8. Bates, D.M., Hamilton, D.C., Watts, D.G.: Calculation of intrinsic and parameter-effects curvatures for nonlinear regression models. *Commun. Stat. Simulat.* **12**(4), 469–477 (1983). doi:[10.1080/03610918308812333](https://doi.org/10.1080/03610918308812333)
9. Battogtokh, D., Asch, D., Case, M., Arnold, J., Schüttler, H.B.: An ensemble method for identifying regulatory circuits with special reference to the qa gene cluster of *Neurospora crassa*. *Proc. Natl. Acad. Sci. USA* **99**(26), 16904–16909 (2002). doi:[10.1073/pnas.262658899](https://doi.org/10.1073/pnas.262658899)
10. Beale, E.M.L.: Confidence regions in non-linear estimation. *J. Roy. Stat. Soc. B* **22**(1), 41–88 (1960)
11. Beaumont, M.A., Zhang, W., Balding, D.J.: Approximate Bayesian computation in population genetics. *Genetics* **162**(4), 2025–2035 (2002)
12. Birnbaum, A.: On the foundations of statistical inference. *J. Am. Stat. Assoc.* **57**(298), 269–306 (1962). doi:[10.2307/2281641](https://doi.org/10.2307/2281641)
13. Brown, K., Sethna, J.: Statistical mechanical approaches to models with many poorly known parameters. *Phys. Rev. E* **68**(2), 021904 (2003). doi:[10.1103/PhysRevE.68.021904](https://doi.org/10.1103/PhysRevE.68.021904)

14. Brown, K.S., Hill, C.C., Calero, G.A., Myers, C.R., Lee, K.H., Sethna, J.P., Cerione, R.A.: The statistical mechanics of complex signaling networks: nerve growth factor signaling. *Phys. Biol.* **1**(3–4), 184–195 (2004). doi:[10.1088/1478-3967/1/3/006](https://doi.org/10.1088/1478-3967/1/3/006)
15. Casey, F.P., Baird, D., Feng, Q., Gutenkunst, R.N., Waterfall, J.J., Myers, C.R., Brown, K.S., Cerione, R.A., Sethna, J.P.: Optimal experimental design in an epidermal growth factor receptor signalling and down-regulation model. *IET Syst. Biol.* **1**(3), 190–202 (2007). doi:[10.1049/iet-syb](https://doi.org/10.1049/iet-syb)
16. Chachra, R., Transtrum, M.K., Sethna, J.P.: Comment on Sloppy models, parameter uncertainty, and the role of experimental design. *Mol. Biosyst.* **7**(8), 2522; author reply 2523–4 (2011). doi:[10.1039/c1mb05046j](https://doi.org/10.1039/c1mb05046j)
17. Chib, S., Greenberg, E.: Understanding the Metropolis Hastings algorithm. *Am. Stat.* **49**(4), 327–335 (1995). doi:[10.1080/00031305.1995.10476177](https://doi.org/10.1080/00031305.1995.10476177)
18. Daniels, B.C., Chen, Y.J., Sethna, J.P., Gutenkunst, R.N., Myers, C.R.: Sloppiness, robustness, and evolvability in systems biology. *Curr. Opin. Biotech.* **19**(4), 389–395 (2008). doi:[10.1016/j.copbio.2008.06.008](https://doi.org/10.1016/j.copbio.2008.06.008)
19. De Smet, R., Marchal, K.: Advantages and limitations of current network inference methods. *Nat. Rev. Microbiol.* **8**(10), 717–729 (2010). doi:[10.1038/nrmicro2419](https://doi.org/10.1038/nrmicro2419)
20. Del Moral, P., Doucet, A., Jasra, A.: Sequential Monte Carlo samplers. *J. Roy. Stat. Soc. B. Met.* **68**(3), 411–436 (2006). doi:[10.1111/j.1467-9868.2006.00553.x](https://doi.org/10.1111/j.1467-9868.2006.00553.x)
21. Demidenko, E.: Criteria for global minimum of sum of squares in nonlinear regression. *Comput. Stat. Data An.* **51**(3), 1739–1753 (2006). doi:[10.1016/j.csda.2006.06.015](https://doi.org/10.1016/j.csda.2006.06.015)
22. Efron, B., Hinkley, D.V.: Assessing the accuracy of the maximum likelihood estimator: observed versus expected Fisher information. *Biometrika* **65**(3), 457–483 (1978). doi:[10.1093/biomet/65.3.457](https://doi.org/10.1093/biomet/65.3.457)
23. Erguler, K., Stumpf, M.P.H.: Practical limits for reverse engineering of dynamical systems: a statistical analysis of sensitivity and parameter inferability in systems biology models. *Mol. Biosyst.* **7**(5), 1593–1602 (2011). doi:[10.1039/c0mb00107d](https://doi.org/10.1039/c0mb00107d)
24. Eydgahi, H., Chen, W.W., Muhlich, J.L., Vitkup, D., Tsitsiklis, J.N., Sorger, P.K.: Properties of cell death models calibrated and compared using Bayesian approaches. *Mol. Syst. Biol.* **9**(644), 644 (2013). doi:[10.1038/msb.2012.69](https://doi.org/10.1038/msb.2012.69)
25. Fernández Slezak, D., Suárez, C., Cecchi, G.A., Marshall, G., Stolovitzky, G.: When the optimal is not the best: parameter estimation in complex biological models. *PLoS One* **5**(10), e13,283 (2010). doi:[10.1371/journal.pone.0013283](https://doi.org/10.1371/journal.pone.0013283)
26. Fisher, R.A.: On the mathematical foundations of theoretical statistics. *Philos. T. Roy. Soc. Lond.* **222**(594–604), 309–368 (1922). doi:[10.1098/rsta.1922.0009](https://doi.org/10.1098/rsta.1922.0009)
27. Flaherty, P., Radhakrishnan, M.L., Dinh, T., Rebres, R.A., Roach, T.I., Jordan, M.I., Arkin, A.P.: A dual receptor crosstalk model of G-protein-coupled signal transduction. *PLoS Comput. Biol.* **4**(9), e1000185 (2008). doi:[10.1371/journal.pcbi.1000185](https://doi.org/10.1371/journal.pcbi.1000185)
28. Girolami, M., Calderhead, B.: Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *J. Roy. Stat. Soc. B. Met.* **73**, 123–214 (2011). doi:[10.1111/j.1467-9868.2010.00765.x](https://doi.org/10.1111/j.1467-9868.2010.00765.x)
29. Gunawardena, J.: Models in systems biology: the parameter problem and the meaning of robustness. In: Lodhi H.M., Muggleton, S.H. (eds.) *Elements of Computational Systems Biology*, pp. 19–47. Wiley Hoboken (2010). doi:[10.1002/9780470556757.ch2](https://doi.org/10.1002/9780470556757.ch2)
30. Gutenkunst, R.: Sloppiness, Modeling, and Evolution in Biochemical Networks. Ph.D. thesis, Cornell University (2008). <http://www.ecommons.cornell.edu/handle/1813/8206>
31. Gutenkunst, R.N., Casey, F.P., Waterfall, J.J., Myers, C.R., Sethna, J.P.: Extracting falsifiable predictions from sloppy models. *Ann. NY Acad. Sci.* **1115**, 203–211 (2007a). doi:[10.1196/annals.1407.003](https://doi.org/10.1196/annals.1407.003)
32. Gutenkunst, R.N., Waterfall, J.J., Casey, F.P., Brown, K.S., Myers, C.R., Sethna, J.P.: Universally sloppy parameter sensitivities in systems biology models. *PLoS Comput. Biol.* **3**(10), e189 (2007b). doi:[10.1371/journal.pcbi.0030189](https://doi.org/10.1371/journal.pcbi.0030189)
33. Hagen, D.R., Apgar, J.F., White, F.M., Tidor, B.: Molecular BioSystems reply to comment on Sloppy models, parameter uncertainty, and the role of experimental design. *Interface Focus* pp. 2523–2524 (2011). doi:[10.1039/c1mb05200d](https://doi.org/10.1039/c1mb05200d)

34. Hagen, D.R., White, J.K., Tidor, B.: Convergence in parameters and predictions using computational experimental design. *Interface Focus* **3**(4), 20130,008–20130,008 (2013). doi:[10.1098/rsfs.2013.0008](https://doi.org/10.1098/rsfs.2013.0008)
35. Haines, L.M., O'Brien, T.E., Clarke, G.P.Y.: Kurtosis and curvature measures for nonlinear regression models. *Stat. Sinica* **14**(2), 547–570 (2004)
36. Hamilton, D.C., Watts, D.G., Bates, D.M.: Accounting for intrinsic nonlinearity in nonlinear regression parameter inference regions. *Ann. Stat.* **10**(38), 393 (1982)
37. Hotelling, H.: Analysis of a complex of statistical variables into principal components. *J. Educ. Psychol.* **24**(6), 417–441 (1933)
38. Hug, S., Schmidl, D., Li, W.B., Greiter, M.B., Theis, F.J.: Bayesian model selection methods and their application to biological ODE systems. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
39. Ivancevic, T.T.: *Applied Differential Geometry: a Modern introduction*. World Scientific, Singapore (2007)
40. Jaqaman, K., Danuser, G.: Linking data to models: data regression. *Nat. Rev. Mol. Cell Bio.* **7**(11), 813–819 (2006). doi:[10.1038/nrm2030](https://doi.org/10.1038/nrm2030)
41. Kass, R.E.: The geometry of asymptotic inference. *Stat. Sci.* **4**(3), 188–219 (1989)
42. Kirk, P., Silk, D., Stumpf, M.P.H.: Reverse engineering under uncertainty. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Cham (2016, this volume)
43. Kirk, P., Thorne, T., Stumpf, M.P.: Model selection in systems and synthetic biology. *Curr. Opin. Biotech.* **24**(4), 767–774 (2013). doi:[10.1016/j.copbio.2013.03.012](https://doi.org/10.1016/j.copbio.2013.03.012)
44. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by simulated annealing. *Science* **220**(4598), 671–680 (1983). doi:[10.1126/science.220.4598.671](https://doi.org/10.1126/science.220.4598.671)
45. Kitano, H.: Systems biology: a brief overview. *Science* **295**(5560), 1662–1664 (2002). doi:[10.1126/science.1069492](https://doi.org/10.1126/science.1069492)
46. Kitano, H.: Biological robustness. *Nat. Rev. Genet.* **5**(11), 826–837 (2004). doi:[10.1038/nrg1471](https://doi.org/10.1038/nrg1471)
47. Kreutz, C., Timmer, J.: Systems biology: experimental design. *FEBS J.* **276**(4), 923–942 (2009). doi:[10.1111/j.1742-4658.2008.06843.x](https://doi.org/10.1111/j.1742-4658.2008.06843.x)
48. Le Novère, N., Bornstein, B., Broicher, A., Courtot, M., Donizelli, M., Dharuri, H., Li, L., Sauro, H., Schilstra, M., Shapiro, B., Snoep, J.L., Hucka, M.: BioModels Database: a free, centralized database of curated, published, quantitative kinetic models of biochemical and cellular systems. *Nucleic Acids Res.* **34**(Database issue), D689–91 (2006). doi:[10.1093/nar/gkj092](https://doi.org/10.1093/nar/gkj092)
49. Machta, B., Chachra, R., Transtrum, M., Sethna, J.: Parameter space compression underlies emergent theories and predictive models. *Science* **342**(6158), 604–607 (2013). doi:[10.1126/science.1238723](https://doi.org/10.1126/science.1238723)
50. Marino, S., Hogue, I.B., Ray, C.J., Kirschner, D.E.: A methodology for performing global uncertainty and sensitivity analysis in systems biology. *J. Theor. Biol.* **254**(1), 178–196 (2008). doi:[10.1016/j.jtbi.2008.04.011](https://doi.org/10.1016/j.jtbi.2008.04.011)
51. Marjoram, P., Molitor, J., Plagnol, V., Tavaré, S.: Markov chain Monte Carlo without likelihoods. *Proc. Natl. Acad. Sci. USA* **100**(26), 15,324–8 (2003). doi:[10.1073/pnas.0306899100](https://doi.org/10.1073/pnas.0306899100)
52. Meyer, P., Cokelaer, T., Chandran, D., Kim, K.H., Loh, P.R., Tucker, G., Lipson, M., Berger, B., Kreutz, C., Raue, A., Steiert, B., Timmer, J., Bilal, E., Sauro, H.M., Stolovitzky, G., Saez-Rodriguez, J.: Network topology and parameter estimation: from experimental design methods to gene regulatory network kinetics using a community based approach. *BMC Syst. Biol.* **8**(1), 13 (2014). doi:[10.1186/1752-0509-8-13](https://doi.org/10.1186/1752-0509-8-13)
53. Moles, C.G., Mendes, P., Banga, J.R.: Parameter estimation in biochemical pathways: a comparison of global optimization methods. *Genome Res.* **13**(11), 2467–2474 (2003). doi:[10.1101/gr.1262503](https://doi.org/10.1101/gr.1262503)
54. Murray, M.K., Rice, J.W.: *Differential Geometry and Statistics, Monographs on statistics and applied probability*, vol. 48. Chapman & Hall, London (1993)
55. Myers, C.R., Gutenkunst, R.N., Sethna, J.P.: Python unleashed on systems biology. *Comput. Sci. Eng.* **9**(3), 34–37 (2007). doi:[10.1109/MCSE.2007.60](https://doi.org/10.1109/MCSE.2007.60)

56. Pittendrigh, C.: On temperature independence in the clock system controlling emergence time in *Drosophila*. *Proc. Natl. Acad. Sci. USA* **40**(10), 1018–1029 (1954)
57. Rand, D.A., Shulgin, B.V., Salazar, D., Millar, A.J.: Design principles underlying circadian clocks. *J. Roy. Soc. Interface* **1**(1), 119–130 (2004). doi:[10.1098/rsif.2004.0014](https://doi.org/10.1098/rsif.2004.0014)
58. Robertson, H.: The solution of a set of reaction rate equations. In: Walsh, J. (ed.) *Numerical Analysis, an Introduction*, pp. 178–182. Academ Press, London (1966)
59. Rodriguez-Fernandez, M., Mendes, P., Banga, J.R.: A hybrid approach for efficient and robust parameter estimation in biochemical pathways. *Biosyst.* **83**(2–3), 248–265 (2006). doi:[10.1016/j.biosystems.2005.06.016](https://doi.org/10.1016/j.biosystems.2005.06.016)
60. Cedersund, G., Samuelsson, O., Ball, G., Tegnér, J., Gomez-Cabrero, D.: Optimization in biology parameter estimation and the associated optimization problem. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
61. Savageau, M.A., Coelho, P.M.B.M., Fasani, R.A., Tolla, D.A., Salvador, A.: Phenotypes and tolerances in the design space of biochemical systems. *Proc. Natl. Acad. Sci. USA* **106**(16), 6435–6440 (2009). doi:[10.1073/pnas.0809869106](https://doi.org/10.1073/pnas.0809869106)
62. Seber, G.A.F., Wild, C.J.: *Nonlinear Regression*. Wiley, New York (1988)
63. Shah, M., Chitforoushzadeh, Z., Janes, K.A.: Statistical data analysis and modeling. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
64. Sisson, S.A., Fan, Y., Tanaka, M.M.: Sequential Monte Carlo without likelihoods. *Proc. Natl. Acad. Sci. USA* **104**(6), 1760–1765 (2007). doi:[10.1073/pnas.0607208104](https://doi.org/10.1073/pnas.0607208104)
65. Spivak, M.: *A Comprehensive Introduction to Differential Geometry*. Publish or Perish (1979)
66. Sunnåker, M., Stelling, J.: Model extension and model selection. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
67. Toni, T., Welch, D., Strelkowa, N., Ipsen, A., Stumpf, M.P.: Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *J. Roy. Soc. Interface* **6**(31), 187–202 (2009). doi:[10.1098/rsif.2008.0172](https://doi.org/10.1098/rsif.2008.0172)
68. Tönsing, C., Timmer, J., Kreutz, C.: Cause and cure of sloppiness in ordinary differential equation models (2014). [arXiv:1406.1734](https://arxiv.org/abs/1406.1734)
69. Transtrum, M.K.: Geodesic Levenberg-Marquardt source code (2012). <http://sourceforge.net/projects/geodesiclm/>
70. Transtrum, M.K., Hart, G., Qiu, P.: Information topology identifies emergent model classes. [arXiv:1409.6203](https://arxiv.org/abs/1409.6203) (2014)
71. Transtrum, M.K., Machta, B.B., Sethna, J.P.: Why are nonlinear fits to data so challenging? *Phys. Rev. Lett.* **104**(6), 060,201 (2010). doi:[10.1103/PhysRevLett.104.060201](https://doi.org/10.1103/PhysRevLett.104.060201)
72. Transtrum, M.K., Machta, B.B., Sethna, J.P.: Geometry of nonlinear least squares with applications to sloppy models and optimization. *Phys. Rev. E* **83**(3), 036,701 (2011). doi:[10.1103/PhysRevE.83.036701](https://doi.org/10.1103/PhysRevE.83.036701)
73. Transtrum, M.K., Qiu, P.: Optimal experiment selection for parameter estimation in biological differential equation models. *BMC Bioinf.* **13**, 181 (2012). doi:[10.1186/1471-2105-13-181](https://doi.org/10.1186/1471-2105-13-181)
74. Transtrum, M.K., Qiu, P.: Model reduction by manifold boundaries. *Phys. Rev. Lett.* **113**(9), 098,701 (2014). doi:[10.1103/PhysRevLett.113.098701](https://doi.org/10.1103/PhysRevLett.113.098701)
75. Transtrum, M.K., Sethna, J.P.: Improvements to the Levenberg-Marquardt algorithm for nonlinear least-squares minimization. [arXiv:1201.5885](https://arxiv.org/abs/1201.5885) (2012)
76. Waterfall, J., Casey, F., Gutenkunst, R., Brown, K., Myers, C., Brouwer, P., Elser, V., Sethna, J.: Sloppy-Model Universality Class and the Vandermonde Matrix. *Phys. Rev. Lett.* **97**(15) (2006). doi:[10.1103/PhysRevLett.97.150601](https://doi.org/10.1103/PhysRevLett.97.150601)
77. Van Schepdael, A., Carlier, A., Geris, L.: Sensitivity analysis in the design of experiments. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
78. von Dassow, G., Meir, E., Munro, E.M., Odell, G.M.: The segment polarity network is a robust developmental module. *Nature* **406**(6792), 188–92 (2000). doi:[10.1038/35018085](https://doi.org/10.1038/35018085)
79. Xu, T.R., Vyshechirsky, V., Gormand, A., von Kriegsheim, A., Girolami, M., Baillie, G.S., Ketley, D., Dunlop, A.J., Milligan, G., Houslay, M.D., Kolch, W.: Inferring signaling pathway topologies from multiple perturbation measurements of specific biochemical species. *Sci. Signal.* **3**(113), ra20 (2010). doi:[10.1126/scisignal.2000517](https://doi.org/10.1126/scisignal.2000517)

Chapter 12

Modeling and Model Simplification to Facilitate Biological Insights and Predictions

Olivia Eriksson and Jesper Tegnér

Abstract Mathematical dynamical models of intracellular signaling networks are continuously increasing in size and model complexity due in large part to the data explosion in biology. However, the sheer complexity of the relationship between state-variables through numerous parameters constitutes a significant barrier against obtaining insight into which parts of a model govern a certain read-out, and the uncertainty in model structure and especially model parameters is here a further complicating factor. To meet these two challenges of complexity and uncertainty, systematic construction of simplified models from complex models is a central area of investigation within systems biology as well as for personalized medicine. Model complexity makes the task of deriving predictions difficult in general and in particular when different read-outs depend on combinations of parameters, since exhaustive computer simulations are not sufficient for understanding nor feasible in practice. Construction of simplified models is therefore an important complementary approach to this end, while also facilitating the identifiability of over-parameterized models. Within this chapter we discuss different methods for model simplification, and we specifically summarize a recently developed simplification method based on an iterative “tearing, zooming and simplifying” approach. We also look into the modeling process in general. In the “tearing, zooming and simplifying” approach the original model is divided into modules (tearing), the modules are considered as input-output systems (zooming), which then are replaced by simplified transfer functions (simplifying). The idea behind the simplification is to utilize biological features such as modularity and robustness as well as abundance of typical dynamical behaviors in biology such as switch-like responses. The methodology is illustrated using a relatively complex model of the cell division cycle, where the resulting simplification corresponds to a piecewise linear system with delay, facilitating an understanding of

O. Eriksson (✉)

Department of Numerical Analysis and Computer Science, Stockholm University,
Stockholm, Sweden
e-mail: olivia@mech.kth.se

J. Tegnér

Unit of Computational Medicine, Department of Medicine, Center for Molecular Medicine,
Karolinska Institutet, Karolinska University Hospital, Stockholm, Sweden
e-mail: jesper.tegner@ki.se

© Springer International Publishing Switzerland 2016

L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_12

the underlying core dynamics and enabling the prediction of combinations of parameters that can change different model features like the size of the cell. Hence, the existence of biological organization principles enables a simplified description of intracellular dynamics.

Keywords Model simplification · Model reduction · Data integration · Dynamical models · Ordinary differential equations · Piecewise linear · Delay · Dynamical modules · Switch-like dynamics · Model decomposition · Lumping · Timescale separation · Sensitivity analysis · Identifiability · Tearing-zooming-simplifying

12.1 Introduction

The enormous increase in cellular experimental data of the last decades, resulting from the sequencing of the human genome [76, 81], microarray techniques [63], the FANTOM projects [43], 1000 genome project [74], ENCODE [75] and other high-throughput methods such as proteomics and metabolomics, has provided us with detailed lists on the constituents of the cell as well as their putative interactions. This has however not yet enabled an understanding of the functionality of the cell on the full genome-scale level. We have static descriptions on networks of interactions, in the same way as we have road maps, but still we do not know much about traffic flow [46]. In order to investigate and describe the dynamical behaviour of the huge system that a cell constitutes, we do not only need experimental developments but also improved computational and mathematical methods [9, 21].

The structure and dynamics of cellular networks are inherently complex, containing a large lattice of redundancy and intertwined feedback loops, where the interactions can be described by detailed biochemical reactions. Imagine the complexity of a single signaling pathway, such as the cell cycle, and then put together all possible pathways of the cell and the resulting picture becomes overwhelming. There are however simplifying circumstances in this complexity. Cells seem to have a modular organization in space and time, with a sparse number of interactions between constituents. Cells often display quite simple functions (e.g. on/off circuits) and have a robust behaviour. The key idea of this chapter is that provided that we could use these simple levels of regulation hidden in the mesh of details, then maybe we could get a step closer to retrieving descriptions of the functionality of whole cells.

In order to give precise unambiguous descriptions, and quantitative predictions, mathematically formalized relationships and parameters are needed [39, 77], and dynamical cellular network models, describing the time evolution of e.g. protein concentrations, are the focus of this chapter. A few examples of dynamical models describing different cellular pathways and phenomena include [12, 33, 50, 77, 82]. The size of these models is, on the full genome-scale, quite modest, describing only a small subset of the proteins or genes of the cell e.g. a signaling pathway or regulatory module. However, although being relatively small models they display large complexity, with intertwined feedback-loops and functional redundancy.

A future vision, which is becoming increasingly realistic [42], is to construct dynamical models of whole cells. One approach towards this goal is to combine small mechanistically detailed models into larger and larger contexts. Such combined models tend to have a large number of parameters and complicated rate laws. Within the emerging complexity of these models it can be a difficult task to figure out the functional relationships of the model constituents, e.g. which parts that are essential for a certain read-out. It is therefore challenging to derive predictions on important features without performing extensive *brute-force* computer simulations. This problem becomes even more urgent as most models of biochemical systems today are highly over-parameterized with respect to available *in vivo* data [65], i.e. there exists a very large number of parameter combinations that give an equally good agreement with the data.

One important tool to deal with larger systems is *model simplification or reduction*. A model simplification process can illuminate the dynamical mechanism behind certain behaviour of the system and identify functional relationships between variables that are not obvious from inspection of a large model. Furthermore, a simplified model could remove over-parameterization and thereby generate an efficient description of the system with improved identifiability.

In this chapter we review different methods for simplifying complicated models of intracellular signaling cascades. We will also discuss in detail an important biological signaling cascade within the cell, namely the network regulating the cell cycle. This well characterized biological system is a useful test case to develop a model simplification process based on the identification, characterization and simplification of dynamical modules, here resulting in a piecewise linear model with delay. Before we go into model simplification and reduction in general and this specific example in detail we will first discuss how intracellular models actually are constructed.

12.2 Data Integration, Experimental Setting and Model Uncertainty

To develop models of intracellular networks data from different sources have to be integrated. The direction of data integration is often described by a *top-down* or *bottom-up* approach [10]. The top-down approach considers data from the whole cell and by an iterative cycle between experiments and modeling retrieves a better and better model resolution of the system. In contrast the bottom-up approach considers detailed descriptions of subparts of the cell and combine these into larger and larger models. The first approach describes the system more phenomenologically while the other line of investigation is geared towards a more mechanistic objective.

The genome sequencing projects [76, 81] and other high-throughput experimental techniques are examples of a top-down approach. Together with computational methods within bioinformatics for combining and analyzing several sources of data these experiments have provided us with a detailed *parts-list* of genes, proteins and

other constituents within the cell. We will denote this type of data as *genome-scale* data. During the last decades we have witnessed how these parts have been combined to large *static networks* of possible interactions, e.g. protein–protein interaction [79], transcriptional regulation [49] and metabolic reaction [20] networks, retrieved from experiments or by computational methods. Finally the *activity* of the parts can also be recorded by e.g. microarrays [63].

In contrast, bottom–up approaches depend on characterizing the interactions between individual components of the system and then integrating them into a larger reaction network [10]. Detailed kinetic and physiochemical properties of the interactions are obtained from experiments, where the experimental setting differ from each other in their level of simplification as compared to native conditions. At one end there are test tube experiments with purified components in optimized conditions, whereas at the other end experiments on living cells. We will denote these settings the dissociated versus the embedded setting, respectively. An example of the dissociated setting is the estimation of the parameters K_M (the Michaelis Menten constant) and k_{cat} (catalytic rate) of an enzyme, using purified components (e.g. proteins) in a test tube in optimized conditions for their function [1]. Examples of the embedded setting include the monitoring of time-course data for protein interaction markers and/or morphological cellular features obtained from the living system under perturbation [62].

Data from several different dissociated characterizations has been combined, in a bottom–up approach, to try to reconstruct the dynamical behavior of the larger system. This has shown some success in the case of modeling of the glycolytic pathway of unicellular organisms like *Saccharomyces cerevisiae* [66]. The efforts to model this universally conserved pathway in this way began several years ago using each enzymes own optimal conditions for its characterization (and thereby different test tube conditions for different enzymes despite some of them sharing the same intracellular environment). Just recently standardized conditions that resemble the intracellular milieu has been used to characterize all enzymes in the pathway resulting in a more accurate model [1]. While a similar dissociated characterization, bottom–up, approach has been proposed for components of signaling cascades [41] there has not been any systematic effort in this direction. Unlike glycolytic enzymes in unicellular organisms, the components of intracellular signaling cascades have a higher degree of compartmentalization and more interacting partners so that even an approximate recreation of physiological conditions in vitro is far more demanding [32]. Thus, there are relatively few interactions of signaling cascades characterized with purified components in vitro. Besides the limited physiological relevance of the experimental conditions used in these characterizations, most of the resulting estimates correspond to steady state conditions and parameters, despite the fact that fast transient signaling is occurring in many systems. While dissociated characterization is appropriate for ranking the effects of mutations, as well as developing inhibitors with pharmaceutical applications, and studying mechanisms of enzyme action, its usefulness for modeling the dynamics of signaling pathways is limited. The parameter values estimated through dissociated characterization should be taken in most cases as soft constrains [12]. The embedded characterization is becoming prevalent

with the growing trend in the development of cell-based measurement methods. This development is driven by progress in protein labeling with genetically encoded markers which allows to track interactions and compartmentalization in real time [53], high-resolution and two photon microscopy which makes it possible to track discrete events in small cellular compartments with reduced photo damage [84] and high-content high-throughput techniques which enables monitoring the effect of several conditions on a few protein markers and second messengers in a single run [27]. Despite all this progress in data acquisition from cell-based experiments, the amount of available phenotypic data is still far from sufficient for fully constraining the model so that every parameter is identified. More information about experimental data for modeling can be found in [64].

12.3 Cellular Organization and Model Structure

There are different rationales driving modeling approaches. A *predictive* model is mainly used to make predictions about the system while a *descriptive* model stores knowledge. Within technical applications a predictive model often does not need to have any structural similarities with the real system as long as it can anticipate the behavior of the system, thus it is essentially viewed as a black-box with respect to mechanisms. However, within biology, models are often of a more descriptive nature, utilized to precisely collect knowledge about the system, and as a picture to discuss around, even though the final goal may well be predictions. Biological models are thus constructed to have a structural correspondence with the real system. Before we go into details on different model formalisms and simplification methods in the coming sections, we will here first discuss the structural and functional organization of cellular networks.

As described above, high-throughput genome-scale experiments have provided huge amounts of data concerning the cellular constituents, like genes and proteins, which can be organized into different static networks. The topology of such networks has been analyzed by graph theoretical methods [7] and intriguing similarities have been found between different biological systems. Numerous networks display a power-law distribution in the number of connections (edges) a node can have, i.e. most nodes have only a few edges to other nodes at the same time as there exist a few nodes having a large number of edges, denoted *hubs*. Many networks have also been found to have a relatively high clustering-coefficient, indicating the existence of groups of highly interconnected nodes, *topological modules*. Cellular functions have long been suggested to be carried out in a highly modular manner [7, 39, 78], where modular refers to a group of molecules, physically and/or functionally linked together, which perform a distinct function. Most studies so far concern static topological modules. Identifying topological modules is however a non-trivial task. The fact that clustering and hubs coexist indicates that topological modules are not independent and well isolated from each other, but rather that the network has a hierarchical organization [7]. Since the cell is in fact a dynamical system, *dynamical*

modules could have a closer correspondence to the functional or biological properties of the system. Tyson et al [78] have reviewed a set of different dynamical circuits typical for cellular networks.

There is one specific type of circuit behavior we would like to consider in a bit more in detail here, namely, *switch-like dynamics*. This is a recurrent phenomenon in cellular networks and we will illustrate how to exploit this feature in the ‘tearing-zooming-simplifying’ example provided at the end of this chapter, effectively defining a simplified model. Switch-like dynamics can be found in many biological systems [6, 36, 38, 40], and is often modeled using a steep sigmoid function e.g. a Hill-function with a high Hill coefficient. Biologically, this can be due to cooperative processes, positive feedback, or enzymes operating near saturation [28]. Two explanations have been suggested by James Ferrell as to why a steep sigmoid input/output response should be useful for the cell [28]. The resting state of a cell could be near the upstroke of the input-output curve, or it could be far away. In the first case a small change in input can give a large change in output, an amplification of the signal. In the second case the system would filter out small stimuli and allow the cell to respond decisively to stimuli of a sufficient magnitude. The second case could therefore support cellular robustness against noise.

12.4 Modeling Formalisms and Choice of Model Detail

An important part in the art of modeling is the decision of modeling formalism. This choice is strongly influenced by the nature of the question and the data at hand, and therefore the art in choosing a suitable abstraction or representation of the system at hand. In order to give precise unambiguous descriptions, and quantitative predictions, mathematically formalized relationships and parameters are often needed [39, 77]. In this section we will mainly focus on dynamical models, but in the interest of completeness we will first briefly discuss static models.

12.4.1 Static Models

As described earlier, high-throughput techniques and computational methods have produced an incredible quantity of biological interaction data. Data that is conveniently represented by large networks, or *graphs*. Such graphs consist of *nodes*, representing e.g. genes, proteins or metabolites, connected by *edges*, representing interactions or other relationships. These graphs can be *directed* as in the case of gene-regulatory networks, where the product of one gene is regulating another gene or *undirected* as is the case for protein–protein networks, when only a binding possibility is recorded. As mentioned earlier, the topology of different biological networks have been analyzed by graph theoretical methods and similarities between different types of networks have been discovered. It has, however, become increasingly clear

that these static networks only describe interaction possibilities, and that not all of the edges are active at a certain time and in a certain context (cellular location, external signal) [35, 54]. The data of possible interactions therefore need to be combined with node activity in order to approach a dynamical description of the cellular network.

12.4.2 Dynamical Models

Depending on the purpose of the modeling, the approximations that can be made on the system, and the experimental data available, there are several different mathematical formalisms to choose from when it comes to *dynamical* models of cellular networks [13, 16]. Examples include *Boolean models*, *non-linear ordinary differential equations*, *piecewise linear differential equations*, *partial differential equations*, *delay differential equations*, *stochastic master equations* and *rule-based formalism*. We here describe ordinary differential equation (ODE) and piecewise linear (PL) differential equation models in more detail, and also touch upon delayed differential equation (DDE) models. We look at ODE models because this is the model formalism most widely used within biochemical modeling, and DDE and PL differential equation models since the simplified model in the example at the end of this chapter is a PL model with delay. More information about modeling, and then especially modeling under uncertainty, can be found in [45].

Non-linear ordinary differential equations Modeling cellular networks by ordinary differential equations (ODE:s) uses non-negative, time-dependent variables to describe e.g. the concentrations of proteins or other molecules. Interactions between molecules correspond to functional or differential relations between the variables. The concentrations are thus described by *rate equations*, describing the rate of production of a component of the system as a function of this and other components of the system. We use \dot{x} to denote time-derivative. The rate equations correspond to

$$\dot{x}_i = f_i(\mathbf{x}, \mathbf{u}), \quad 1 \leq i \leq n, \quad (12.1)$$

where $\mathbf{x} = (x_1, \dots, x_n) \geq \mathbf{0}$ is a vector of protein or other molecule concentrations internal to the system, $\mathbf{u} = (u_1, \dots, u_m)$ is a vector of external input signals, e.g. nutrients and f_i a function.

The above formalism can be extended with discrete time-delays, $\mathbf{x}_\tau = (\mathbf{x}(t - \tau_1), \dots, \mathbf{x}(t - \tau_p))$, where τ_i are positive constants, to deal with e.g. the time required for transcription and translation. The system is thus transformed into a DDE model.

Piecewise linear differential equations Several cellular networks have been modelled by piecewise linear differential equations, e.g. [17, 29, 30]. We here follow the notation of Gonçalves [31]. Piecewise linear systems (PLS) are constructed from a set of affine linear systems,

$$\dot{\mathbf{x}} = \mathbf{A}_\alpha \mathbf{x} + \mathbf{B}_\alpha \quad (12.2)$$

where $\mathbf{x} \in \mathbb{R}^n$, \mathbf{A}_α an $n \times n$ matrix, \mathbf{B}_α an $n \times 1$ input vector, and α is a *switching rule*,

$$\alpha(\mathbf{x}) \in \{1, \dots, M\} \tag{12.3}$$

which describes when to switch between the linear systems. The switching rule depends on the present state $\mathbf{x}(t)$ and might also depend on earlier states e.g. $\mathbf{x}(t - \tau)$. We denote a switching rule that only depends on the present state as *memoryless*. A solution to (12.2)–(12.3) are functions (\mathbf{x}, α) , satisfying (12.2)–(12.3), where α is piecewise constant. A *switching time* of a solution (\mathbf{x}, α) is a time t where $\alpha(t)$ is discontinuous. A trajectory *switches* at a time t if t is a switching time. Switching occurs at *switching surfaces* in the state space of \mathbf{x} . If the switching rule is memoryless and consists of linear inequalities, then these surfaces are hyperplanes of dimension $n - 1$,

$$\mathbf{S}_j = \{\mathbf{x} | \mathbf{C}_j \mathbf{x} + d_j = 0\} \tag{12.4}$$


where \mathbf{C}_j is a $1 \times n$ vector and $j \in \{1, \dots, N\}$. For a PLS let us define a partition of the state space, where a separate linear system is used, as

$$\mathbf{X}_i = \{\mathbf{x} | \alpha(\mathbf{x}) = i\} \tag{12.5}$$

for $i = \{1, \dots, M\}$. When the switching rule has no memory, $\mathbf{X}_i \cap \mathbf{X}_j = \emptyset, i \neq j$, and in each partition the dynamics is given by the linear system $\dot{\mathbf{x}} = \mathbf{A}_i \mathbf{x} + \mathbf{B}_i$. A phase portrait of a two dimensional piecewise linear system is displayed in Fig. 12.1.

Switching rule with memory If the switching rule has a memory, i.e. it does not only depend on $\mathbf{x}(t)$, but also earlier states, e.g. $\mathbf{x}(t - \tau)$, then the intersection of different \mathbf{X}_i might not give the empty set. One example of this is given in [23], where the switching rule not only depends on the present state but also on a past state $\alpha = \alpha(\mathbf{x}(t), \mathbf{x}(t - \tau))$. In [23], at each phase point \mathbf{x} , one out of two linear systems can be used, depending on the value of $\mathbf{x}(t - \tau)$.

Fig. 12.1 A phase portrait of a piecewise linear system consisting of two linear systems defined in $\mathbf{X1}$ and $\mathbf{X2}$ respectively. The fixed points $\hat{\mathbf{x}}_1$ and $\hat{\mathbf{x}}_2$ of the respective linear systems are also indicated. From [23]



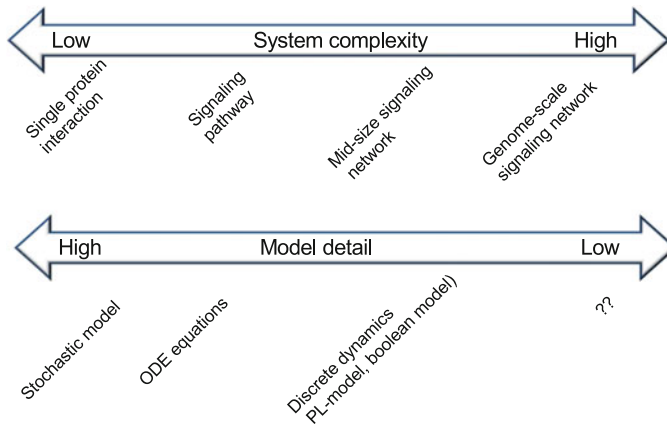


Fig. 12.2 Model detail is determined by system complexity. A single protein interaction can be modelled in great detail using stochastic models. For larger and thus more complex biological systems less detailed models are more informative. Dynamical models on a genome scale are not yet feasible. Modified from [9]

12.4.2.1 Choice of Model Detail

As our focus shifts from modeling single pathways to increasingly complex cellular networks, the computational methodology and formalism used must be carefully considered. Extrapolating the traditional ODE model used for modeling single pathways to whole-cell systems would make the model prohibitively complicated. This is discussed in [9], where it is suggested that finding a “course-grained” level of model description where the molecular details are left out when possible and focusing on the system behaviour could be one strategy to solve this problem, see Fig. 12.2. This theme is also elaborated by de Jong [16] who suggests that whole-cell models could be organized in a hierarchical manner, based on the inherent modularity of the cell. On different levels of abstraction different modeling formalisms could be used, thus on a higher level a more “course-grained” formalism would be appropriate.

12.5 Model Simplification and Reduction

Model simplification and reduction in order to enable analysis is not a new concept within biology. The most famous example is perhaps the Michaelis-Menten equation where a separation in time-scales justifies a “quasi-steady-state” approximation. Throughout this chapter, we use the term model simplification as a more loose description that a model is simplified in some sense, whereas the term model reduction to describe simplifications where the number of degrees of freedom in the model is reduced, while the model formalism remains the same.

Which method to use when simplifying a particular full-scale model depends on the objective of the simplification process and the equation structure of the full-scale model. Before a model simplification process is performed it is useful to decide on which features of the original model to be retained, in [68] named *reduction target* and in [65] denoted as *reference data*. The simplified model must of course be consistent with at least some aspect of the full-scale model. It can, however, in most cases, by definition, not reproduce all possible dynamics of the full model under all possible circumstances. This was investigated in [26], where it was found, that reduced models that reproduced predictions of the full model for a particular set of parameters loosed their predictive capacity when parameters were varied over two-fold ranges.

One can note that the model simplification process, in the same way as ordinary modeling, is an approach to essentially obtain insights on how the underlying system works. Model simplification can therefore be seen as the “modeling of an already existing model”. The proper approximations made during the model simplification process could therefore be instrumental in learning about the underlying processes. As was noted earlier, biological systems are *sparse* and seem to have a *modular* structure [39], features that could aid a reduced model description.

There are a number of reduced models described for a variety of biological systems and some early examples include [48, 67, 85]. The model reduction procedures often include *ad hoc* components, which require that the modeler is intimately familiar with the dynamics of the original system. Attempts of more systematic model reduction approaches have been performed e.g. [34, 47, 65, 68]. However, also such approaches are often only applicable to a small subset of biological systems of a particular structure (e.g. systems having one ubiquitous variable) and/or dynamical behavior (e.g. hopf bifurcation). Hence, there is a need for development of more general methods that would enable systematic simplification and evaluation of complex non-linear systems.

12.5.1 Model Simplification and Reduction Approaches

We here give examples of some approaches towards model simplification and reduction. This list is by no means exhaustive. A review can be found in [58]. Some model simplification strategies include a first step of decomposing the system into subsystems, whereas others remove or simplify individual parts (e.g. equations or reactions) one-by-one. As described earlier there is also a difference between methods that remains within the same formalism and methods that translate the model into a new formalism. Examples of the first case include lumping of variables, separation of timescales, and removal of variables based on sensitivity analysis or identifiability, whereas examples of the second case include boolean approximations, introduction of explicit time delay, and hybrid approaches; as the simplified model described at the end of this chapter, a delayed piecewise linear approximation of an ordinary differential model.

Module based model decomposition Cellular functions have, as was described earlier, long been suggested to be carried out in a modular manner [7, 39, 78], and different methodological approaches have tried to utilize this for model decomposition, in order to retrieve model subsystems that can be easier analyzed than the original system. One example can be found in [60, 61] where the original system is divided into subsystems or modules of low *retroactivity*. Loosely defined, retroactivity describes the phenomenon in which a downstream event affects upstream reactions. This concept emerged in the field of electric engineering. When the output of one electric unit is connected as input to another electric unit, this can affect the first electric unit retroactively, for example by draining electrons too fast [4]. A unit would be without retroactive effects if both input and output are unidirectional. This means that the behavior of the unit only depends on its input, and that ‘connecting’ it to another unit does not change its input-output behavior. The behaviour of retroactive-free modules can be studied uncoupled with the system and analyzed by systems theory’s tools. Kinetic insulation is another concept related to modularity and a mechanism suggested to increase “isolation” between modules by the use of different timescales [18] and segregate between different signals that are using common pathway components [8].

Combining variables, or lumping refers to the process of reducing the number of dimensions of a system by merging states (e.g. protein concentrations) together. This approach is well suited for biological of the often occurring modular structure, and as the remaining states can have a biological meaning, like sums of protein concentrations. This can be illustrated by the following system



where S is a substrate that turns into P through two different intermediates I_1 or I_2 . The species I_1 and I_2 can be lumped together to the new pseudo-species $I = I_1 + I_2$, resulting in



Lumping of variables leads to a simplification in terms of the number of states and reactions of the systems but this might come to the cost of a higher complexity of the remaining rate equations. A description of lumping procedures can also be found in [58]. Lumping can be divided into two categories, in *proper lumping* each of the species of the original model contributes only to one of the lumps of the reduced model, whereas in *improper lumping* species can contribute to more than one lump.

The choice of which variables to be lumped together is often decided on an intuitive basis corresponding to an understanding of the specifics of the system. An example of a more systematic method is the use of *equivalent potentials* [44, 68] used to reduce neuron models. Here the special structure of these neuron models enabled the design of an automated method. This method first uses a nonlinear transformation to

the state variables, in order to find similarities between them. Finally, the variables are sorted into groups based on these similarities. Next, each group is replaced by a new variable approximating the function of the original state variables. In [72, 73] the back-translation properties of a reduction process for biochemical models are emphasized, and they present a method based on lumping of clusters of fast variables, where it is possible to map from reduced model variables and parameters back to original model variables and parameters. A systematic method for proper lumping of the systems state is also described in [19] which is based on a search through all possible combinations of lumps. The combinatorial complexity of the problem is bypassed through a heuristic, greedy algorithm.

Separation of timescales, is another method to reduce the complexity of models. As a rule, biological processes occur over a broad range of time-scales, from milliseconds (e.g. phosphorylation reactions), minutes (e.g. transcription) to hours (e.g. cell cycle), days and years, and the principle of time-scale separation has been widely used within biology. One example is the *quasi-steady-state approximation* (QSS), which is used within enzyme kinetics to derive the well known *Michalis-Menten* equation. In QSS, the variables x are decomposed into two blocks x_f , fast variables and x_s slow variables: $\frac{d}{dt} \begin{pmatrix} x_f \\ x_s \end{pmatrix} = \begin{pmatrix} F_f(x_f, x) \\ F_s(x_f, x_s) \end{pmatrix}$. The reactions including the fast variable x_f are approximated as to be instantaneous compared to the reactions including the slow variable x_s . For given fast species we can simply set $\frac{dx_f}{dt} = 0$, which results in an algebraic relation

$$F_f(x_f, x_s) = 0.$$

If the functions F_f is nice, the Implicit Function Theorem can be applied, that is, there is a unique solution $x_f = G(x_s)$. Substituting it into the original system yields a lower dimension system $\frac{dx_s}{dt} = F_s(G(x_s), x_s)$. This means that we can study the lower dimensional system on the slow manifold $F_f(x_f, x_s) = 0$. The generalization of this technique is difficult since it is no trivial task to divide variables into fast and slow. There are also conditions that have to be satisfied for the method to work. In [37] a nice extension to the QSS is described based on the zero-derivative principle. For a more detailed description on timescale separation see [58] or [25].

Introduction of explicit time delay, is used in a few studies as a mean to simplify biological models e.g. [22, 69]. Consecutive interactions in an ODE model often give rise to a time delay in the system. In the above studies, intermediate variables are removed and instead represented by an explicit delay, reducing the number of variables, and turning the system into a system of delayed differential equations (DDE). It can be noted that DDE systems, however, being infinite dimensional, are in general harder to analyze than ODE systems.

Aside from model reduction, the question whether to use explicit delay or slow intermediate variables, when modeling biological systems has been discussed [55]. Oscillations in biological systems are often assumed to be due to a delayed feedback loop and the delay is often modeled by slow intermediary variables. This may have

the consequence that processes such as transcription and translation are assumed to be instantaneous. In [55] it is shown in a model of the oscillatory expression of Hes1, p53 and NF- κ B that no intermediate variables are needed to get oscillations if an explicit transcriptional delay is introduced. This is a principally important result, since it is preferable to include an anonymous delay in the system (that of course could be due to an intermediate variable) than a variable that correspond to a protein that might not exist. If a truly delayed process is modeled as instantaneous (even though slow) this can result in erroneous parameter estimates.

Sensitivity analysis based methods A sensitivity analysis (SA) investigates how the model output depends on the model parameters and can be performed through a *local* SA, describing infinitesimal changes around a nominal point in parameter space, or through a *global* SA, investigating larger parts of the parameter space, often by statistical analysis. More information about sensitivity analysis can be found in [80]. In model reduction SA is often used through two steps. First all parts (e.g. terms, reactions, variables) of a model are ranked according to their influence on the model output. In the next step parts that are considered to be unimportant for the considered behaviour are removed (corresponding to low ranked parts) [11, 58]. This is not a trivial task, however, since individual parts can have a low sensitivity, but combinations of parts can be crucial. Also for a local sensitivity analysis the result depend on the nominal parameter value that the sensitivity analysis starts off from. Therefore, an individual ranking list can only be considered as a guide. Also notable is that the original model structure must be considered during the removal of variables, so that the remaining model is biochemically consistent. One example of model reduction based on SA can be found in [51], where SA together with flux analysis and principal component analysis is used to reduce a model of epidermal growth factor (EGF) mediated signaling and trafficking. SA is also used in [5] to guide the order of removal of parameters in the model.

Identifiability. In [65] the concept of unidentifiability is used to reduce rate expressions. The rationale is that an unidentifiable rate expression has more than one parameter set that can describe data equally well and that simpler expressions therefore could be used. In the systematic method of [65], the rate expressions have to be in rational form, i.e. a fraction between two polynomials. The reduction is next performed in a reaction-wise manner, so that the complexity of the individual rate expressions is reduced, while the structure of the cellular network is conserved. The reduction is performed in relation to a *reference data* set, corresponding to in silico simulations of the original full-scale model, and parts of the reaction rates that are unidentifiable with respect to this data are removed. By this method terms of the reaction rates that are less important for the model behaviour, as represented by the reference data, are removed. This procedure thereby identifies the functionally important interactions. Another important study exploits the profile likelihood to detect structural and practical non-identifiabilities and suggests the use of this methodology for model reduction [59]. We also refer to [45] where issues related to identifiability are discussed.

Topological filtering and viable parameter spaces. Uncertainties in model structure often makes it necessary to evaluate more than one model topology, where different topologies correspond to different biological hypothesis (see also [70]). In a recent study [71] a “supermodel” is constructed that incorporates all hypothetical mechanisms concerning the underlying biological system. This includes the region of the parameter space of this supermodel that is compatible with the experimental data, called the viable space. The exploration of such a large multidimensional space is done by a procedure that combines an initial coarse grained global sampling of the viable space with a subsequent finer grained exploration [86]. This supermodel is then reduced to a set of new models, with reduced topology, by an iterative procedure where parameters are eliminated.

Another related study [5] is also based on the exploration of a parameter region where the model yields some required output, here called admissible region. Based on the shape of this region, parameters and variables are removed and lumped as long as the dynamical behavior of some target species are preserved. The authors further use sensitivity analysis as guidance for ordering the parameters during the reduction procedure.

Change of modeling formalism. The decision on which modeling formalism to use when describing a system, includes a decision on model detail [9, 16] and a decision on which underlying processes that are important to include in a model. A stochastic model representing all molecules of the system is more detailed than an ODE model which assumes that molecular concentrations are sufficiently large in order to be approximated with a continuous and deterministic description. An ODE model in turn is more detailed than a boolean model. Therefore, as an example, can a boolean model be seen as a reduction of a certain class of ODE models where the variables can be described as being on or off. This technique was used by Albert and co-workers to identify conditions for robustness in a large kinetic ODE model of the *Drosophila* segmentation [2], and by Davidich and Bornholdt to reduce an ODE model of the cell cycle [15]. Other ways to reduce model complexity by a change of formalism is to transform a nonlinear ODE model to a piecewise linear ODE model [22] or a hybrid model [34].

12.5.2 “Tearing, Zooming and Simplifying”—an Example of a Model Simplification Procedure

We close this chapter by describing a specific example of a model simplification procedure developed in [23]. The core concept is to utilize biological properties such as modularity in order to identify a simplified description of the system and secondly to develop this methodology using a well characterized biological system. Hence, we start off from a medium sized non-linear ODE-model of the cell division cycle which has been pioneered by the work of Tyson and Novak [57]. The method effectively produces a smaller (in terms of variables and parameters) model in the

mathematical form of a delayed piecewise linear (DPL) system. The fact that the simplified model is piecewise linear facilitates the analysis and enables detailed predictions on parameter relationships that regulate model output such as the cell size. The tearing-zooming-simplifying approach consists of three steps:

1. tearing or subdividing the original model into subsystems (modules),
2. zooming out and characterizing the modules by input/output transfer functions
3. replacing the modules by simplified descriptions.

This is an iterative process where e.g. information from the second step can suggest a new subdivision in the first step. This simplification method is akin to the system theoretical approach to model large engineering systems by tearing and zooming, see e.g. [14, 83]. The concept of zooming has also been used recently in [72, 73].

The variables of the cell cycle model correspond to protein concentrations, and before we go into the details of the tearing-zooming-simplifying approach we will first shortly describe the cell division cycle.

12.5.2.1 The Cell Division Cycle

Cells reproduce by duplicating their contents and then dividing into two. This process contains two parts, the chromosomic cycle and the cytoplasmic cycle [3]. The chromosomic cycle consists of the exact duplication of DNA, *DNA synthesis*, and the subsequent separation of the two copies, *mitosis*. In parallel to this process, during the cytoplasmic cycle, all other constituents are doubled in quantity, the cell grows and the whole cell is divided into two cells. There needs to be a coordination between, and interactive control of, the chromosomic and cytoplasmic cycles. This is achieved by the cell-cycle genes and proteins. Failure in regulation of the cell cycle can result in uncontrolled cell growth and the initiation of cancer.

Several mathematical models have been constructed to account for this system e.g. the pioneering work encoded in the experimentally constrained models of Novak and Tyson [77], one of which [57] is used in this model simplification example. During the eukaryotic cell cycle in fission yeast, the cell grows, DNA is replicated (*S-phase*), and divided into two daughter cells (*M-phase*). Between the S-phase and the M-phase there are also two gap-phases, referred to as *G1* and *G2*.

12.5.2.2 Tearing or Subdividing the Original Model into Subsystems

The final goal of the simplification procedure is to find more or less isolated units (modules), with a dynamical behavior that can be replaced by simpler descriptions. The first step in achieving this consists of dividing the original model into subsystems, based on the topology of the interactions. These potential modules—in the first iteration—consists of a subset of original variables, and have a corresponding set of coupled ODEs. It is important that the coupling between the ODEs within the module is intact, and not torn apart by the subdivision. Therefore a graph describing

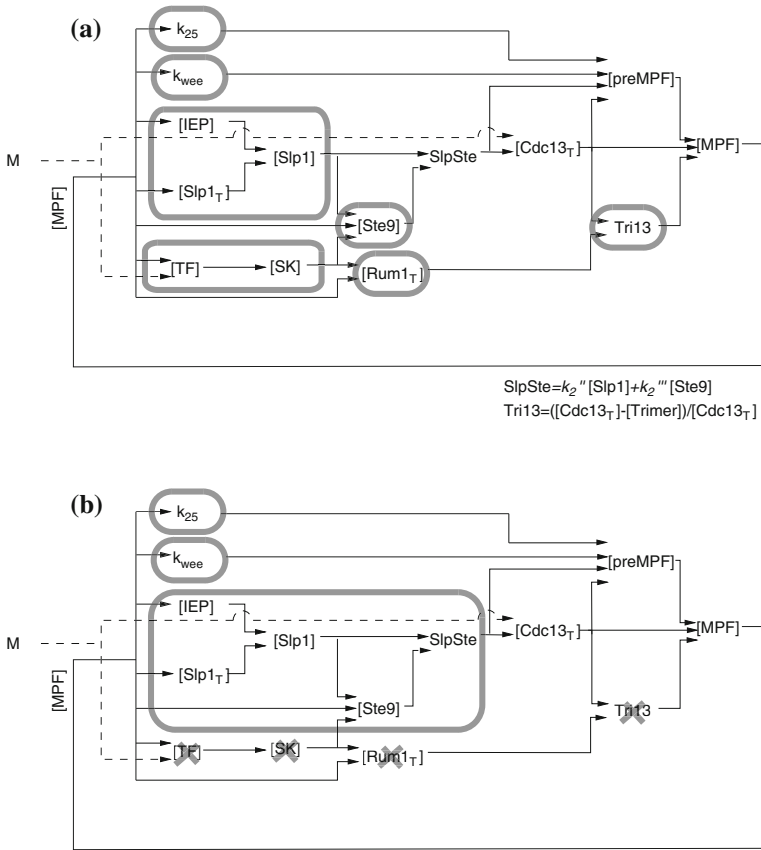


Fig. 12.3 Graph describing the coupling between ODEs, and subdivision into modules (from [23]). A node in this graph corresponds to a variable of the original model, and there is an edge from node j to i if j directly affects i (i.e. is on the right hand side of a differential equation defining i). This graph was used to divide the system into potential modules. **a** Full DPL-model, subdivision into switching modules when $M \geq 0$. **b** Small DPL-model, subdivision into switching modules when $M > 0.8$. Some of the variables can be replaced by constant parameters (indicated by crosses), when $M > 0.8$

the coupling structure of the ODEs is constructed, see Fig. 12.3. The graph illustrates how the dynamics of different variables depend (directly) on other variables in the model. Interestingly, this kind of graphical approach for decomposition was recently utilized in a study on observability of biochemical systems [52].

Putative modules are chosen so that (i) each module has one output-variable, and (ii) all nodes within the module are connected to the output by a connected path. Some of the nodes within a module correspond to ODEs which depend on variables coming from outside the module, denoted input-variables. In addition to defining modules, the graph is also useful for representing the coupling relevant




Fig. 12.4 Steady-state response and step function approximation. **a** The steady state response of the output $[SK]$ (concentration of Starter Kinase) to the inputs M (*Mass*) and $[MPF]$ (Mitosis Promoting Factor). **b** The same function as in **a** but approximated by the step function S_{sk} . From [23], for a colored version see the online version of the book

for elucidating dynamical motifs. In the case of the cell cycle model, the feedback structure within the network and the central role of the protein complex $[MPF]$ is made more transparent by our procedures (Fig. 12.3).

12.5.2.3 Zooming Out and Characterizing the Module Dynamics by Input/output Functions

In the next step the dynamical behavior of the potential modules is analyzed by investigating the input/output relationship. This is performed for each module in isolation by two complementary procedures. First by describing the *steady-state response* of the output in response to different inputs (for an example see Fig. 12.4), and secondly to consider the *response time*, i.e. the time it takes for the module to reach steady state after a significant change of input.

Based on such a characterization and the additional constraint that we avoid modules to be overlapping, leads to the final modules being selected so that (including the earlier defined criteria) (i) each module has one output-variable, (ii) all variables within the module is connected to the output by a connected path, (iii) the modules are non-overlapping, and (iv) the output has a switching input/output behaviour. By switching input/output behaviour we mean a steep, sigmoid-like, response curve (e.g. Fig. 12.4). Most of the variables in this cell cycle example have a dynamics based on *Michaelis-Menten* or *Hill* kinetics, and modules can therefore readily be identified which have a steep sigmoid almost step-like steady-state dependency on the input.

During this process we also remove variables for which the signal response curve are constant or almost constant, by re-defining them into constant parameters. In the cell cycle example this can be done since some of the variables are not “active” in the regime we are interested in, i.e. when the cell mass is larger than 0.8 ($M > 0.8$). In Fig. 12.3 the final modules are illustrated, and in Fig. 12.4 an example of a input-output curve can be seen.

12.5.2.4 Replacing the Modules by Simplified Descriptions

In the final step, the modules are replaced by step functions (Fig. 12.4) and a time delay corresponding to the measured response time (for details see supplementary material of [23]). This transforms the original system into a piecewise linear system (with delay). Not all variables are included into modules, some are removed as described in the previous section, and some become state-variables of the new simplified system. It can be noted that it is not obvious that replacing the modules with step functions would turn the remaining system into a DPL-model. This is possible in this case due the special form of the remaining ODEs of the original model. Whether this is true also for other models remains to be explored.

12.5.2.5 Elucidating the Core Dynamics by a Simplified Model and Making Predictions

The simplified (DPL) model of the cell cycle emphasizes some important features of the system. In the DPL-model, one normal division of the cell, corresponds to the subsequent move between four linear subsystems (Fig. 12.5). Interestingly, these

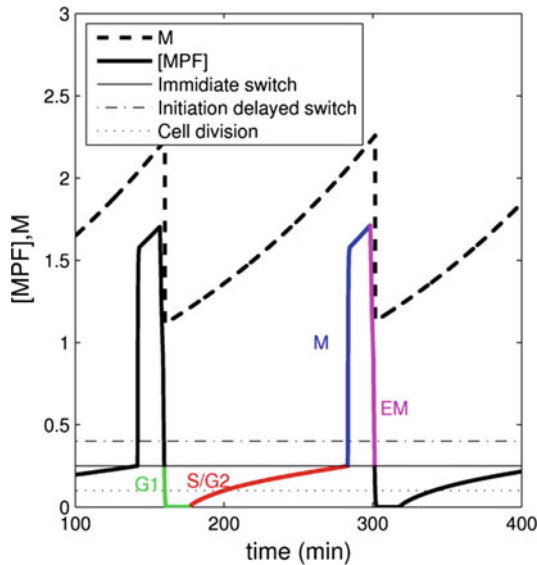



Fig. 12.5 Numerical simulation of the of the DPL-model, showing the time evolution of the cell mass (M), as well as the concentration of the protein complex $[MPF]$ (Mitosis Promoting Factor). During different parts of the cell division cycle, different linear systems are used, as indicated on the time course of the variable $[MPF](t)$ with green, red, blue and magenta. The linear systems correspond to the four cell cycle phases G1, S/G2, M and EM, where EM is the ending of Mitosis. Which linear system that is used at a certain time t depends on $[MPF](t)$ and $[MPF](t - \tau)$ as detailed in [24]. From [24]

subsystems can in turn be directly mapped to different phases of the cell cycle (G1, S/G2, Mitoses, and ending of Mitoses). This means that during each cell cycle phase only one (linear) subsystem is active, corresponding to a subpart of the whole system, and thus, to a subset of system parameters. Therefore, during that phase only a few of the system parameters are “active”, and can influence the behavior of the cell. The other parameters are “silent”, and can be changed without changing the behavior, as long as they are changed back when it is their turn, i.e. when their linear system is active. This also means that different parameters can only act as system-controls during specific timings of the cell cycle, there is a separation in time between different functional modules.

The change between linear system (and thus “active” system parameters), occurs when the system trajectory (i.e. protein concentrations), passes certain thresholds (phase space switching lines), either immediately or after a certain delay. Based on the DPL model it can be calculated when this occurs, and which and how different parameters can affect this. As an example, the switch between the G1 and the S/G2 phase is important for the final size of the cell, and by the DPL-model it can be predicted how large the cell will be if there is a change in one or more of the model parameters. It can also be predicted which parameters that can change the size of the cell, and which combination should be most effective. Other mechanistic principles can also be found, like the observation that the length of the G1 phase corresponds to the time delay of one of the system modules. These predictions are qualitatively, and semi-quantitatively verified in the original model.

The analysis of the DPL-model is based on the calculation of different important dynamical features like stability, fixed points, switching thresholds, eigenvalues and eigenvectors, and analytical conditions for these. The fact that the DPL-model consists of linear systems enables this, even though the inclusion of delay is a complicating factor. In the study presented in [23] it is shown that the dynamics of the DPL-model very well could be approximated by the slow eigenvectors of the different linear systems (Fig. 12.6), and this facilitated a proof on the global stability of the system. Finally, a further interesting feature were observed in the original model and explained in the simplified model [24]. The effect of different parameter perturbations was explored through two different characteristics; an essential effect of a perturbation made the system stop working (stop oscillating), whereas a modulatory

Fig. 12.6 Validation of the simplified DPL-model as well as the slow eigenvector approximation. From [23], for a colored version see the online version of the book



effect had a more minor but significant effect (change in cell size). It was noted that several parameters with no modulatory effect whatsoever for smaller perturbations could in fact be essential to the system, at larger perturbations.

12.5.2.6 What Is Lost by the Simplification?

Three important assumptions are made in the simplification process described above, (i) variables have time to get sufficiently close to their steady-state before there is a significant change of input, and (ii) that the transient behaviour is not important and (iii) that the exact form of the steady-state is not critical for the system, for example, a sigmoid function can be substituted by a step-function. If this is not the case then the simplification will fail. One idea behind this is that biological systems appear to be robust to many biochemical details [56, 82], and that this can be utilized to retrieve simplified coarse-grained descriptions.

The simplification of the original model was performed in a nominal point in the parameter space. In an extended investigation [24] it was analyzed how well the simplified model reproduced the dynamical behavior of the original model for other parameter values than the nominal point. It was noted that the models agreed well for smaller parameter perturbation, but for larger there was a disagreement.

12.6 Conclusions

We have within this chapter described how the huge increase in experimental cellular data of the last decades, and the following increase in the size and number of dynamical models, has been followed by an increasing need for and development of methods for model analysis through simplification. Since dynamical cellular models can consist of several hundreds of parameters, representing interactions between large numbers of species, connected in a nonlinear way by intertwined feedback-loops it is difficult to determine which parts of the models are important for different output. Another complicating factor is the model uncertainty, both when it comes to parameter values as well as topology. New experimental data are both qualitative, e.g. describing existence of interactions between species, as well as quantitative, e.g. describing interaction strength, but there is a mismatch between the amounts of qualitative versus quantitative data, resulting in not fully constrained or non-identifiable models.

Model simplification and reduction has herein been presented as a method to meet these challenges of complexity and identifiability of large-scale cellular models. The idea is to retrieve a smaller or in some other sense simpler model, with increased transparency (easier to understand intuitively), increased identifiability and/or increased predictability. It must also be possible to map the parts of the simplified model back to the biological entities or features of the original model, for predictions to be relevant.

Other utilizations of simpler models, which are not discussed so much in this chapter, are reduced computational times, as well as a means to compare models to each other.

Many different approaches have been taken towards simplification and reduction of models. There are ‘horizontal’ approaches, working on the same scale or level in the system hierarchy, for example lumping entities of the same kind together, like proteins or reactions, or decomposing the system into subsystems and then simplifying or analyzing these separately. Other ‘vertical’ approaches are approximating features on other scales than the one considered; like approximating a stochastic model by the average number of species to receive an ODE model; or approximating continuous functions with discrete functions, when transforming an ODE model to a Boolean or piecewise linear model; or the common theme of reducing variables working on faster time-scales. The different approaches differ in methodology, where, for example, horizontal approaches use model topology, sensitivity, identifiability, etc., vertical approaches often use approximations of different kinds, like averages.

It can also be noted that there is a difference between “local” and “global model simplification” approaches, i.e. some approaches are performing the simplification at a specific point in parameter space (e.g. time scale separation), whereas others are taking into account the full (or parts of the full) admissible parameter space, i.e. the parameter values for which the model output is consistent with experimental data (e.g. methods based on identifiability).

Traditional means of analyzing nonlinear dynamical models, like bifurcation analysis, do not suffice in order to understand these new large-scale models, rather new approaches have to be taken. Here model simplification and reduction have an important part to play.

Acknowledgments We would like to thank Yishao Zhou and Omar Gutierrez-Arenas for valuable discussions and important input, Sara Maad Sasane for commenting on the manuscript and the Swedish e-Science Research Center (SeRC) and the the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no 604102 (HBP) for financial support.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

1. Adamczyk, M., van Eunen, K., Bakker, B.M., Westerhoff, H.V.: Enzyme kinetics for systems biology: When, why and how. *Methods. Enzymol.* **500**, 233–257 (2011)
2. Albert, R., Othmer, H.G.: The topology of the regulatory interactions predicts the expression pattern of the *Drosophila* segment polarity genes. *J. Theor. Biol.* **223**, 1–18 (2003)
3. Alberts, B., Bray, D., Lewis, J., Raff, M., Roberts, K., Watson, J.D.: *Molecular biology of the cell*. Garland, New York (1994)

4. Alexander, R.P., Kim, P.M., Emonet, T., Gerstein, M.B.: Understanding modularity in molecular networks requires dynamics. *Sci. Signal.* **2**, pp. pe44 (2009)
5. Apri, M., de Gee, M., Molenaar, J.: Complexity reduction preserving dynamical behavior of biochemical networks. *J. Theor. Biol.* **304**, 16–26 (2012)
6. Bagowski, C.P., Besser, J., Frey, C.R., Ferrell Jr, J.E.: The JNK cascade as a biochemical switch in mammalian cells: ultrasensitive and all-or-none responses. *Curr. Biol.* **13**, 315–320 (2003)
7. Barabasi, A.L., Oltvai, Z.N.: Network biology: understanding the cell's functional organization. *Nat. Rev. Genet.* **5**, 101–113 (2004)
8. Behar, M., Dohlman, H.G., Elston, T.C.: Kinetic insulation as an effective mechanism for achieving pathway specificity in intracellular signaling networks. *Proc. Nat. Acad. Sci.* **104**, 16146–16151 (2007)
9. Bornholdt, S.: Less is more in modeling large genetic networks. *Science* **310**, 449–451 (2005)
10. Bruggeman, F.J., Westerhoff, H.V.: The nature of systems biology. *Trends Microbiol.* **15**, 45–50 (2007)
11. Cedersund, G.: Core-box modelling, theoretical contributions and applications to glucose homeostasis related systems. Ph.D. thesis, Chalmers University of Technology, Sweden, 2006
12. Chen, W.W., Schoeberl, B., Jasper, P.J., Niepel, M., Nielsen, U.B., Lauffenburger, D.A., Sorger, P.K.: Input-output behavior of ErbB signaling pathways as revealed by a mass action model trained against dynamic data. *Mol. Syst. Biol.* **5**, 239 (2009)
13. Christensen, C., Thakar, J., Albert, R.: Systems-level insights into cellular regulation: inferring, analysing, and modelling intracellular networks. *IET Syst. Biol.* **1**, 61–77 (2007)
14. Csete, M.E., Doyle, J.C.: Reverse engineering of biological complexity. *Science* **295**, 1664–1669 (2002)
15. Davidich, M., Bornholdt, S.: The transition from differential equations to boolean networks: A case study in simplifying a regulatory network model. *J. Theor. Biol.* **255**, 269–277 (2008)
16. de Jong, H.: Modeling and simulation of genetic regulatory systems: a literature review. *J. Comput. Biol.* **9**, 67–103 (2002)
17. de Jong, H., Gouzé, J.L., Hernandez, C., Page, M., Sari, T., Geiselmann, J.: Qualitative simulation of genetic regulatory networks using piecewise-linear models. *Bull. Math. Biol.* **66**, 301–340 (2004)
18. Del Vecchio, D., Ninfa, A.J., Sontag, E.D.: Modular cell biology: Retroactivity and insulation. *Mol. Syst. Biol.* **4**, (2008)
19. Dokoumetzidis, A., Aarons, L.: Proper lumping in systems biology models. *IET Syst. Biol.* **3**, 40–51 (2009)
20. Duarte, N.C., Becker, S.A., Jamshidi, N., Thiele, I., Mo, M.L., Vo, T.D., Srivas, R., Palsson, B.Ø.: Global reconstruction of the human metabolic network based on genomic and bibliomic data. *Proc. Nat. Acad. Sci.* **104**, 1777–1782 (2007)
21. Ehrenberg, M., Elf, J., Aurell, E., Sandberg, R., Tegnér, J.: Systems biology is taking off. *Genome Res.* **13**, 2377–2380 (2003)
22. Eriksson, O., Zhou, Y., Tegnér, J.: 'Modeling complex cellular networks - robust switching in the cell cycle ensures a piecewise linear reduction of the regulatory network'. *CDC: IEEE Conf. Decis. Contr.* **2004**(1), 117–123 (2004)
23. Eriksson, O., Brinne, B., Zhou, Y., Björkegren, J., Tegnér, J.: Deconstructing the core dynamics from a complex time-lagged regulatory biological circuit. *IET Syst. Biol.* **3**, 113–129 (2009)
24. Eriksson, O., Andersson, T., Zhou, Y., Tegnér, J.: Decoding complex biological networks - tracing essential and modulatory parameters in complex and simplified models of the cell cycle. *BMC Syst. Biol.* **5**, 123 (2011)
25. Ermentrout, B.: Simplifying and reducing complex models. In Bower, J.M., Bolouri, H. (Eds.) *Computational Modeling of Genetic and Biochemical Networks*, MIT Press, Cambridge, MA, pp. 307–323
26. Faeder, J.R., Blinov, M.L., Goldstein, B., Hlavacek, W.S.: Combinatorial complexity and dynamical restriction of network flows in signal transduction. *Syst. Biol. (Stevenage)* **2**, 5–15 (2005)

27. Feng, Y., Mitchison, T.J., Bender, A., Young, D.W., Tallarico, J.A.: Multi-parameter phenotypic profiling: using cellular effects to characterize small-molecule compounds. *Nat. Rev. Drug Discov.* **8**, 567–578 (2009)
28. Ferrell Jr, J.E.: Tripping the switch fantastic: how a protein kinase cascade can convert graded inputs into switch-like outputs. *Trends Biochem. Sci.* **21**, 460–466 (1996)
29. Glass, L.: Classification of biological networks by their qualitative dynamics. *J. Theor. Biol.* **54**, 85–107 (1975)
30. Glass, L., Kauffman, S.A.: The logical analysis of continuous, non-linear biochemical control networks. *J. Theor. Biol.* **39**, 103–129 (1973)
31. Gonçalves, J.: Constructive global analysis of hybrid systems. Ph.D. Thesis, Massachusetts Institute of Technology, 2000
32. Gureasko, J., Galush, W.J., Boykevisch, S., Sondermann, H., Bar-sagi, D., Groves, J.T., Kuriyan, J.: Membrane-dependent signal integration by the Ras activator Son of sevenless. *Nat. Struct. Mol. Biol.* **15**, 452–461 (2008)
33. Gutierrez-Arenas, O., Eriksson, O., and Hellgren Kotaleski, J.: Segregation and crosstalk of D1 receptor-mediated activation of ERK in striatal medium spiny neurons upon acute administration of psychostimulants. *PLoS Comput. Biol.* **10**, pp. e1003445 (2014)
34. Halász, A., Kumar, V., Imieliński, M., Belta, C., Sokolsky, O., Pathak, S., Rubin, H.: Analysis of lactose metabolism in *E.coli* using reachability analysis of hybrid systems. *IET Syst. Biol.* **1**, 130–148 (2007)
35. Han, J.D., Bertin, N., Hao, T., Goldberg, D.S., Berriz, G.F., Zhang, L.V., Dupuy, D., Walhout, A.J., Cusick, M.E., Roth, F.P., Vidal, M.: Evidence for dynamically organized modularity in the yeast protein-protein interaction network. *Nature* **430**, 88–93 (2004)
36. Hardie, D.G., Salt, I.P., Hawley, S.A., Davies, S.P.: AMP-activated protein kinase: an ultrasensitive system for monitoring cellular energy charge. *Biochem. J.* **338**, 717–722 (1999)
37. Härdin, H.M., Zagaris, A., Krab, K., Westerhoff, H.V.: Simplified yet highly accurate enzyme kinetics for cases of low substrate concentrations. *FEBS J.* **276**, 5491–5506 (2009)
38. Harper, J.W.: A phosphorylation-driven ubiquitination switch for cell-cycle control. *Trends Cell Biol.* **12**, 104–107 (2002)
39. Hartwell, L.H., Hopfield, J.J., Leibler, S., Murray, A.W.: From molecular to modular cell biology. *Nature* **402**, c47–c52 (1999)
40. Huang, C.Y., Ferrell, J.E. Jr.: Ultrasensitivity in the mitogen-activated protein kinase cascade. *Proc. Natl. Acad. Sci. U.S.A.* **93**, pp. 10078–10083 (1996)
41. Iyengar, R.: Why we need quantitative dynamic models. *Sci. Signal.* **2**, eg3 (2009)
42. Karr, J.R., Sanghvi, J.C., Macklin, D.N., Gutschow, M.V., Jacobs, J.M., Bolival Jr, B., Assad-Garcia, N., Glass, J.I., Covert, M.W.: A whole-cell computational model predicts phenotype from genotype. *Cell* **150**, 389–401 (2012)
43. Kawai, J., et al.: Functional annotation of a full-length mouse cDNA collection. *Nature* **409**, 685–690 (2001)
44. Kepler, T.B., Abbott, L.F., Marder, E.: Reduction of conductance-based neuron models. *Biol. Cybern.* **66**, 381–387 (1992)
45. Kirk, P., Silk, D., Stumpf, M.P.H.: Reverse engineering under uncertainty. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
46. Kitano, H.: Systems biology: A brief overview. *Science* **295**, 1662–1664 (2002)
47. Koschorreck, M., Conzelmann, H., Ebert, S., Ederer, M., Gilles, E.D.: Reduced modeling of signal transduction - a modular approach. *BMC Bioinform.* **8**, 336 (2007)
48. LeBeau, A.P., Robson, A.B., McKinnon, A.E., Sneyd, J.: Analysis of a reduced model of corticotroph action potentials. *J. Theor. Biol.* **192**, 319–339 (1998)
49. Lee, T.I., et al.: Transcriptional regulatory networks in *Saccharomyces cerevisiae*. *Science* **298**, 799–804 (2002)
50. Leloup, J.C., Goldbeter, A.: A model for circadian rhythms in *Drosophila* incorporating the formation of a complex between the PER and TIM proteins. *J. Biol. Rhythms.* **13**, 70–87 (1998)
51. Liu, G., Swihart, M.T., Neelamegham, S.: Sensitivity, principal component and flux analysis applied to signal transduction: the case of epidermal growth factor mediated signaling. *Bioinformatics* **21**, 1194–1202 (2005)

52. Liu, Y.-Y., Slotine, J.-J., Barabási, A.-L.: Observability of complex systems. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 2460–2465 (2013)
53. Lohse, M.J., Nikolaev, V.O., Hein, P., Hoffmann, C., Vilardaga, J.P., Bünnemann, M.: Optical techniques to analyze real-time activation and signaling of G-protein-coupled receptors. *Trends Pharmacol. Sci.* **29**, 159–165 (2008)
54. Luscombe, N.M., Babu, M.M., Yu, H., Snyder, M., Teichmann, S.A., Gerstein, M.: Genomic analysis of regulatory network dynamics reveals large topological changes. *Nature* **431**, 308–312 (2004)
55. Monk, N.A.M.: Oscillatory expression of Hes1, p53, and NF- κ B driven by transcriptional time delays. *Curr. Biol.* **13**, 1409–1413 (2003)
56. Morohashi, M., Winn, A.E., Borisuk, M.T., Bolouri, H., Doyle, J., Kitano, H.: Robustness as a measure of plausibility in models of biochemical networks. *J. Theor. Biol.* **216**, 19–30 (2002)
57. Novak, B., Pataki, Z., Ciliberto, A., Tyson, J.J.: Mathematical model of the cell division cycle of fission yeast. *Chaos* **11**, 277–286 (2001)
58. Okino, M.S., Mavrouniotis, M.L.: Simplification of mathematical models of chemical reaction systems. *Chem. Rev.* **98**, 391–408 (1998)
59. Raue, A., Kreutz, C., Maiwald, T., Bachmann, J., Schilling, M., Klingmüller, U., Timmer, J.: Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics* **25**, 1923–1929 (2009)
60. Saez-Rodriguez, J., Kremling, A., Conzelmann, H., Bettenbrock, K., Gilles, E.D.: Modular analysis of signal transduction networks. *IEEE Control Syst. Mag.* **24**, 35–52 (2004)
61. Saez-Rodriguez, J., Kremling, A., Gilles, E.D.: Dissecting the puzzle of life: modularization of signal transduction networks. *Comput. Chem. Eng.* **29**, 619–629 (2005)
62. Sanghvi, J.C., Regot, S., Carrasco, S., Karr, J.R., Gutschow, M.V., Bolival Jr, B., Covert, M.W.: Accelerated discovery via a whole-cell model. *Nat. Methods* **10**, 1192–1195 (2013)
63. Schena, M., Shalon, D., Davis, R., Brown, P.: Quantitative monitoring of gene expression patterns with a complementary dna microarray. *Science* **270**, 467–470 (1995)
64. Schliemann-Bullinger, M., Fey, D., Bastogne, T., Findeisen, R., Scheurich, P., Bullinger, E.: The experimental side of parameter estimation. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
65. Schmidt, H., Madsen, M.F., Danø, S., Cedersund, G.: ‘Complexity reduction of biochemical rate expressions’. *Bioinformatics* **24**, 848–854 (2008)
66. Smallbone, K., et al.: A model of yeast glycolysis based on a consistent kinetic characterisation of all its enzymes. *FEBS Lett.* **587**, 2832–2841 (2013)
67. Smolen, P., Baxter, A.D., Byrne, J.H.: A reduced model clarifies the role of feedback loops and time delays in the *Drosophila* circadian oscillator. *Biophys. J.* **83**, 2349–2359 (2002)
68. Sorensen, M.E., DeWeerth, S.P.: An algorithmic method for reducing conductance-based neuron models. *Biol. Cybern.* **95**, 185–192 (2006)
69. Srividhya, J., Gopinathan, M.S., Schnell, S.: The effects of time delays in a phosphorylation-dephosphorylation pathway. *Biophys. Chem.* **125**, 286–297 (2007)
70. Sunnåker, M., Stelling, J.: Model extension and model selection. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
71. Sunnåker, M., Zamora-Sillero, E., Dechant, R., Ludwig, C., Busetto, A. G., Wagner, A., and Stelling, J.: Automatic generation of predictive dynamic models reveals nuclear phosphorylation as the key Msn2 control mechanism. *Sci. Signal.* **6**, ra41 (2013)
72. Sunnåker, M., Schmidt, H., Jirstrand, M., Cedersund, G.: Zooming of states and parameters using a lumping approach including back-translation. *BMC Syst. Biol.* **4**, 28 (2010)
73. Sunnåker, M., Cedersund, G., Jirstrand, M.: A method for zooming of nonlinear models of biochemical systems. *BMC Syst. Biol.* **5**, 140 (2011)
74. The 1000 Genomes Project Consortium: An integrated map of genetic variation from 1,092 human genomes. *Nature* **491**, 56–65 (2012)
75. The ENCODE Project Consortium: An integrated encyclopedia of DNA elements in the human genome. *Nature* **489**, 57–74 (2012)

76. The International Human Genome Mapping Consortium: A physical map of the human genome. *Nature* **409**, 934–941 (2001)
77. Tyson, J.J., Chen, K., Novak, B.: Network dynamics and cell physiology. *Nat. Rev. Mol. Cell. Biol.* **2**, 908–916 (2001)
78. Tyson, J.J., Chen, J.C., Novak, B.: Sniffers, buffers, toggles and blinkers: dynamics of regulatory and signaling pathways in the cell. *Curr. Opin. Cell. Biol.* **15**, 221–231 (2003)
79. Uetz, P., et al.: A comprehensive analysis of protein-protein interactions in *Saccharomyces cerevisiae*. *Nature* **403**, 623–627 (2000)
80. Van Schepdael, A., Carlier, A., Geris, L.: Sensitivity analysis by design of experiments. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
81. Venter, J., et al.: The sequence of the human genome. *Science* **291**, 1304–1351 (2001)
82. von Dassov, G., Meir, E., Munro, E.M., Odell, M.O.: The segment polarity network is a robust developmental module. *Nature* **406**, 188–192 (2000)
83. Willems, J.C.: Paradigm and puzzles in the theory of dynamical systems. *IEEE Trans. Automatic Control* **36**, 259–294 (1991)
84. Yasuda, R.: Imaging intracellular signaling using two-photon fluorescent lifetime imaging microscopy. *Cold Spring Harb. Protoc.* **2012**, 1121–1128 (2012)
85. Yildirim, N., Santillan, M., Horike, D., Mackey, M.C.: Dynamics and bistability in a reduced model of the lac operon. *Chaos* **14**, 279–292 (2004)
86. Zamora-Sillero, E., Hafner, M., Ibig, A., Stelling, J., Wagner, A.: Efficient characterization of high-dimensional parameter spaces for systems biology. *BMC Syst. Biol.* **5**, 142 (2011)

Chapter 13

Sensitivity Analysis by Design of Experiments

An Van Schepdael, Aurélie Carlier and Liesbet Geris

Abstract The design of experiments (DOE) is a valuable method for studying the influence of one or more factors on the outcome of computer experiments. There is no limit to the number of times a computer experiment can be run, but they are often time-consuming. Moreover, the number of parameters in a computer model is often very large and the range of variation for each of these parameters is often quite extensive. The DOE provides the statistical tools necessary for choosing a minimum amount of parameter combinations resulting in as much information as possible about the computer model. In this chapter, several designs and analysing methods are explained. At the end of the chapter, these designs and methods are applied to a mechanobiological model describing tooth movement.

Keywords Design of experiments · Sampling parameter space · Sensitivity analysis

13.1 Introduction

The design of experiments (DOE) is a valuable method for studying the influence of one or more factors on physical experiments (see tutorial [21]). Physical experiments can often only be run a limited number of times and can be expensive and time-consuming. Therefore, when performing a sensitivity analysis on a model with many parameters, limiting the number of parameter combinations to be studied is very

A. Van Schepdael · L. Geris (✉)
Biomechanics Research Unit, University of Liège, Chemin des Chevreuils 1-BAT 52/3,
4000 Liège, Belgium
e-mail: Liesbet.Geris@ulg.ac.be

A. Van Schepdael
e-mail: An.VanSchepdael@mech.kuleuven.be

A. Carlier
Biomechanics Section, University of Leuven, Celestijnenlaan 300C,
Box 2419, 3001 Heverlee, Belgium
e-mail: aurelie.carlier@kuleuven.be

© Springer International Publishing Switzerland 2016
L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_13

important. The basic problem of designing such experiments is deciding which factor combinations to examine. The design of experiments (DOE)—introduced by Fisher [7]—was developed for this purpose.

There is no limit to the number of times a computer experiment can be run, but they are often time-consuming. Moreover, the number of parameters in a computer model is often very large and the range of variation for each of these parameters is often larger than in physical experiments. Although there are fundamental differences between physical experiments and computer simulations, the techniques of DOE that were originally developed for physical experimentation can also be used to investigate the sensitivity of a computer model to its parameters with a minimum of computational time.

13.2 Theory

Running a sensitivity analysis of a computer model using the DOE consists of three steps. Firstly, a suitable design, meaning a number of parameter combinations for which we will run the model, has to be set up. The purpose of this design is to get as much information as possible about the influence of the relevant parameters on the outcome of the model at minimal cost. In computer models, this cost is usually the computational time, which is kept low by limiting the number of parameter combinations that is studied. Next, simulations are run with these parameter combinations and finally, the results are analysed and conclusions are drawn [26, 27].

13.2.1 Available Designs

A number of designs are available to conduct a sensitivity analysis [26, 27, 34]. This section provides an overview of the different techniques that are most commonly found in the biomedical literature.

13.2.1.1 OAT-Design

The simplest design is a *one-at-a-time (OAT) analysis*, where each parameter is varied individually. A standard OAT-design uses a reference condition and then changes each parameter individually to a higher and a lower value, while keeping other parameters at the reference value. The difference between the outcome for the high and the low value is then used as a measure of the influence of the parameter on the system. The main advantage of this design is its simplicity and the fact that it only requires $2M$ experiments, with M being the number of parameters studied. It is however impossible to study interactions between parameters, and the effect of the parameters resulting from this analysis might be different when choosing a different reference

condition [13]. The OAT analysis was used by Lacroix [19] and Geris et al. [9] to assess the influence of the value and duration of the initial and boundary conditions on the simulation results of a fracture healing model.

13.2.1.2 Factorial Designs

In factorial designs, the parameters are assigned different values. In two-level designs, two different levels, a high and a low level, are chosen. Several combinations of parameter values are then compared, changing various parameters at the same time. In a *two-level full factorial design*, all possible combinations are examined, requiring 2^M experiments (see Fig. 13.1a). In three-level designs, requiring 3^M runs, the outcome of the model is also studied with parameters at an intermediate level (Fig. 13.1b). The advantage is that the effect of each parameter can be studied, and that interactions between the factors can be examined. Furthermore, no reference condition is required, giving the results more reliability. The main disadvantage is the computational cost [26, 34]. The design requires 2^M runs, which becomes very high when the model contains many parameters. With 30 parameters, this would require 1.07×10^9 runs.

In *fractional factorial designs*, not all of these 2^M or 3^M combinations are examined (Fig. 13.1c). In a two-level full fractional factorial design with six parameters,

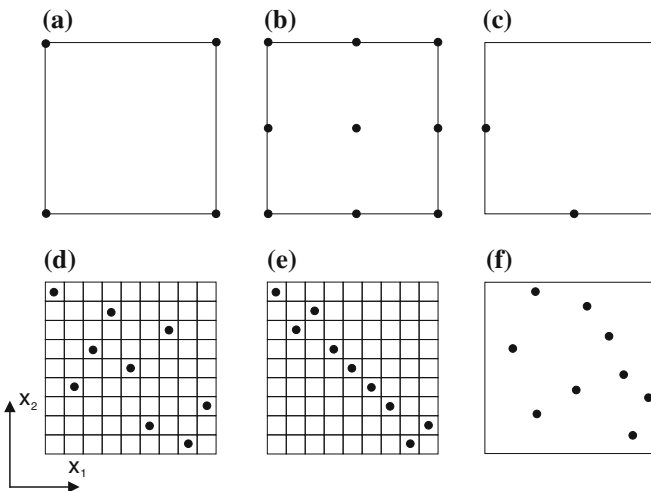


Fig. 13.1 Schematic overview of different designs for two factors x_1 and x_2 . **a** A two-level full factorial design. **b** A three-level full factorial design. **c** A three-level fractional factorial design. **d** A latin hypercube design with nine runs. The parameter space is divided into $9^2 = 81$ cells, and one cell on each *row* and *column* is chosen. **e** A latin hypercube design with nine runs. This example shows that a LHD design is not necessarily space-filling. **f** A uniform design. Note that the factorial designs used discrete values of the parameters, while the LHD and uniform designs spread out the points in space

only six out of 64 runs are used to estimate the main effects, and 15 are used to estimate the two-factor interactions. The remaining runs are used to calculate higher order interactions [27]. Fractional factorial designs are based on the principle that, most likely, at some point the higher order interactions become negligible. It is thus not necessary to examine all possible combinations of parameters, but it is sufficient to choose a suitable set of combinations [27, 34]. By omitting several combinations compared to the full factorial design, the amount of information gained from the sensitivity analysis decreases. The number of runs remains however limited resulting in a significant computational gain. Depending upon the number of experiments, several interactions will become indistinguishable. When using a minimum amount of experiments, only the effect of each parameter separately, the main effects, can be determined. When increasing the number of runs, two-factor interactions can be examined. Generally speaking, few experiments worry about higher order interactions. Fractional factorial designs are classified according to the level of information they provide. A resolution III design is set up in such way that the main effects are distinguishable, but may be affected by one or more two-factor interactions. These must thus assumed to be zero in order for the results to be meaningful. In a resolution IV design, the main effects can be distinguished from the other main effects and the two-factor interactions, but the two-factor interactions are confounded with each other [26, 27, 34]. A fractional factorial design is thus a trade-off between computational cost and accuracy, and are most frequently used to identify a subset of parameters that is most important and needs to be studied more extensively [27]. The main disadvantage is that the parameters are only studied at several levels and the values are not spread out over the entire parameter space.

Several other factorial designs are possible; Plackett-Burman designs, Cotter designs and mixed-level designs offer alternatives to standard fractional factorial designs, each having its own specific advantages and disadvantages. Isaksson et al. [12] determined, for example the most important cellular characteristics for fracture healing using a resolution IV fractional factorial design. Such design was also used by Malandrino et al. [23] to analyse the influence of six material properties on the displacement, fluid pore pressure and velocity fields in the L3–L4 lumbar intervertebral disc.

13.2.1.3 Taguchi's Design

Taguchi's design was originally developed to assist in quality improvement during the development of a product or process. In a manufacturing process, for example, there are control factors and noise factors. The latter cause a variability in the final products and are usually uncontrollable. The goal of robust parameter design is to find the levels of the control factors that are least influenced by the noise factors [26]. In the Taguchi parameter design methodology one orthogonal design is chosen for the control factors (inner array) and one design is selected for the noise factors (outer array).

Taguchi's methodology has received a lot of attention in statistical literature. His philosophy was very original, but the implementation and technical nature of data analysis has received some criticism. Firstly, it does not allow the estimation of interaction terms. Secondly, some of the designs are empirically determined, but are suboptimal compared to rigorous alternatives such as fractional factorial designs [26]. Finally, if the Taguchi approach works and yields good results, it is still not clear what caused the result because of the aliasing of critical interactions. In other words, the problem may be solved short-term, without gaining any long-term process knowledge. Despite this criticism, Taguchi's approach is often used in biomedical literature because of its simplicity [2, 20, 41].

13.2.1.4 Space-Filling Designs

In space-filling designs, the parameter combinations are spread out over the entire parameter space, enabling the design to capture more complex behaviour [36]. This approach is particularly useful for deterministic or near deterministic systems, such as computer simulations. To achieve an effective spreading of the parameters, several sampling methods are available. One of the most used methods is *latin hypercube sampling (LHD)*. This method can be most easily explained by using the very simple example of a 2D experimental region, representing a system with 2 parameters x_1 and x_2 (Fig. 13.1d). For a design with N runs, the region is divided into N equally spaced rows and columns, creating N^2 cells. The points are then spread out, so that each row and column contains exactly one point. The main advantage of this method is that a latin hypercube design is computationally cheap to generate and that it can deal with a high number of parameters [6, 36]. The main disadvantage however, is that the design is not flexible with regard to adding or excluding runs. By changing the number of runs, the condition that each row and column contains exactly one point is no longer met. Furthermore, LHD is well suited for monotonic functions, but might not be adequate for other systems [6]. Finally, LHD designs are not necessarily space-filling (Fig. 13.1e). More elaborate algorithms which aim at ensuring the space-filling property of latin-hypercube designs, are described by Fang et al. [6].

Another method to achieve an effective spreading in space-filling designs is uniform sampling [5, 6]. In *uniform designs*, the parameters are spread out over space as uniformly as possible (Fig. 13.1f). The higher the number of runs, the better the spreading will be. Uniform designs are found to be efficient and robust, easy to understand and convenient, but computationally very demanding. Although this is a disadvantage, the fact that uniform designs cope well with the adding and removing of parameter combinations to the design makes them very useful in biomedical applications. For example, Carrier et al. [1] used a latin hypercube and uniform design to determine the most influential parameters of a calcium model that describes the effect of CaP biomaterials on the activity of osteogenic cells.

13.2.2 Methods for Analysing the Results of a Design

Once a suitable design has been set up and computer simulations are run with the different parameter combinations, the results have to be analysed. Depending on the design and the goal of the analysis several methods are available. Analysis of variance (ANOVA) is particularly suited for analysing the outcome of a (full or fractional) factorial design, giving an indication of the importance of the investigated parameters. For the more complex space-filling designs, Gaussian processes are more appropriate, as they not only determine the importance of a parameter but also give an estimate of the exact effect of varying a particular parameter on the outcome of the model. That way, more complex and non-linear effects can be revealed.

13.2.2.1 Analysis of Variance (ANOVA)

Analysis of variance (ANOVA) can be used to investigate the result of a full or fractional factorial design. Firstly, the total variation in the output is modelled by calculating the total sum of squares of the deviation about the mean (SS_T) [13].

$$SS_T = \sum_{i=1}^N [y_i - \bar{y}]^2 \quad (13.1)$$

In this equation, N is the number of runs, y_i the output for the i th run, and \bar{y} the overall mean of the output. The influence of one parameter is determined by SS_F :

$$SS_F = \sum_{i=1}^L N_{F,i} [\bar{y}_{F,i} - \bar{y}]^2, \quad (13.2)$$

where L is the number of levels used for each parameter, $N_{F,i}$ is the number of runs at each level of each factor and $\bar{y}_{F,i}$ is the mean output at each level of each factor. The percentage of the total sum of square,

$$\%TSS = [SS_F/SS_T] \times 100 \% \quad (13.3)$$

is a measure of importance for the parameter to the defined outcome [2].

13.2.2.2 Gaussian Process

Gaussian processes not only estimate the importance of individual parameters, but also the influence of the parameters on the outcome of a model. Given the output data $\mathbf{t}_N = \{t_i\}_{i=1}^N$ resulting from a combination $\mathbf{X}_N = \{\mathbf{x}_i\}_{i=1}^N$ of input parameters, determined in the set-up of the design, Gaussian processes are used to predict the

output t_* for a certain combination \mathbf{x}_* of input parameters [22]. To make this prediction, the output data are studied (Fig. 13.2), and a function $y(\mathbf{x})$ is searched, so that $y(\mathbf{x}_i)$ approaches the measured data t_i as closely as possible. In linear regression, the function $y(\mathbf{x})$ is assumed to be linear and usually least square methods are applied to find the most likely result for $y(\mathbf{x})$. This analysis method however implies that assumptions have to be made regarding the form of the function, prior to analysing the data.

A Gaussian process starts from the following posterior probability function:

$$P(y(\mathbf{x})|\mathbf{t}_N, \mathbf{X}_N) = \frac{P(\mathbf{t}_N|y(\mathbf{x}), \mathbf{X}_N)P(y(\mathbf{x}))}{P(\mathbf{t}_N|\mathbf{X}_N)}. \quad (13.4)$$

The first factor on the right hand side of (13.4), $P(\mathbf{t}_N|y(\mathbf{x}), \mathbf{X}_N)$, is the probability of the measured data given the function $y(\mathbf{x})$, and the second factor $P(y(\mathbf{x}))$ is the prior distribution on functions assumed by the model. In linear regression, this prior specifies the form of the function (e.g.: $y = ax + b$), and might put some restrictions on the parameters (e.g.: $a \neq 0$). The idea of Gaussian process modelling is to place a prior $P(y(\mathbf{x}))$ directly on the space of functions, without making assumptions on the form of the function. Just as a Gaussian distribution is fully defined by the mean and the covariance matrix, a Gaussian process is defined by a mean function and a covariance function. The mean is thus a function $\mu(\mathbf{x})$, which is often assumed to be the zero function, and the covariance is a function $\mathbf{C}(\mathbf{x}, \mathbf{x}')$. The only restriction on the covariance function is that it must be positive semi-definite. Several functions have been used widely, and proven valuable in literature. To get a better grasp on what exactly the covariance function represents and how a choice between the different available functions has to be made, an intuitive approach to developing the covariance function is explained below. A fully detailed and more theoretical approach can be found in Mackay [22].

Consider a system dependent upon one parameter x , with N parameter values x_i , and a parametrisation of $y(x)$ using a set of basis functions $\{\phi_h(x)\}_{h=1}^H$. The function




Fig. 13.2 Schematic representation of a Gaussian process on a system with output t , depending on one parameter x . The system is analysed for parameter values $\{x_1, x_2, \dots, x_N\}$, for which output values $\{t_1, t_2, \dots, t_N\}$ are obtained. In order to find the output value t^* resulting from parameter value x^* , the Gaussian process searches for a function $y(x)$ which can explain the output values $\{t_1, t_2, \dots, t_N\}$ the best. t^* is then found as $t^* = y(x^*)$

$y(x)$ can then be written as:

$$y(x, \mathbf{w}) = \sum_{h=1}^H w_h \phi_h(x). \tag{13.5}$$

As basis functions $\{\phi_h(x)\}$, radial basis functions centred at fixed points $\{c_h\}$ are chosen (Fig. 13.3).

$$\phi_h(x) = e^{-\frac{[x-c_h]^2}{2r^2}} \tag{13.6}$$

Using the input points $\{x_i\}$ and the H basis functions ϕ_h , an $N \times H$ matrix \mathbf{R} can be defined.

$$R_{ih} = \phi_h(x_i) \tag{13.7}$$

For a certain set of parameters \mathbf{w} , the function $y(x)$ then has the values $\mathbf{y}_N = \{y_i\}$ at the input points x_i .

$$y_i = \sum_{h=1}^H w_h \phi_h(x_i) = \sum_{h=1}^H w_h R_{ih} \tag{13.8}$$

In parametric regression methods, it is normally assumed that the prior distribution of the parameters \mathbf{w} is Gaussian with a zero mean.

$$P(\mathbf{w}) \sim \mathcal{N}(0, \sigma_w^2 \mathbf{I}) \tag{13.9}$$




Fig. 13.3 Radial basis functions centred at $c_h = 10$, used for the parametrisation of $y(x)$. The parameter r is a length scale defining the width of the basis function. A larger value of r results in a wider basis function and a smoother approximation of the function $y(x)$

In that case, \mathbf{y}_N , being a linear function of \mathbf{w} is also Gaussian distributed with mean zero and covariance matrix \mathbf{Q} .

$$\mathbf{Q} = \left\langle \mathbf{y}_N \mathbf{y}_N^T \right\rangle = \sigma_w^2 \mathbf{R} \mathbf{R}^T \quad (13.10)$$

If each data point t_n is assumed to differ by additive Gaussian noise of variance σ_v^2 from the corresponding function value $y(x_i)$, then:

$$\begin{aligned} P(\mathbf{t}) &\sim \mathcal{N}(0, \mathbf{Q} + \sigma_v^2 \mathbf{I}) \\ &\sim \mathcal{N}(0, \mathbf{C}). \end{aligned} \quad (13.11)$$

Now the assumption is made that the radial basis functions are uniformly spaced, that $H \rightarrow \infty$ and that $\sigma_w^2 = S/(\Delta H)$, where ΔH is the number of base functions per unit length of the x -axis. The summation over h then becomes an integral and the (i, j) entry of \mathbf{Q} equals:

$$Q_{ij} = \sqrt{\pi r^2} S e^{-\frac{|x_j - x_i|^2}{4r^2}}. \quad (13.12)$$

The covariance function $\mathbf{C}(\mathbf{x}, \mathbf{x}')$ of the Gaussian process is thus related to the basis functions chosen in the model. The parameter r is a length parameter describing the width of the basis function. For a high value of r , the basis functions are wider, implying a higher correlation of the values of $y(x)$ at input points x_i and x_j , resulting in a smoother function.

As mentioned before, several forms of the covariance function are possible. The first one used is the Gaussian or squared exponential covariance function.

$$Q_{ij} = \sigma^2 e^{-\sum_{m=1}^M \theta_m [x_{im} - x_{jm}]^2} = \sigma^2 \prod_{m=1}^M e^{\theta_m [x_{im} - x_{jm}]^2} \quad (13.13)$$

The summation in (13.13) is a result of the M -dimensional nature of the parameter combinations, which has not been taken into account in the intuitive approach, but is reintroduced here. The parameter $1/\theta_m$, corresponding to parameter m , is related to the length scale r described above. A very large number of θ_m implies a short length scale, indicating the function value will change significantly when changing the parameter. A value of $\theta_m = 0$ implies an infinite length scale, meaning y is a constant function of that input.

The second form for the covariance function is the cubic correlation covariance function. The covariance matrix is composed of the following elements:

$$Q_{ij} = \sigma^2 \prod_{m=1}^M \rho(d_{ijm} \theta_m), \quad (13.14)$$

where

$$d_{ijm} = x_{im} - x_{jm} \tag{13.15}$$

$$\rho(d\theta) = \begin{cases} 1 - 6[d\theta]^2 + 6[|d|\theta]^3, & |d| \leq \frac{1}{2\theta} \\ 2[1 - |d|\theta]^3, & \frac{1}{2\theta} \leq |d| \leq \frac{1}{\theta} \\ 0, & \frac{1}{\theta} \leq |d|. \end{cases} \tag{13.16}$$

Equation (13.16) shows that for a certain distance d_{ijm} between the points x_{im} and x_{jm} the covariance function becomes zero, meaning that the values of the output at these locations are not correlated. This allows the cubic covariance function to capture variations on smaller length scales and include outliers in the model.

The regression process now consists of selecting the parameters $\{\theta_m, \sigma^2, \sigma_v^2\}$. This is done by maximizing the following equation:

$$\mathcal{M} = \log p(\mathbf{t}_N | \mathbf{X}_N, \{\theta_m, \sigma^2, \sigma_v^2\}) = -\frac{1}{2} \mathbf{t}_N^T \mathbf{C}^{-1} \mathbf{t}_N - \frac{1}{2} \log |\mathbf{C}| - \frac{n}{2} \log 2\pi \tag{13.17}$$

It should be noted that Gaussian processes are not designed to actually find a specific function $y(\mathbf{x})$ that fits the data best. They are designed to make predictions of the outcome for a new input point \mathbf{x}_* , without specifying the actual function $y(\mathbf{x})$ that was used to reach that conclusion. For every input \mathbf{x}_* , the result of the Gaussian process will be a mean \bar{t}_* and a variance $\sigma_{\bar{t}_*}^2$ on the prediction. Using

$$\mathbf{C}_* = [C(x_*, x_1) \ C(x_*, x_2) \ \dots \ C(x_*, x_N)] \quad \text{and} \quad \mathbf{C}_{**} = C(x_*, x_*), \tag{13.18}$$

the mean and variance of t_* become [3]:

$$\begin{aligned} P(t_* | \mathbf{t}_N) &\sim \mathcal{N}(\bar{t}_*, \sigma_{\bar{t}_*}^2) \\ &\sim \mathcal{N}(\mathbf{C}_* \mathbf{C}^{-1} \mathbf{t}_N, \mathbf{C}_{**} - \mathbf{C}_* \mathbf{C}^{-1} \mathbf{C}_*^T). \end{aligned} \tag{13.19}$$

The results of a Gaussian process regression are usually visualised by plotting the mean \bar{t}_* as a function of the different inputs x_m and adding a confidence interval to the graph, calculated by using the variance $\sigma_{\bar{t}_*}^2$ (Fig. 13.4).

13.2.3 Interpretation of the Results

After statistically processing the simulation outcomes with ANOVA or Gaussian processes, the most influential parameters and possible interactions between parameters can be determined. Moreover, a biological interpretation of the predictions might lead to a greater understanding of the modelled processes at hand. Remark, however, that the results of the sensitivity analysis are valid within the chosen parameter as well as response space. Indeed, if the parameter ranges or the responses




Fig. 13.4 Visualisation of the results of a Gaussian process. The *horizontal axis* shows the parameter value, the *vertical axis* shows the output value predicted by the Gaussian process. For every parameter value, the predicted output has a mean (*solid line*) and a confidence interval (*dashed line*)

would be altered, different results could be obtained. Moreover, DOE is a statistical tool implying that a larger number of runs will provide the statistical model with more data thereby allowing more accurate predictions. It is instructive to test the statistical predictions by running the original model with the according parameter values in order to test whether the examined region was correctly sampled. Finally, every design and analysis method has its pros and cons. If little is known about possible non-linearities, interactions and monotonicity of the model being studied, it might be informative to compare the results of different designs and analysis methods.

13.3 Application to a Mechanobiological Model of Tooth Movement

This section demonstrates the techniques described previously by applying them to a mechanobiological model describing tooth movement. Progress in medicine and higher expectation of quality of life have led to a higher demand for several dental and medical treatments [33], making it more and more common for other medical conditions needing to be taken into account by the orthodontist when planning orthodontic treatment. During treatment, tooth displacement is achieved by applying orthodontic forces to the tooth. Under the influence of these forces, the pressure side of the tooth root will experience bone resorption while bone formation will take place on the tension side. The coordination of these two processes through cellular communication results in permanent tooth displacement through the alveolar bone. Together with experiments, computer models might lead to a better understanding of orthodontic treatment and the pathologies affecting the outcome. The model analysed by DOE in this section is a mechanobiological model using partial differential equations to describe cell densities, growth factor concentrations and matrix densities occurring during orthodontic tooth movement and was presented previously by the authors [37].

13.3.1 Model Equations

The mechanobiological model consists of a set of nine coupled non-linear partial differential equations, of the taxis-diffusion-reaction (TDR) type. The equations describe the concentration of various cells, growth factors, cytokines and matrix-components. The periodontal ligament (PDL) consists of collagen fibres (m_f) and contains a large amount of fibroblasts (c_f). The alveolar bone consists of mineralized collagen, with m_m representing the degree of mineralization of the collagen. The bone has a small concentration of osteoblasts (c_b) and osteoclasts (c_l), constantly remodelling and renewing the bone. To coordinate bone remodelling, osteoclasts, osteoblasts and fibroblasts communicate through the RANKL-RANK-OPG signalling pathway. In the model, RANKL ($g_r = g_{rb} + g_{rf}$) is produced by fibroblasts (g_{rf}) and osteoblasts (g_{rb}), while OPG (g_o) is produced by osteoblasts only. The osteogenic differentiation of mesenchymal stem cells into osteoblasts is regulated by active TGF- β (g_b), also produced by osteoblasts and fibroblasts. Multinucleated osteoclasts are formed through the fusion of hematopoietic stem cells, which are present in the vascular matrix in the PDL and the bone.

Fibroblasts are modelled to respond to mechanical stretching by producing the osteogenic growth factor TGF- β , along with other osteogenic factors of the TGF- β superfamily [15, 25, 30, 39]. The upregulation of the TGF- β production results in the appearance of a high number of osteoblasts in and around the PDL. This leads to bone formation in the tension zones. Fibroblasts respond to compression by upregulating the production of RANKL [14, 18, 28, 40]. This results in a higher number of osteoclasts, which start resorbing the alveolar bone, making it possible for the tooth to move. A schematic overview of the processes captured by the model can be found in Fig. 13.5. More information concerning the biological assumptions made in this model can be found in Van Schepdael et al. [37, 38] and a more comprehensive overview of the biology of tooth movement can be found in Garant [8], Krishnan and Davidovitch [17, 18] and Henneman et al. [11].

The specific equations for all nine variables are represented below. More information on the parameters, equations and initial conditions can be found in Van Schepdael et al. [37]. An overview of the origin and value of all parameters can be found in Table 13.1.

$$\frac{\partial m_c}{\partial t} = \underbrace{P_{cs}[1 - \kappa_c m_c]c_b}_{\text{production by osteoblasts}} + \underbrace{P_{csf}[1 - \kappa_{cf} m_c]c_f}_{\text{production by fibroblasts}} \quad (13.20)$$

$$\frac{\partial m_m}{\partial t} = \underbrace{P_{ms}[1 - m_m]c_b}_{\text{mineralisation by osteoblasts}} - \underbrace{Q_{md}c_l H(m_m)}_{\text{demineralsation by osteoclasts}} \quad (13.21)$$

$$\frac{\partial c_b}{\partial t} = \underbrace{\frac{Y_{11}g_b}{H_{11} + g_b}[1 - m_m]H(\bar{m}_m - m_{bt})}_{\text{differentiation from MSCs}} + \underbrace{A_{b0}m_m c_b [1 - \alpha_b c_b]}_{\text{proliferation}} - \underbrace{d_b c_b}_{\text{apoptosis}} \quad (13.22)$$

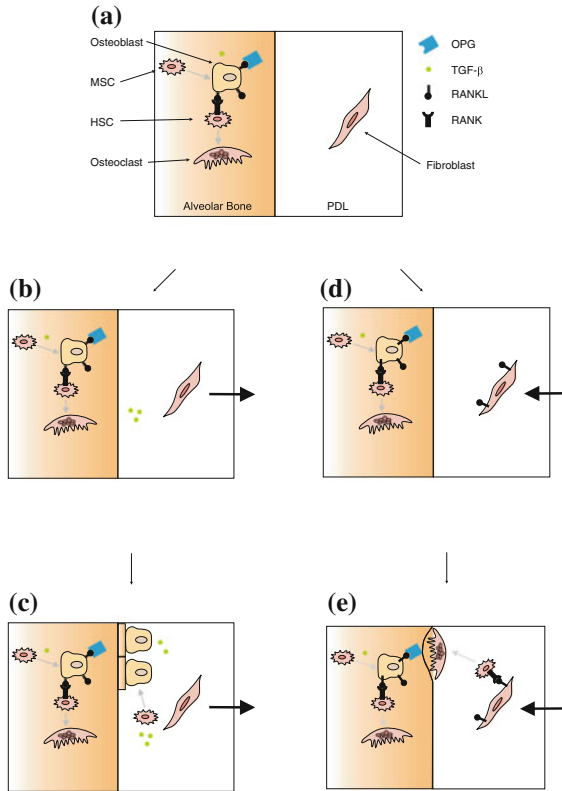


Fig. 13.5 Schematic overview of the most important cells and signalling pathways involved in orthodontic tooth movement. **a** In homeostasis, the fibroblast is the most abundant cell in the PDL, while osteoclasts and osteoblasts are present in the alveolar bone. Osteoblasts are derived from MSC's, which differentiate under the influence of TGF- β . Osteoclasts are derived from HSC's, which express RANK on their membranes. When RANK binds to RANKL, HSC's are stimulated to fuse into multinucleated osteoclasts. OPG is a soluble decoy operator that also binds to RANKL, thus preventing the formation of osteoclasts. **b, c** In response to tension, fibroblasts express TGF- β . As a response, MSC's differentiate into osteoblasts, which start forming new bone. **d, e** In response to compression, fibroblasts express RANKL, stimulating osteoclast formation. Osteoclasts attach to the bone surface and bone resorption starts

$$\frac{\partial c_l}{\partial t} = \underbrace{Y_2 g_r}_{\text{fusion from HSCs}} - \underbrace{[D_2 + H_2 g_b] d_{l0} c_l}_{\text{apoptosis}} - \underbrace{\nabla \cdot [C_{mh} c_l \nabla m_m]}_{\text{attachment to bone matrix}} \quad (13.23)$$

$$\frac{\partial c_f}{\partial t} = \underbrace{A_{f0} [1 + \alpha_{fs} |S|] m_c c_f [1 - \alpha_f c_f]}_{\text{proliferation}} - \underbrace{d_f m_m c_f}_{\text{apoptosis}} + \underbrace{\Phi \nabla \cdot [D_f \nabla c_f]}_{\text{diffusion}} \quad (13.24)$$

$$\frac{\partial g_b}{\partial t} = \underbrace{G_{gb} [1 - \alpha_g g_b] c_b}_{\text{production by osteoblasts}} + \underbrace{E_{gb} [S \cdot H(S)] c_f}_{\text{production by fibroblasts}} - \underbrace{d_{gb} g_b}_{\text{denaturation}} + \underbrace{\nabla \cdot [D_{gb} \nabla g_b]}_{\text{binding to OPG}} \quad (13.25)$$

Table 13.1 Overview of the parameters of the mechanobiological model, their value, unit and origin

Parameter	Value	Unit	Origin
P_{ms}	3.42×10^{-5}	ml cells ⁻¹ day ⁻¹	From steady state conditions
Q_{md}	3.6×10^{-5}	ml cells ⁻¹ day ⁻¹	(1)
P_{cs}	2×10^{-7}	g cells ⁻¹ day ⁻¹	(1)
κ_c	13.55	ml g ⁻¹	(1)
P_{csf}	2×10^{-8}	g cells ⁻¹ day ⁻¹	(1)
κ_{cf}	10	ml g ⁻¹	(1)
A_{b0}	0.54	day ⁻¹	(1) and stability analysis
α_b	2×10^{-5}	ml cells ⁻¹	(1) and stability analysis
d_b	0.18	day ⁻¹	(1)
Y_{11}	3.27×10^8	cells ml ⁻¹ day ⁻¹	(1)
H_{11}	10	ng ml ⁻¹	(1)
m_{bt}	0.3	[-]	Estimated
C_{mh}	3.06×10^{-2}	mm ² day ⁻¹	Estimated
Y_2	551.6	cells ng ⁻¹ day ⁻¹	(2)
d_{l0}	0.7	day ⁻¹	Using life span of osteoclast
D_f	0.25	mm ² day ⁻¹	(1)
A_{f0}	1.06	ml g ⁻¹ day ⁻¹	(1)
A_{fs}	10	[-]	Estimated
α_f	1×10^{-6}	ml cells ⁻¹	(1)
d_f	0.11	day ⁻¹	From steady state conditions
D_{gb}	6.13×10^{-2}	mm ² day ⁻¹	(1)
G_{gb}	6.03×10^{-5}	ng cells ⁻¹ day ⁻¹	(1)
α_g	0.1	ml ng ⁻¹	(3), (4)
E_{gb}	1×10^{-4}	ng cells ⁻¹ day ⁻¹	Estimated
d_{gb}	100	day ⁻¹	(1)
P_{rs}	3440	ng ml ⁻¹ day ⁻¹	(2)
R_1	9.15×10^{-5}	ng cells ⁻¹	(2)
d_{gr}	10.05	day ⁻¹	(2)
B_{1r}	2.5×10^{-3}	ml ng ⁻¹ day ⁻¹	(2)
B_{1o}	1.67×10^{-3}	ml ng ⁻¹ day ⁻¹	(2)
E_{grf}	1×10^{-3}	ng cells ⁻¹ day ⁻¹	Estimated

(continued)

Table 13.1 (continued)

Parameter	Value	Unit	Origin
D_{go}	4.58×10^{-2}	$\text{mm}^2 \text{day}^{-1}$	Using molecular weight of OPG
P_{os}	6.83×10^{-3}	$\text{ng cells}^{-1} \text{day}^{-1}$	(2)
κ_o	8.3×10^{-8}	ml ng^{-1}	(2)
d_{go}	35	day^{-1}	(2)
D_2	248.5	[-]	From H_2
H_2	48.6	ml ng^{-1}	(2)

(1) Derived from Geris et al. [10]. (2) Derived from Pivonka et al. [31]. (3) Derived from Pfeilschifter et al. [29] (4) Derived from Sandberg et al. [35]

$$\frac{\partial g_{rb}}{\partial t} = \underbrace{P_{rs} \left[1 - \frac{g_{rb}}{R_1 c_b} \right]}_{\text{production by fibroblasts}} - \underbrace{d_{gr} g_{rb}}_{\text{denaturation}} - \underbrace{B_{1r} g_{rb} g_o}_{\text{binding to OPG}} \quad (13.26)$$

$$\begin{aligned} \frac{\partial g_{rf}}{\partial t} = & \underbrace{E_{grf} [[H(S) - 1] S] c_f}_{\text{production by fibroblasts}} - \underbrace{d_{gr} g_{rf}}_{\text{denaturation}} - \underbrace{B_{1r} g_{rf} g_o}_{\text{binding to OPG}} \\ & + \underbrace{\nabla \cdot \left[\frac{D_f}{c_f} g_{rf} \nabla c_f \right]}_{\text{moving with fibroblasts}} - \underbrace{d_{fm} g_{rf}}_{\text{apoptosis of fibroblasts}} \end{aligned} \quad (13.27)$$

$$\frac{\partial g_o}{\partial t} = \underbrace{P_{os} [1 - \kappa_o g_o] c_b}_{\text{production by osteoblasts}} - \underbrace{d_{go} g_o}_{\text{denaturation}} - \underbrace{B_{1o} g_r g_o}_{\text{binding to RANKL}} + \underbrace{\nabla \cdot [D_{go} \nabla g_o]}_{\text{diffusion}} \quad (13.28)$$

The numerical simulations were performed on a domain that consists of two rectangular parts that represent small sections of the tooth root as shown in Fig. 13.6, and are located about halfway between the tooth crown and the tooth apex.

Table 13.2 shows the initial values of all variables in the PDL and the bone. To prevent the appearance of numerical instabilities, continuous initial conditions were used to model the boundary between PDL and alveolar bone. To simulate an abrupt, but continuous, transition from PDL to alveolar bone, the following function was used.

$$c_i(t=0) = \frac{c_{i,bone}^0 - c_{i,PDL}^0}{\pi} \arctan(D \cdot [x - x_S]) + \frac{c_{i,bone}^0 - c_{i,PDL}^0}{2} + c_{i,PDL}^0 \quad (13.29)$$

The parameter D is dependent upon the desired steepness of the function, x_S is the x -coordinate of the boundary between PDL and alveolar bone, and $c_{i,bone}^0$ and $c_{i,PDL}^0$ are the initial conditions of the variable in the alveolar bone and the PDL.

The model parameters and variables were non-dimensionalised for the numerical calculations. A typical length scale during orthodontic tooth movement is the thickness of the periodontal ligament, $L_0 = 0.2 \text{ mm}$ [32] and a typical time scale of $T_0 = 1 \text{ day}$ was chosen. A representative concentration of collagen content in the




Fig. 13.6 Schematic representation of the model domain. **a** The model domain consists of two rectangular parts that are located on the *left* and the *right side* of the tooth root. For the simulations, a horizontal translation \mathbf{u} is applied to the root. The *left part* of the domain will thus experience bone formation, the *right part* experiences bone resorption. **b** Detail of the formation side of the model domain at the start of force application. **c** Detail of the formation side of the model domain at the end of the simulation. Due to bone formation, more alveolar bone is now present, and the PDL has shifted to the *right*

Table 13.2 Initial conditions applied to the model domain. The initial values and units of all variables in the PDL, the alveolar bone and the root are shown

	PDL	Alveolar bone	Root	Unit
m_m	0	0.9	0	[-]
m_c	1	0.075	0	g ml^{-1}
c_b	0	3.2×10^4	0	cells ml^{-1}
c_l	0	2.3×10^3	0	cells ml^{-1}
c_f	1×10^6	0	0	cells ml^{-1}
g_b	0	2×10^{-2}	0	ng ml^{-1}
g_{rb}	0	2.9	0	ng ml^{-1}
g_{rf}	0	0	0	ng ml^{-1}
g_o	0	6.2	0	ng ml^{-1}

tissue is $m_0 = 0.1 \text{ g/ml}$ [10]. Typical growth factor concentrations are in the order of magnitude of $g_0 = 100 \text{ ng/ml}$, and a non-dimensionalisation value of $c_0 = 10^6 \text{ cells/ml}$ was used for the cell densities [10]. All results and parameter values are presented in their undimensionalised value, unless mentioned otherwise.

13.3.2 Model Analysis

The complexity of the biological processes occurring during orthodontic tooth movement results in a high number of parameters in the mechanobiological model. When performing numerical analysis using such models, it is important to study the behaviour of the model, and its parameters, more closely. That way, stability, numerical accuracy and convergence can be ensured (see also [16]).

The model described in the previous section depends upon 37 parameters, although not all of these are independent. To make good predictions, all the parameters have to be determined as accurately as possible, which is not an easy task (see also [24]). Some of the parameters do not represent physical processes and cannot be measured. Others represent a combination of several processes, making it more difficult to measure them experimentally. Some of the parameter values are based on experimental values, but although some experiments were conducted *in vivo*, most of the estimations are based on *in vitro* experimental results.

A sensitivity analysis can be used to determine the importance of the parameters, as well as important interactions between them. Secondly, a sensitivity analysis can also be used to simplify the model by eliminating insignificant model parameters (see also [4]). Finally, if unexpected behaviour of the model surfaces during the sensitivity analysis, this behaviour can be studied more closely and corrected if necessary.

Since a space-filling design for the full model would be computationally very demanding and very difficult to interpret, a fractional factorial design was performed first to estimate the importance of the parameters. The results were analysed using ANOVA analysis.

The most important parameters were then studied more closely using space-filling designs. As will be explained in Sect. 13.3.2.1, some of the parameter combinations will create instabilities in the mechanobiological model or predict physically impossible situations and have to be excluded from the analysis. Furthermore, the tooth model is highly non-linear. Therefore, latin hypercube designs are probably not the best choice for the analysis, while the fact that uniform designs cope well with the adding and removing of parameter combinations makes them a very good choice to examine the tooth model.

Three different uniform designs were performed: one containing the ten most important parameters, one containing the 12 parameters that were most important to the equilibrium conditions in the PDL and the bone, and one containing the 12 parameters that had the highest influence on bone formation and resorption. The results of all uniform designs were analysed using a Gaussian process with a Gaussian and a cubic covariance function. The designs were generated and analysed using JMP (SAS Institute Inc., Cary, North Carolina, USA).

13.3.2.1 Sensitivity Analysis Using a Fractional Factorial Design

A two-level fractional factorial design was used to identify the most important parameters of the model. Of the 37 parameters of the model, 28 were studied, as shown in Table 13.3. All parameters describing matrix densities at which matrix productions stops (κ_c, κ_{cf}), growth factor concentrations at which growth factor production stops (α_g, κ_o) or cell concentrations at which cell proliferation is halted (α_b, α_f), were excluded from the analysis, since they do not have a specific physiological meaning and influence the model in a very predictable way [13]. The parameter B_{1o} , describing the reaction rate of OPG and RANKL, is related to B_{1r} and varied accordingly.

The parameters were investigated at two levels, a high and a low level (see Table 13.3), using a resolution IV array with 64 runs. The use of a resolution IV array guarantees that the main effects will not be confounded with two-factor interactions, while limiting the number of runs required for the analysis. For each parameter combination, the model was run three times, once in steady state conditions, once modelling bone formation and once modelling bone resorption. To assess the results obtained from the study, several output variables were studied.

During steady state, the concentrations of nine variables in both the alveolar bone and the PDL were measured. To assess the ability of the model to represent a meaningful situation, a variable *Real* was introduced, shown by (13.30). All 18 variables (nine in the alveolar bone, nine in the PDL), were assigned a boolean parameter γ_i which is zero or one. The value $\gamma_i = 1$ was given when a predicted concentration approaches the expected concentration (e.g. $m_m(\text{bone}) = 0.7$), $\gamma_i = 0$ was assigned to a predicted concentration which is not realistic (e.g. $m_m(\text{bone}) = 0$). Secondly, when one of the steady states has a positive eigenvalue, the boolean parameter γ was set to $\gamma = 0$, for two stable steady states, γ was equal to one.

$$Real = \frac{\gamma + \sum_{i=1}^{18} \gamma_i}{19} \quad (13.30)$$

A typical parameter combination for which the value of *Real* is small, is one for which the predicted mineralization of the alveolar bone is equal to zero. The same is true for parameter combinations for which one of the steady states is unstable or for which the mineralization of the PDL is very high.

Next, the ability of the model to combine the alveolar bone and the PDL was also studied. This was done by assessing the variable v_{SS} which is the movement of the boundary between bone and PDL. When $v_{SS} \neq 0$, the boundary between bone and ligament will move, even when no orthodontic force is applied. In that case, it is not possible to model co-existence of the alveolar bone and the PDL, which is not a physiological situation.

During bone formation, three variables were recorded: the speed of tooth movement (v_F), the mineralization of the newly formed bone (m_{mnew}) and the concentration of TGF- β in the PDL (g_{bF}). During bone resorption, the speed of tooth movement (v_R) and the concentration of RANKL in the PDL (g_{rR}) were monitored.

Table 13.3 Non-dimensionalised parameter ranges used in the different designs of the sensitivity analysis

Parameter	FF		U		USS		UM	
	Low	High	Low	High	Low	High	Low	High
P_{ms}	15	60	15	60	–	–	15	60
Q_{md}	18	72	–	–	18	72	18	72
P_{cs}	1	4	–	–	–	–	–	–
P_{csf}	0.1	0.4	–	–	–	–	0.1	0.4
A_{b0}	0.25	1	–	–	–	–	0.25	1
d_b	0.07	0.3	0.07	0.2	0.07	0.3	0.07	0.25
Y_{11}	180	600	180	600	180	600	180	600
H_{11}	0.05	0.2	0.05	0.2	0.05	0.2	–	–
m_{bt}	0.15	0.6	0.15	0.6	–	–	0.15	0.4
C_{mh}	0.3	1.4	0.5	1.4	0.3	1.4	–	–
Y_2	0.02	0.1	–	–	0.02	0.1	–	–
d_{l0}	0.35	1.5	–	–	–	–	–	–
D_f	3	9	–	–	–	–	–	–
A_{f0}	0.05	0.2	–	–	0.05	0.2	–	–
A_{fs}	5	20	–	–	5	20	–	–
d_f	0.05	0.2	–	–	0.05	0.2	–	–
D_{gb}	0.6	3	–	–	–	–	–	–
G_{gb}	0.3	1.2	0.5	1.2	0.3	1.2	–	–
E_{gb}	0.5	2	–	–	–	–	0.5	2
d_{gb}	50	200	50	200	50	200	50	150
P_{rs}	15	70	–	–	–	–	–	–
R_1	0.5	2	–	–	0.5	2	–	–
d_{gr}	5	20	5	20	–	–	5	20
B_{1r}	0.07	0.3	–	–	–	–	–	–
E_{grf}	5	20	–	–	–	–	5	20
D_{go}	0.5	2	–	–	–	–	–	–
P_{os}	40	120	–	–	–	–	–	–
d_{go}	20	70	–	–	–	–	20	70

For the fractional factorial design, the parameters are varied between a low and a high value, as can be seen in the first column. For the other designs, several parameters were selected and varied over approximately the same range. Highlighted numbers indicate ranges that have been corrected in order to avoid non-physiological situations. FF: parameter ranges for the fractional factorial design of Sect. 13.3.2.1. U parameter ranges for the uniform design of Sect. 13.3.2.2. USS: parameter ranges for the uniform design on the equilibrium concentrations (Sect. 13.3.2.3). UM: parameter ranges for the uniform design on bone formation and resorption (Sect. 13.3.2.4)

Steady State

The data were investigated using analysis of variance (ANOVA) and the percentage of the total sum of square, %TSS, for each parameter and output variable can be found

Table 13.4 Results of the fractional factorial design concerning the parameters *Real* and v_{SS}

	Real	v_{SS}		Real	v_{SS}
A_{b0}	8.40	0.00	E_{gb}	1.27	0.00
A_{f0}	0.03	0.00	E_{grf}	0.03	0.71
A_{fs}	0.00	0.71	G_{gb}	4.38	5.87
B_{1r}	0.65	0.00	H_{11}	0.03	5.91
C_{mh}	3.14	0.00	m_{bt}	0.03	5.87
d_b	24.93	0.00	P_{cs}	0.93	0.00
D_f	1.27	0.72	P_{csf}	0.03	0.70
d_f	2.59	0.00	P_{ms}	1.27	0.74
d_{gb}	2.59	5.99	P_{os}	1.66	0.00
d_{go}	0.93	5.91	P_{rs}	0.00	0.00
d_{gr}	0.23	5.91	Q_{md}	4.38	0.00
D_{gb}	1.66	0.00	R_1	0.93	5.95
D_{go}	0.10	0.70	Y_2	3.14	0.00
d_{l0}	0.10	0.00	Y_{11}	0.10	5.87

The percentage of total sum of squares (%TSS) is listed for each parameter and output variable

in Tables 13.4 and 13.5. The parameters that most influenced the alveolar bone were mostly related to the concentration of TGF- β . Both the production rate G_{gb} and the degradation rate d_{gb} have a high influence on many output variables. The same is true for the parameter H_{11} , related to the sensitivity of MSC's to TGF- β . The rate at which osteoclasts demineralize bone (Q_{md}) and the haptotactic parameter C_{mh} also seemed to influence the alveolar bone.

The periodontal ligament was mainly influenced by the proliferation and apoptosis rate of the fibroblasts (A_{f0} and d_f), the fusion and apoptosis of osteoclasts (Y_2 and d_{l0}) and the production of TGF- β (G_{gb}). The artificial variable *Real* was most influenced by parameters A_{b0} and d_b , describing the proliferation and apoptosis of osteoblasts. The influence of the parameters on v_{SS} was limited.

Out of the 64 parameter combinations chosen for the screening design, 17 gave results where the mineralization m_m in the alveolar bone was smaller than 0.1. In that case, the modelled situation is assumed to be non-physiological, and no results were obtained for bone formation and bone resorption. To prevent this in future analyses, the parameter combinations in which this problem occurred were examined more closely. Without making definite conclusions, the problem seemed to occur more frequently for high values of Q_{md} , d_b and Y_2 and for low values of C_{mh} and G_{gb} . In future analyses care must be taken not to allow parameter combinations that promote bone resorption too strongly, especially when the object is to study bone formation or resorption.

Table 13.5 (continued)

	Osteoblasts		Fibroblasts		Osteoclasts		TGF- β		OPG		Ob-RANKL		F-RANKL		Collagen		Mineralization	
	Bone	PDL	Bone	PDL	Bone	PDL	Bone	PDL	Bone	PDL	Bone	PDL	Bone	PDL	Bone	PDL	Bone	PDL
P_{os}	0.10	0.00	0.46	1.68	0.00	0.00	0.03	0.00	3.42	3.70	0.10	0.00	0.00	0.00	0.55	2.80	0.88	2.58
P_{rs}	0.11	0.00	0.64	1.86	0.42	0.00	0.13	0.00	1.26	0.00	0.26	0.00	0.00	0.00	0.79	1.39	0.55	2.64
Q_{md}	0.28	0.00	9.12	5.54	1.13	0.00	0.11	0.00	0.01	0.00	0.06	0.00	0.00	0.00	7.90	0.02	9.71	0.13
R_1	0.01	0.00	3.72	0.01	11.36	0.00	0.05	0.00	2.12	0.00	7.71	0.00	0.00	0.00	4.06	0.51	4.67	0.24
Y_2	0.00	14.29	5.55	3.10	7.92	8.33	0.06	0.00	0.50	6.17	0.11	0.00	0.00	0.00	4.40	5.18	6.67	4.42
Y_{11}	10.10	0.00	0.07	0.11	0.10	0.00	9.96	0.00	4.99	2.47	7.06	0.00	0.00	0.00	0.79	0.24	0.04	1.58

The percentage of total sum of squares (%TSS) is listed for each parameter and output variable. The most influential parameter for each output variable is highlighted. Note that Ob-RANKL indicates the RANKL bound on osteoblast membranes (g_{rb}) and F-RANKL represents RANKL bound on fibroblast membranes (g_{rf})

Bone Formation and Bone Resorption

The same process was applied to analyse the data concerning bone formation and bone resorption. The percentage of the total sum of squares of the relevant output variables can be found in Table 13.6. As expected, the speed of bone formation was mostly influenced by the parameter m_{br} , determining when a mesenchymal stem cell is close enough to the bone to differentiate into an osteoblast. The mineralization rate of collagen by osteoblasts also influenced the speed of bone formation. The

Table 13.6 Results of the fractional factorial design concerning bone formation and bone resorption during orthodontic tooth movement

	Formation			Resorption	
	v_F	gbF	m_{mnew}	v_R	g_rR
A_{b0}	0.00	1.03	7.77	4.51	3.96
A_{f0}	0.01	0.05	0.25	0.70	0.03
A_{fs}	0.33	0.53	0.45	0.42	3.11
B_{1r}	0.00	0.09	1.51	0.56	0.05
C_{mh}	0.43	0.12	0.47	0.46	1.87
d_b	0.06	0.02	4.30	10.08	3.20
D_f	0.66	0.13	0.08	0.07	0.21
d_f	0.73	0.29	1.15	3.47	0.10
d_{gb}	0.58	37.78	9.93	7.04	0.05
d_{go}	0.55	1.44	0.33	0.79	5.85
d_{gr}	0.16	0.06	0.38	10.07	28.37
D_{gb}	0.11	0.30	0.11	0.13	0.45
D_{go}	1.10	0.16	4.21	1.42	0.72
d_{l0}	0.10	0.03	0.15	0.79	0.08
E_{gb}	0.10	42.07	0.02	0.56	0.01
E_{grf}	2.17	0.05	0.38	5.94	23.76
G_{gb}	0.51	0.20	1.35	3.57	0.20
H_{11}	0.41	0.63	0.64	2.97	0.80
m_{bt}	63.24	0.23	36.30	1.38	3.92
P_{cs}	0.00	0.15	0.27	0.04	0.01
P_{csf}	0.48	0.17	0.04	5.60	1.11
P_{ms}	8.14	0.00	4.14	1.85	2.01
P_{os}	0.20	0.02	1.29	0.10	1.80
P_{rs}	0.09	0.10	0.71	3.33	0.99
Q_{md}	0.68	0.33	6.41	3.17	0.13
R_1	0.33	0.43	1.87	0.02	0.22
Y_2	0.67	0.07	1.42	1.13	0.01
Y_{11}	1.08	2.30	1.84	5.66	0.01

The percentage of total sum of squares (%TSS) is listed for each parameter and output variable. The two most influential parameters for each output variable are highlighted

concentration of TGF- β in the PDL during tooth movement was mainly influenced by the production of TGF- β by fibroblasts (E_{gb}) and the degradation rate d_{gb} . This degradation rate also influenced the mineralization of the newly formed bone, just like the parameter m_{br} and the proliferation rate A_{b0} of the alveolar bone.

The speed of bone resorption was mainly influenced by the degradation rate d_b of osteoblasts and the denaturation rate d_{gr} of RANKL, but many other parameters also had an influence on bone resorption. The amount of RANKL in the PDL was mainly influenced by the denaturation rate d_{gr} of RANKL and the production rate of RANKL by fibroblasts (E_{grf}).

Preliminary Conclusions

Although the ANOVA process is able to identify the most important parameters in the process, it gives no insight into the effect of these parameters on the outcome. Furthermore, although some of the results are very straightforward, others need some more explaining. The effect of TGF- β on alveolar bone and fibroblasts on the PDL is quite clear, but the influence of the parameter C_{mh} on the alveolar bone is more puzzling. In the following sections, the results of the fractional factorial design will be used to set up several uniform designs, analysed using Gaussian processes, that study the influence of these parameters more closely. That way, the exact effect of the parameters can be investigated, and the interaction between parameters can be studied.

13.3.2.2 Sensitivity Analysis Using a Uniform Design

To solve the questions that arose in the previous section, it was decided to look more closely at the ten most influential parameters. Those parameters were chosen by selecting every parameter for which $\%TSS > 10\%$. As the high influence of the artificial parameter m_{br} on the speed of bone formation dominated the results, the parameter P_{ms} was added to this list. From this list, those parameters who influenced several output variables or output variables with a large variation, represented by a high value of SS_T , were chosen. The parameter ranges of d_b , G_{gb} and C_{mh} were limited to avoid non-physiological situations in which the mineralization of the alveolar bone was too low (Table 13.3).

The influence of the ten parameters was investigated as previously explained using a uniform design with 100 runs. 100 parameter combinations were generated, evenly spread out over the ten-dimensional parameter space. The uniformity of the distribution can be checked using scatterplots and histograms (Fig. 13.7).

The results of a uniform process can no longer be analysed using ANOVA. Since the parameter points are now spread out over the parameters space, the exact value of the parameters has to be taken into account. A Gaussian process with a Gaussian covariance function was therefore used. Afterwards, the results were compared to an analysis using a Gaussian process with a cubic covariance function.

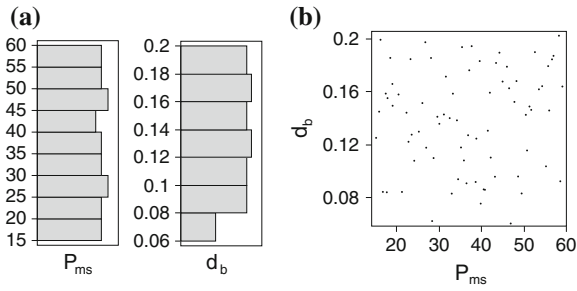


Fig. 13.7 Uniformity of the distribution in a uniform design. **a** Histogram showing the distribution of the parameters P_{ms} and d_b in a uniform design. In this particular example, it can be seen that the distribution of P_{ms} is uniform, while for the parameter d_b , only a small number of low values are included. **b** Scatterplot of all parameter combinations. The parameter combinations are spread out over the entire parameter space

One remark concerning the results of the Gaussian process is that one should take care when interpreting the predictions. When assessing the accuracy of the predictions, two methods can be used. Firstly, looking at the value of the function \mathcal{M} , defined in (13.17), will give a good idea of the efficiency of the process. Secondly, JMP plots the actual results, computed by your model, and compares these to the results predicted by the Gaussian process. Figure 13.8 illustrates this for the specific examples of mineralization of the alveolar bone and the concentration of fibroblasts in the PDL. As can be seen from this figure, the results for the mineralization of the bone are predicted far more accurately than the results for the concentration of fibroblasts. Since the horizontal axis plots the predicted value and the vertical axis the actual value, all the points would be located on a 45-degree line in a perfect model. A deviation from this line can have various causes. Firstly, because a more

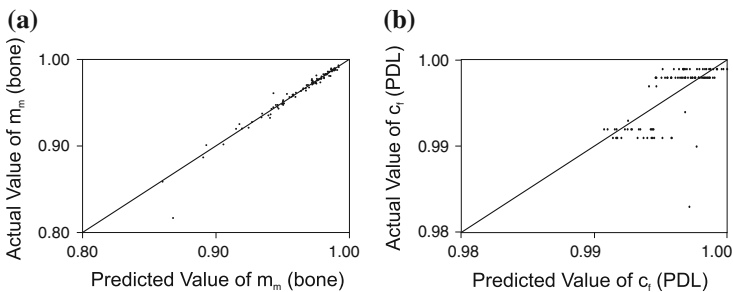


Fig. 13.8 Actual by predicted plot for the mineralization of the alveolar bone and the fibroblast density in the PDL. The horizontal axis shows the value of the output predicted by the Gaussian process, the vertical axis shows the actual output value of the mechanobiological model. **a** Actual by predicted plot for the mineralization of the alveolar bone. **b** Actual by predicted plot for the fibroblast density in the PDL

sophisticated covariance function may be needed, but secondly, as is the case here, because the variation in the output variable is so small the numerical errors become too important. When interpreting the analysis, care should be taken only to include those results with accurate predictions.

Steady State

The predictions for the most important output variables are summarized in Fig. 13.9. From these results it can be seen that the mineralization rate P_{ms} of the collagen fibres by osteoblasts positively influences the mineralization of both the alveolar bone and the PDL, and has a negative influence on the amount of fibroblasts in the bone. Increasing the apoptosis rate of osteoblasts d_b decreases the number of osteoblasts and the concentrations of RANKL and OPG, both produced by osteoblasts. Raising the production rate G_{gb} of TGF- β by osteoblasts increases the mineralization of the bone and the TGF- β concentration while decreasing the number of osteoclasts. Increasing the denaturation rate d_{gb} has the exact opposite effect. From this uniform design, it is now clear that increasing C_{mh} seems to favour bone formation by decreasing the number of osteoclasts.

The variable v_{SS} only becomes non-zero for low values of d_{gb} , but this result should be taken lightly, since out of 100 parameter combinations, the parameter was non-zero for only two parameter combinations. The results also show that limiting the parameter ranges of d_b , G_{gb} and C_{mh} had its desired effect, keeping the mineralization in reasonable bounds. As a result, the output *Real* has a value of one over the entire parameter space (results not shown).

Bone Formation and Bone Resorption

As was already shown in the ANOVA analysis, the main influence on the speed of bone formation was the parameter m_{bt} , although increasing the mineralization rate also increased bone formation. As Fig. 13.10 shows, a value of $m_{bt} > 0.4$ reduced the speed of bone formation to zero. Decreasing the denaturation rate d_{gb} greatly increased the concentration of TGF- β in the PDL.

The concentration of RANKL is clearly influenced by the production rate E_{grf} and the denaturation rate d_{gr} , but the speed of bone resorption shows some complex behaviour that will be looked at more closely in the next analysis. In summary, the osteogenic parameters P_{ms} and G_{gb} slow down bone resorption, while increasing d_b , E_{grf} or d_{gb} allowed bone to be resorbed faster.

The Effect of the Choice of the Covariance Function

In order to study the influence of the covariance function, the analysis was repeated using a cubic covariance function. As discussed in Sect. 13.2.2.2, the cubic covariance function includes smaller length scale, allowing smaller features to be incorporated in the results.




Fig. 13.9 Results of the uniform design concerning the equilibrium conditions in the alveolar bone and the PDL, and the parameters *Real* and v_{SS} . Only significant results are shown. *Solid lines* represent the prediction of the non-dimensionalised output variable by the Gaussian process, *dashed lines* represent the confidence intervals. The *first column* of figures shows the influence of P_{ms} on several output variables, assuming all other parameters have values in the middle of their range (e.g. $d_b = 0.135$, $G_{gb} = 0.85$, etc.). The same holds for the other columns showing the influence of d_b , G_{gb} , d_{gb} and C_{mh}




Fig. 13.10 Results of the uniform design concerning bone formation and resorption during orthodontic tooth movement. Only significant results are shown. *Solid lines* represent the prediction of the non-dimensionalised output variable by the Gaussian process, *dashed lines* represent the confidence intervals

The majority of the results for the cubic covariance function did not differ from the results with the Gaussian covariance function. However, there were three major differences. Firstly, when looking at the results concerning the parameter d_b (the apoptosis rate of osteoblasts), θ_m becomes very large ($\theta_m > 50$) for several output variables. Since $1/\theta_m$ represents the length scale of the modelled features, this means that very small variations will be incorporated into the results. As can be seen from Fig. 13.11 this does not change the general trend of the results. Since the parameter d_b (and many others) represent a physical process, the small variations are most likely due to numerical errors. Although non-linearities occur in nature, changing one parameter will more than likely result in a smooth change of the outcome variable.

Secondly, when studying the results for the speed of bone formation and the mineralization of the newly formed bone (Fig. 13.12), some highly non-linear behaviour can be seen. Once m_{bt} exceeds $m_{bt} = 0.4$, the predicted results for the mineralization of the newly formed bone show a very erratic behaviour. Furthermore, the predicted effect of H_{11} on the speed of bone formation shows similar variations. The cause of this effect can be seen when looking at the marginal model plots for H_{11} and m_{bt} , produced by JMP. These are plots that show the actual and the predicted values of an output variable as a function of only one parameter, disregarding the influence of other parameters. Figure 13.13 shows the actual and predicted values of v_F as a function of these two parameters. There is an apparent split into two classes: one for which the speed of bone formation is positive, and one for which it is zero. The latter





Fig. 13.11 Some results of the uniform design obtained using the cubic covariance function. Small fluctuations can be seen in the predicted influence of the osteoblast apoptosis rate d_b on the osteoblast density in the alveolar bone, as well as in the influence of d_b on the RANKL and OPG concentrations in the alveolar bone

Fig. 13.12 The influence of H_{11} and m_{bt} on bone formation, using the cubic covariance function, shows some highly non-linear behaviour and very erratic behaviour with very large confidence intervals



class appears when $m_{bt} > 0.4$, regardless of the values of other parameters. Since the cubic covariance function allows for very small length scales, the Gaussian process attempts to include both classes into one model, when actually they represent very different processes.

Finally, it should be noted that in general, the confidence intervals are much larger when using the cubic covariance function, compared to the Gaussian covariance function. This is particularly apparent in the case of the speed of bone formation. In that case, the plot of the actual value versus the predicted value already indicates a problem when using the cubic covariance function, while for the Gaussian covariance function the results were good (Fig. 13.14).

Preliminary Conclusions

From this uniform design, some preliminary conclusions can be drawn. Firstly, the use of the cubic covariance function added little to the value of the analysis. On the contrary, it introduced some errors and problems that did not appear when using the Gaussian covariance function. Secondly, when setting up more designs, care should




Fig. 13.13 Marginal model plots for H_{11} and m_{bt} . The horizontal axis shows the value of the parameter, the vertical axis shows the value of the output variable v_F . The dots indicate actual results obtained using the mechanobiological model, the solid line shows the results predicted by the Gaussian process. **b** The influence of m_{bt} , indicating how close MSC's have to be to the bone to differentiate into osteoblasts, on the speed of bone formation. When $m_{bt} > 0.4$, no bone formation occurs. **a** The influence of the sensitivity of MSC's to TGF- β (determined by H_{11}) on the speed of bone formation. The actual results can be split up into two classes: one for which there is bone formation, and one for which $v_F = 0$, located at the bottom




Fig. 13.14 Actual by predicted plot for the speed of bone formation, using the Gaussian and the cubic covariance function. **a** Actual by predicted plot using the Gaussian covariance function. All points are located quite close to the 45-degree line, indicating a good prediction. **b** Actual by predicted plot using the cubic covariance function. In this case, the points are spread out, indicating the predictions made by the sensitivity analysis for the speed of bone formation using the cubic covariance function are less reliable

be taken with the parameter range of m_{bt} . To avoid that the speed of bone formation drops to zero, m_{bt} should be kept smaller than 0.4.

This uniform design was implemented using the ten parameters which proved to be most influential according to the ANOVA analysis of the fractional factorial design explained in Sect. 13.3.2.1. However, the high number of output variables complicated the analysis of the results. Furthermore, most of the output variables were only influenced by one, or sometimes two, of the parameters. This implies that the uniform design will not give any information on the interaction between different parameters. For these reasons, it was decided to run two more uniform designs. The first one studies the influence of 12 parameters on the steady states, the second one looks more closely into the influence of 12 (different) parameters on bone formation and bone resorption.

13.3.2.3 Sensitivity Analysis on the Equilibrium Concentrations Using a Uniform Design

A 12-parameter uniform design with 100 runs was set up. Using the fractional factorial design described in Sect. 13.3.2.1, 12 parameters were chosen that influenced the steady states the most. In Table 13.5, the most important parameter was chosen out of each column related to the steady state. This set of seven parameters was then expanded with five parameters which were not included yet, but had a significant influence on one or more outcome variables. Since no results for bone formation or resorption will be calculated, and in order to study the influence of the parameters over the entire parameter space, the parameter ranges were kept the same as in the fractional factorial design.

Steady State

The parameters can be classified into two sets: those favouring bone formation and those favouring bone resorption. The most significant results of the analysis are shown in Fig. 13.15. Increasing the production rate of TGF- β (G_{gb}) increases the mineralization of the bone, while decreasing the osteoclast concentration and the number of fibroblasts in the bone. The apoptosis rate of osteoblasts (d_b), the decay rate of TGF- β (d_{gb}), the demineralization rate of the bone by osteoclasts (Q_{md}), the sensitivity of HSC's to RANKL (Y_2) and the maximum RANKL carrying capacity of osteoblasts (R_1) have the opposite effect. They decrease mineralization and osteoblast concentration, while increasing the number of osteoclasts, resulting in a lower mineralization of the bone. As in the previous uniform design, the haptotactic parameter C_{mh} seemed to encourage bone formation by decreasing the number of osteoclasts. Finally, increasing the proliferation rate of fibroblasts A_{f0} increased the density of the PDL and the concentration of fibroblasts, although the influence was very small and confidence intervals quite large.

The analysis also shows a significant interaction between the parameters H_{11} , G_{gb} and d_{gb} when modelling the concentrations of osteoblasts and TGF- β . When looking at the results of the analysis more closely, it can be concluded that for low values of H_{11} and d_{gb} (Fig. 13.16), the model is more non-linear. The influence of G_{gb} becomes more irregular and highly non-linear. The marginal model plots (Fig. 13.17) suggest this might be related to the presence of two outliers, in which the concentrations of osteoblasts and TGF- β are much higher than average.

These outliers are located in an area of the parameter space with a high sensitivity of MSC's to TGF- β , described by a low value of H_{11} , combined with a long half-life of TGF- β , expressed by the low value of d_{gb} and an average to high production of TGF- β by osteoblasts (G_{gb}). This particular combination results in a chain reaction in which the concentration of osteoblasts will rise quickly, producing even more TGF- β . The fact that the analysis predicts that the concentrations of osteoblasts will return back to normal for high values of G_{gb} can be explained by the lack of parameter combinations in that particular area of the parameter space. This was confirmed by running the model with parameter combinations in that area of the parameter space.




Fig. 13.15 Results of the uniform design on the equilibrium concentrations. Only significant results are shown. *Solid lines* represent the prediction of the non-dimensionalised output variable by the Gaussian process, *dashed lines* represent the confidence intervals

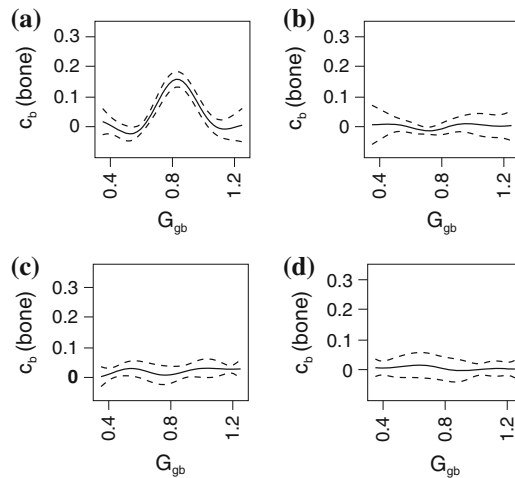


Fig. 13.16 Influence of the production rate of TGF- β (G_{gb}) on the predicted osteoblast density in bone, for different values of H_{11} and d_{gb} . A low value of H_{11} indicates a high sensitivity of MSC's to TGF- β and a low value of d_{gb} coincides with a long half life of TGF- β . **a** Results for $H_{11} = 0.08$ and $d_{gb} = 60$, for which the influence of G_{gb} on the osteoblast density is highly non-linear. **b** Results for $H_{11} = 0.18$ and $d_{gb} = 180$, showing a limited influence of G_{gb} on the osteoblast density. **c** Results for $H_{11} = 0.08$ and $d_{gb} = 180$, showing a limited influence of G_{gb} on the osteoblast density. **d** Results for $H_{11} = 0.18$ and $d_{gb} = 60$, showing a limited influence of G_{gb} on the osteoblast density




Fig. 13.17 Marginal model plots for G_{gb} , d_{gb} and H_{11} . The *horizontal axis* shows the value of the parameter, the *vertical axis* shows the osteoblast concentration in alveolar bone. The *dots* indicate actual results obtained using the mechanobiological model, the *solid line* shows the results predicted by the Gaussian process. On each figure, two outliers (*circled*) with a very high osteoblast density can be noted. **a** The influence of the production rate of TGF- β (G_{gb}) on osteoblast density. **b** The influence of the denaturation rate of TGF- β (d_{gb}) on osteoblast density. **c** The influence of the sensitivity of MSC's to TGF- β (determined by H_{11}) on osteoblast density

For the variable v_{SS} no conclusions can be drawn, since, out of 100 runs, only three values are non-zero. To prevent non-physiological situations in the next design, the results for the variable *Real* were interpreted. It was concluded that the appearance of such a situation was most likely with high values of d_b , d_{gb} , Y_2 , Q_{md} and R_1 , all parameters favouring bone resorption, or with low values of G_{gb} and C_{mh} .

The Effect of the Choice of the Covariance Function

The results of the uniform design were also analysed using a cubic covariance function. As before, results were generally very similar to the Gaussian covariance analysis. One observed difference is the effect of the parameter d_{gb} on the osteoblast concentration and the presence of TGF- β , OPG and RANKL in the bone. Figure 13.18 shows the results obtained using the cubic covariance function, revealing a sharp rise in osteoblast concentrations for small values of d_{gb} . Using the Gaussian covariance function, this rise is only seen to that extent when a low value of d_{gb} is combined with a low value of H_{11} . The analysis with the cubic covariance function predicts that the chain reaction leading to an alveolar bone with a high osteoblast and TGF- β content occurs at low values of d_{gb} , independent of the sensitivity of MSC's to TGF- β , described by H_{11} . To get a more detailed description of this phenomenon, a more elaborate sensitivity analysis with more points in that region of parameter space is needed.

13.3.2.4 Sensitivity Analysis on Bone Formation and Bone Resorption Using a Uniform Design

A 12-parameter uniform design with 100 runs was used to assess the influence of several parameters on tooth movement. Using the fractional factorial design described in Sect. 13.3.2.1, 12 parameters were chosen that influenced the processes of bone formation and resorption during tooth movement the most. In Table 13.5, the most




Fig. 13.18 The influence of d_{gb} on the osteoblast density and TGF- β concentration in the alveolar bone, using the cubic covariance function. When using the cubic covariance function a sharp rise in both osteoblast density and TGF- β concentration is predicted for low values of d_{gb} . This rise is not predicted when using the Gaussian covariance function (Fig. 13.15)

important parameter was chosen out of each column. This set of seven parameters was then expanded with five parameters which were not included yet, but had a significant influence on one or more outcome variables.

To prevent the occurrence of non-physiological situations in which the mineralization of the alveolar bone was too low, the parameter ranges of d_{gb} and d_b were limited. The upper limit of m_{bt} was also changed to 0.4, in order to avoid situations in which the speed of bone formation dropped to zero (see Sect. 13.3.2.2).

As expected, the speed of bone formation was mainly influenced by the parameter m_{bt} (Fig. 13.19), but also increased with increasing mineralization rate P_{ms} . The mineralization of the newly formed bone decreases when the decay rate of TGF- β (d_{gb}) is higher and the presence of TGF- β in the ligament was influenced by the production rate E_{gb} of TGF- β by fibroblasts and the parameter d_{gb} .

The effect of the parameters on the resorption rate v_R is less clear. Most parameters influence the resorption rate to some extent. Increasing the mineralization rate P_{ms} or the apoptosis rate d_{gr} of RANKL slows down bone resorption during tooth movement. Increasing the demineralization rate Q_{md} , the apoptosis rate d_b of osteoblasts and d_{gb} of TGF- β , or the production E_{grf} of RANKL by fibroblasts speeds up bone resorption. As expected, the concentration of RANKL in the PDL during bone resorption is determined by the parameters E_{grf} and d_{gr} . Similar results were obtained using the cubic covariance function.

13.3.3 Discussion

The design of experiments was used to determine the most important parameters of the model, and to investigate their result on the outcome of the model. The sensitivity




Fig. 13.19 Results of the uniform design on bone formation and resorption. Only significant results are shown. *Solid lines* represent the prediction of the non-dimensionalised output variable by the Gaussian process, *dashed lines* represent the confidence intervals

analysis started with a two-level fractional factorial design, followed by an ANOVA analysis of the results. Several parameters were then selected for a more in-depth analysis. Three different uniform designs were performed using these parameters, in order to assess their influence on the alveolar bone, the PDL, and the speed of bone formation and resorption.

The results of these analyses are summarized in Table 13.7. During homeostasis, bone resorption was shown to be promoted by the activity of osteoclasts (Q_{md}) and their sensitivity to RANKL (Y_2). Increasing the decay constant used to model the apoptosis of osteoblasts (d_b), or the denaturation rate of TGF- β (d_{gb}) had similar effects. The disappearance of both osteoblasts and TGF- β from the bone decreased the bone mass. Increasing bone mass was mainly noted as a result of increasing the production of TGF- β (G_{gb}) or the activity of osteoblasts (P_{ms}).

Under normal conditions, the model has two separate steady states. The first one represents the PDL, with a high concentration of fibroblasts and no mineralization, the second one represents the alveolar bone, consisting of mineralized collagen fibres, osteoblasts and osteoclasts. Changing the parameters usually resulted in a small, continuous change of the steady states, preserving their main properties. Alveolar bone and PDL remained recognisable as such. However, in some cases more abrupt changes were noted.

Firstly, with several parameter combinations that promoted bone resorption, the predicted mineralization of the alveolar bone dropped to zero. With those parameter combinations, only one steady state was possible and the alveolar bone could not be modelled. Secondly, when a high sensitivity of MSC's to TGF- β was combined with a high production and a long half-life of TGF- β , a chain reaction occurred. The steady state representing the alveolar bone switched from bone containing a small

Table 13.7 Summary of the parameters that showed the most influence on equilibrium conditions and tooth movement (TM), as indicated by the sensitivity analysis

Parameter	Influence
Q_{md}	Promotes bone resorption
d_b	Promotes bone resorption
	Increases bone resorption during TM
Y_2	Promotes bone resorption
G_{gb}	Promotes bone formation
d_{gb}	Promotes bone resorption
	Increases TGF- β concentration during TM
	Increases mineralization during TM
	Increases bone resorption during TM
P_{ms}	Promotes bone formation
	Increases bone formation during TM
	Decreases bone resorption during TM
m_{bt}	Decreases bone formation during TM
E_{gb}	Increases TGF- β concentration during TM
d_{gr}	Decreases bone resorption during TM
	Decreases RANKL concentration during TM
E_{grf}	Increases bone resorption during TM
	Increases RANKL concentration during TM

amount of osteoblasts to bone containing large amounts of osteoblasts. Since the parameter combinations at which the phenomenon occurred were located at the edge of the parameter space investigated in this analysis, it did not pose any problems. For the same reason, not much information of the phenomenon and its influences on tooth movement was obtained. Performing a sensitivity analysis with a broader parameter space will give a better insight in this chain reaction.

Concerning the parameter C_{mh} , describing the haptotaxis of osteoclasts, an unexpected effect was discovered. The modelling of haptotaxis of osteoclasts, attracting them to the alveolar bone, was included in the model to mimic the function of osteopontin and the working of osteoclasts during bone resorption. The fact that increasing C_{mh} seemed to favour bone formation during steady state was unexpected, and initiated a closer investigation of this process. It was concluded that this effect was related to the implementation of the continuous transition between the PDL and the alveolar

bone during steady state, as described in equation (13.29). Due to this approach, the mineralization in the alveolar bone is not constant, but increases slightly when moving away from the PDL. As a result, with higher values of C_{mh} , more osteoclasts will migrate to the more mineralized part of the domain, where osteoclast apoptosis is slightly higher due to the higher concentration of TGF- β . The number of osteoclasts at the measuring point will thus decrease, resulting in a higher degree of mineralization. It should be noted that, although the effect does occur, it is more subtle than predicted by the sensitivity analysis.

Four parameters had a significant influence on bone formation during orthodontic tooth movement. The activity of osteoblasts (P_{ms}) increased bone formation, while the parameter m_{bt} had the opposite effect. The production and denaturation of TGF- β had a very high influence on the TGF- β concentration in the PDL. The same conclusion can be drawn for the concentration of RANKL in the PDL at the resorption site during tooth movement. This concentration was mainly influenced by the production and denaturation of RANKL. Decreasing the activity of osteoblasts or decreasing their life span speeds up bone resorption during tooth movement, as well as decreasing the half-life of TGF- β .

The results of all uniform designs were analysed using a Gaussian process with both a Gaussian and a cubic covariance function. The general trend of the results was the same, regardless of the applied covariance function. However, in some cases the cubic covariance function picked up on numerical errors, while the Gaussian covariance function was able to filter them out. On the other hand, when dealing with outliers, the cubic covariance function had less problems incorporating them into the analysis. When analysing the results of physical experiments outliers are usually the result of measurement errors and random noise, and incorporating them into the analysis should be avoided. When investigating computer models, such errors are not present, and outliers are indicators of abrupt changes in the outcome of the model. As such, they should not be ignored, but instead be investigated more closely.

With the results of the sensitivity analysis, it was possible to identify the ten parameters that had the highest influence on the outcome of the mechanobiological model presented in this work. Most notable was the high influence of osteoblast apoptosis and the half-life of TGF- β . When using experimental results to refine the estimates of the parameters, these are the parameters that should be looked at first.

Some further observations were made that should be kept in mind when analysing and developing the model further. Some parameter combinations promoted bone resorption too strongly, resulting in a total loss of bone mass. Others resulted in a fast occurrence of a large number of osteoblasts in the bone. The effect of the mechanical stimulus on bone formation proved to be small, indicating a closer look into the assumption that osteoblasts will only form close to bone could be useful.

Some unexpected results did occur, and were investigated closer. In most cases valid explanations for the phenomena were found, and no further actions were necessary, in other cases, suggestions for further analysis could be made. In summary, the extensive sensitivity analysis resulted in a better understanding of model and its parameters.

13.4 Conclusions

The design of experiments is a valuable tool for studying a large variety of computer models. Any type of model can be analysed using a DOE, as long as the experiment can be easily repeated using different parameters. The designs are developed to get as much information as possible at minimal cost, usually referring to the computational time. Each design has its pros and cons and depending on the needs of the researcher a suitable design is generated. Several statistical software packages provide support for DOE.

As with all statistical tools, care has to be taken when interpreting the results of a DOE, certainly with more complex models. Both the designs and the analysing methods have their limitations. However, keeping those in mind, DOE is a great tool, allowing the researcher to gain a better insight into their model and its behaviour in parameter space. Moreover, the results of a DOE analysis can be a good starting point for further model optimization (see [4]).

Acknowledgments The authors acknowledge support from European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement n279100 (An Van Schepdael and Liesbet Geris). Aurélie Carlier is a Fellow of the Research Foundation Flanders (FWO).

Conflict of Interest

The authors declare that they have no conflict of interest.

References

1. Carlier, A., Chai, Y.C., Moesen, M., Theys, T., Schrooten, J., Van Oosterwyck, H., Geris, L.: Designing optimal calcium phosphate scaffold-cell combinations using an integrative model-based approach. *Acta Biomater.* **7**, 3573–3585 (2011)
2. Dar, F.H., Meakin, J.R., Aspden, R.M.: Statistical methods in finite element analysis. *J. Biomech.* **35**, 1155–1161 (2002)
3. Ebdon, M.: Gaussian processes for regression: a quick introduction. *Lecture Notes* (2008)
4. Eriksson, O., Tegnér, J.: Modeling and model simplification to facilitate biological insights and predictions. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
5. Fang, K.T.: The uniform design: application of number-theoretic methods in experimental design. *Acta Mathematicae Applicatae Sinica* **3**, 363–372 (1980)
6. Fang, K.T., Li, R., Sudijanto, A.: *Design and Modeling for Computer Experiments*. Chapman & Hall/CRC, Boca Raton (2006)
7. Fisher, R.A.: *The Design of Experiments*. Oliver & Boyd, Edinburgh (1935)
8. Garant, P.R. (ed.): *Oral Cells and Tissues*. Quintessence Publishing, Hanover Park (2003)
9. Geris, L., Gerisch, A., Maes, C., Carmeliet, G., Weiner, R., Vander Sloten, J., Van Oosterwyck, H.: Mathematical modeling of fracture healing in mice: comparison between experimental data and numerical simulation results. *MBEC* **44**, 280–289 (2006)

10. Geris, L., Gerisch, A., Sloten, J.V., Weiner, R., Van Oosterwyck, H.: Angiogenesis in bone fracture healing: a bioregulatory model. *J. Theory Biol.* **251**(1), 137–158 (2008)
11. Henneman, S., Von den Hoff, J.W., Maltha, J.C.: Mechanobiology of tooth movement. *Eur. J. Orthod.* **30**(3), 299–306 (2008)
12. Isaksson, H., van Donkelaar, C.C., Huiskes, R., Ito, K.: A mechano-regulatory bone-healing model incorporating cell-phenotype specific activity. *J. Theory Biol.* **252**, 230–246 (2008a)
13. Isaksson, H., van Donkelaar, C.C., Huiskes, R., Yao, J., Ito, K.: Determining the most important cellular characteristics for fracture healing using design of experiments methods. *J. Theory Biol.* **255**(1), 26–39 (2008)
14. Kanzaki, H., Chiba, M., Shimizu, Y., Mitani, H.: Periodontal ligament cells under mechanical stress induce osteoclastogenesis by receptor activator of nuclear factor kappaB ligand up-regulation via prostaglandin e2 synthesis. *J. Bone Miner. Res.* **17**(2), 210–220 (2002)
15. Kimoto, S., Matsuzawa, M., Matsubara, S., Komatsu, T., Uchimura, N., Kawase, T., Saito, S.: Cytokine secretion of periodontal ligament fibroblasts derived from human deciduous teeth: effect of mechanical stress on the secretion of transforming growth factor-beta 1 and macrophage colony stimulating factor. *J. Periodontal Res.* **34**(5), 235–243 (1999)
16. Kirk, P., Silk, D., Stumpf, M.P.H.: Reverse engineering under uncertainty. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
17. Krishnan, V., Davidovitch, Z.: Cellular, molecular, and tissue-level reactions to orthodontic force. *Am. J. Orthod. Dentofacial Orthop.* **129**(4), 469.e1-469.32 (2006)
18. Krishnan, V., Davidovitch, Z.: On a path to unfolding the biological mechanisms of orthodontic tooth movement. *J. Dent. Res.* **88**(7), 597–608 (2009)
19. Lacroix, D.: Simulation of tissue differentiation during fracture healing. PhD thesis, University of Dublin (2001)
20. Lin, C.L., Chang, S.H., Chang, W.J., Kuo, Y.C.: Factorial analysis of variables influencing mechanical characteristics of a single tooth implant placed in the maxilla using finite element analysis and the statistics-based taguchi method. *Eur. J. Oral Sci.* **115**, 408–416 (2007)
21. Lundstedt, T., Seifert, E., Abramo, L., Thelin, B., Nyström, A., Pettersen, J., Bergman, R.: Experimental design and optimization. *Chemom. Intell. Lab.* **42**(2), 3–40 (1998)
22. MacKay, D.J.C.: Introduction to gaussian processes. *Lecture Notes* (1998)
23. Malandrino, A., Planell, J., Lacroix, D.: Statistical factorial analysis on the poroelastic material properties sensitivity of the lumbar intervertebral disc under compression, flexion and axial rotation. *J. Biomech.* **42**, 2780–2788 (2009)
24. Mannakee, B.K., Ragsdale, A.P., Transtrum, M.K., Gutenkunst, R.N.: Sloppiness and the geometry of parameter space. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
25. Marotti, G.: The osteocyte as a wiring transmission system. *J. Musculoskelet. Neuronal. Interact.* **1**(2), 133–136 (2000)
26. Montgomery, D.C.: *Design and analysis of experiments*, 7th edn. Wiley, New York (1997)
27. Myers, R.H., Montgomery, D.C.: *Response surface methodology: process and product optimization using designed experiments*. Wiley, New York (1995)
28. Nishijima, Y., Yamaguchi, M., Kojima, T., Aihara, N., Nakajima, R., Kasai, K.: Levels of RANKL and OPG in gingival crevicular fluid during orthodontic tooth movement and effect of compression force on releases from periodontal ligament cells in vitro. *Orthod. Craniofac. Res.* **9**(2), 63–70 (2006)
29. Pfeilschifter, J., Diel, I., Scheppach, B., Bretz, A., Krempien, R., Erdmann, J., Schmid, G., Reske, N., Bismar, H., Seck, T., Krempien, B., Ziegler, R.: Concentration of transforming growth factor beta in human bone tissue: relationship to age, menopause, bone turnover, and bone volume. *J. Bone Miner. Res.* **13**(4), 716–730 (1998)
30. Pinkerton, M.N., Wescott, D.C., Gaffey, B.J., Beggs, K.T., Milne, T.J., Meikle, M.C.: Cultured human periodontal ligament cells constitutively express multiple osteotropic cytokines and growth factors, several of which are responsive to mechanical deformation. *J. Periodontal Res.* **43**(3), 343–351 (2008)

31. Pivonka, P., Zimak, J., Smith, D.W., Gardiner, B.S., Dunstan, C.R., Sims, N.A., Martin, T.J., Mundy, G.R.: Model structure and control of bone remodeling: a theoretical study. *Bone* **43**(2), 249–263 (2008)
32. Provatidis, C.G.: An analytical model for stress analysis of a tooth in translation. *Int. J. Eng. Sci.* **39**, 1361–1381 (2001)
33. Rinchuse, D.J., Rinchuse, D.J., Sosovicka, M.F., Robison, J.M., Pendleton, R.: Orthodontic treatment of patients using bisphosphonates: a report of 2 cases. *Am. J. Orthod. Dentofacial Orthop.* **131**(3), 321–326 (2007)
34. Saltelli, A., Chan, K., Scott, E.M. (eds.): *Sensitivity Analysis*. Wiley, New York (2000)
35. Sandberg, M., Vuorio, T., Hirvonen, H., Alitalo, K., Vuorio, E.: Enhanced expression of *tgf-beta* and *c-fos* mRNAs in the growth plates of developing human long bones. *Development* **102**(3), 461–470 (1988)
36. Santner, T.J., Williams, B.J., Notz, W.I.: *The Design and Analysis of Computer Experiments*. Springer, New York (2003)
37. Van Schepdael, A., Vander Sloten, J., Geris, L.: A mechanobiological model of orthodontic tooth movement. *Biomech. Model Mechanobiol.* **12**(2), 249–265 (2013)
38. Van Schepdael, A., Vander Sloten, J., Geris, L.: Mechanobiological modeling can explain orthodontic tooth movement: three case studies. *J. Biomech.* **46**(3), 470–477 (2013)
39. Wescott, D.C., Pinkerton, M.N., Gaffey, B.J., Beggs, K.T., Milne, T.J., Meikle, M.C.: Osteogenic gene expression by human periodontal ligament cells under cyclic tension. *J. Dent. Res.* **86**(12), 1212–1216 (2007)
40. Yamaguchi, M., Aihara, N., Kojima, T., Kasai, K.: Rankl increase in compressed periodontal ligament cells from root resorption. *J. Dent. Res.* **85**(8), 751–756 (2006)
41. Yang, K., Teo, E.C., Fuss, F.K.: Application of Taguchi method in optimization of cervical ring cage. *J. Biomech.* **40**, 3251–3256 (2007)

Chapter 14

Waves in Spatially-Disordered Neural Fields: A Case Study in Uncertainty Quantification

Carlo R. Laing

Abstract Neural field models have been used for many years to model a variety of macroscopic spatiotemporal patterns in the cortex. Most authors have considered homogeneous domains, resulting in equations that are translationally invariant. However, there is an obvious need to better understand the dynamics of such neural field models on heterogeneous domains. One way to include heterogeneity is through the introduction of randomly-chosen “frozen” spatial noise to the system. In this chapter we investigate the effects of including such noise on the speed of a moving “bump” of activity in a particular neural field model. The spatial noise is parameterised by a large but finite number of random variables, and the effects of including it can be determined in a computationally-efficient way using ideas from the field of Uncertainty Quantification. To determine the average speed of a bump in this type of heterogeneous domain involves evaluating a high-dimensional integral, and a variety of methods are compared for doing this. We find that including heterogeneity of this form in a variety of ways always slows down the moving bump.

Keywords Neural field · Uncertainty quantification · Pattern formation · Heterogeneity · Integration

14.1 Introduction

Neural field models have been used for many years as models of large-scale pattern formation in the cortex [1, 8, 13, 16, 27, 29, 31, 35]. These models are typically formulated as nonlocal partial differential equations in space and time where the nonlocality arises via spatial integrals, meant to represent the influence of neurons at many different spatial locations on the dynamics at a specific location [8, 13]. They have been used to model a variety of neurophysiological phenomena such as working

C.R. Laing (✉)

Institute of Natural and Mathematical Sciences, Massey University,
Private Bag, 102-904 NSMC Auckland, New Zealand
e-mail: c.r.laing@massey.ac.nz

© Springer International Publishing Switzerland 2016
L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_14

367

memory [31], orientation tuning in the visual cortex [4] and EEG rhythms [39]. Much of the analysis of patterns in these models has assumed that the domain is homogeneous and thus the governing equations are translationally invariant. This invariance allows one to, for example, choose the origin of space to simplify analysis. When studying travelling waves, this invariance means that it is relatively easy to construct “bumps” and fronts of activity which move with a constant speed. However, the brain is far from homogeneous and it is of interest to understand how various forms of heterogeneity affect the properties of moving waves in neural field models.

A number of authors have considered including heterogeneity in neural field models by introducing spatially periodic modulation of various components of the model such as connectivity [6, 24] and input currents [14]. This type of heterogeneity is structured rather than random, but a number of other authors have considered the effects of truly random heterogeneity, in either space, time, or both. For example, the authors [5] considered the effects on the speed of a front of a spatially uniform firing rate threshold which randomly fluctuated in time. They found that such fluctuations always increased the average front speed. Coombes et al. [15] briefly considered a variety of forms of heterogeneity such as adding “frozen” spatial noise, and driving the system with temporal noise. Bressloff [7] adapted ideas from PDE theory to study the effects of slowly modulated (in space) synaptic connectivity on the invasion and extinction of activity in a neural field model. Several authors have very recently considered the effects of additive spatio-temporal noise on the dynamics of a neural field [10, 11, 23, 25].

In this chapter we will use ideas from the relatively new field of Uncertainty Quantification (UQ) to investigate the effects of spatial heterogeneity on the dynamics of moving “bumps” in a particular neural field model. Traditionally, numerical models of physical phenomena have been solved under the assumption that both the initial conditions and all values of relevant parameters are known exactly. However, recent increases in computational power have meant that it is now possible to solve a model where one or more parameters are not known exactly, but are known (or assumed) to come from some distribution(s). For our purposes, UQ involves a systematic investigation of the effects of this uncertainty in parameter values on quantities of interest. The field of UQ is large and rapidly growing [33, 42, 45] and here we will only use those aspects of it which are directly relevant.

14.2 Model and Analysis

The model we first consider is governed by the following equations:

$$\frac{\partial u(x, t)}{\partial t} = -u(x, t) + \int_0^{2\pi} G(x - y)F[u(y, t) - a(y, t) + h(y)]dy \quad (14.1)$$

$$\tau \frac{\partial a(x, t)}{\partial t} = Bu(x, t) - a(x, t) \quad (14.2)$$

where $u(x, t)$ represents the average voltage of neurons at position $x \in [0, 2\pi]$ at time t , and $a(x, t)$ represents the value of a slow variable at x and t which provides negative feedback to the dynamics of u . Similar models have been studied elsewhere [9, 15, 28, 30, 36]. Periodic boundary conditions are used, B and τ are positive constants, and the firing rate function is given by

$$F[u] = \frac{1}{1 + e^{-20(u-0.4)}} \quad (14.3)$$

Note that F is bounded between 0 and 1 and is an increasing function. The function $h(y)$, to be specified below, provides the spatial heterogeneity to the system. We choose the coupling function to be $G(x) = 0.09 + 0.45 \cos x$; note that this is even. The physical interpretation of the model is that neurons with average voltage u fire at a frequency $F[u]$, and the strength of connections between neurons at position x and those at position y is $G(x - y)$. Summing (integrating) over all y gives the nonlinear term in (14.1), and this is the influence of all other neurons on those at position x . The variable a is driven up when u is high and down when u is low, with a time-scale of τ . The way that a appears in (14.1) means that it acts a negative feedback mechanism.

For suitable choices of parameters the system (14.1–14.2) is capable of supporting travelling “bumps” of activity. See Fig. 14.1a. A bump is defined to be a state in which one region of the domain is active, i.e. has $F[u] \approx 1$, while the rest of the domain is inactive, i.e. has $F[u] \approx 0$. When $h(y)$ is constant the bumps travel with constant speed and profile, while if $h(y)$ is not constant—but is sufficiently small—bumps continue to travel, but with non-constant speed and profile. See Fig. 14.1b–d. Because the domain has periodic boundary conditions, these modulated bumps are periodic in time. Our goal is to determine, in a computationally-efficient manner, the expected effects of making $h(y)$ a random function of y , in a way to be explained below. In particular we wish to answer the question: given that $h(y)$ is randomly chosen from some distribution of functions, what is the expected value of the average speed of the resulting travelling bump (after transients have decayed)? As mentioned, we will answer this using techniques from the field of uncertainty quantification [33, 41]. Here, the uncertainty arises because we do not exactly know $h(y)$. This uncertainty then affects the dynamics of the neural field model, making measurable quantities such as the bump speed uncertain, i.e. have some distribution of values. Typically, we would like to describe this distribution so that we can calculate, for example, its mean.

The form of the coupling function $G(x)$ allows us to write (14.1) as

$$\begin{aligned} \frac{\partial u(x, t)}{\partial t} = & -u(x, t) + 0.09 \int_0^{2\pi} F[u(y, t) - a(y, t) + h(y)] dy \\ & + 0.45 \cos x \int_0^{2\pi} F[u(y, t) - a(y, t) + h(y)] \cos y dy \\ & + 0.45 \sin x \int_0^{2\pi} F[u(y, t) - a(y, t) + h(y)] \sin y dy \quad (14.4) \end{aligned}$$

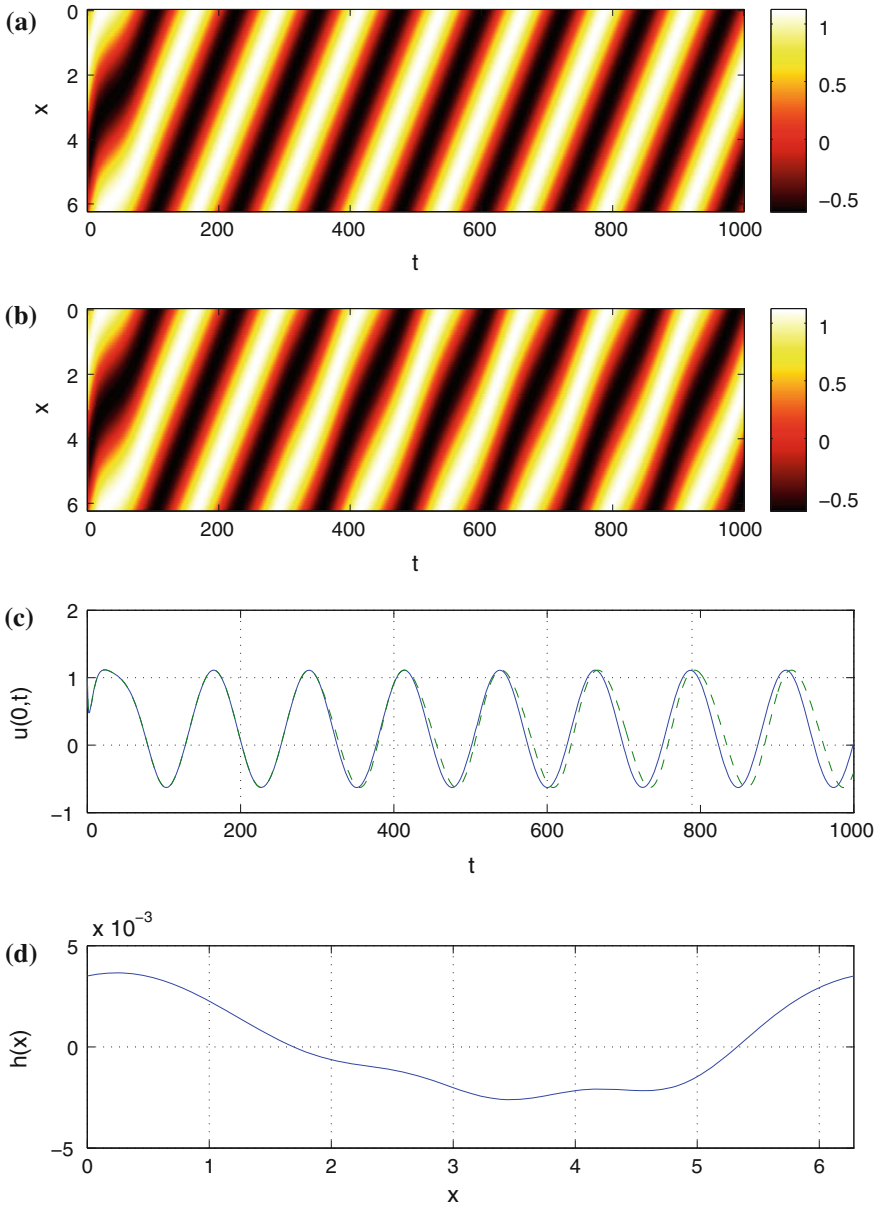


Fig. 14.1 Travelling bumps of activity in the model (14.1–14.2) when (a) $h(x) = 0$ and (b) when $h(x)$ is as shown in (d) for $t > 300$, $h(x) = 0$ for $t < 300$. Panel (c) shows $u(0, t)$ for the solution in panel (a) (blue solid line) and the solution in panel (b) (green dashed line). We see that non-zero $h(x)$ slightly slows the bump. Parameters: $B = 0.1$, $\tau = 14$

As noted [28], if we expand $u(x, t)$ and $a(x, t)$ in Fourier series in x we see that terms of the form $\sin(nx)$ and $\cos(nx)$ for $n > 1$ will decay to zero, and since we are not interested in transients we write

$$u(x, t) = u^0(t) + u^c(t) \cos x + u^s(t) \sin x \quad (14.5)$$

and

$$a(x, t) = a^0(t) + a^c(t) \cos x + a^s(t) \sin x \quad (14.6)$$

Substituting these expansions into (14.2) and (14.4) we find that the modulated bumps of interest are described by the six ordinary differential equations (ODEs)

$$\frac{du^0}{dt} = -u^0 + 0.09 \int_0^{2\pi} F[u^0 - a^0 + (u^c - a^c) \cos x + (u^s - a^s) \sin x + h(x)] dx \quad (14.7)$$

$$\frac{du^c}{dt} = -u^c + 0.45 \int_0^{2\pi} F[u^0 - a^0 + (u^c - a^c) \cos x + (u^s - a^s) \sin x + h(x)] \cos x dx \quad (14.8)$$

$$\frac{du^s}{dt} = -u^s + 0.45 \int_0^{2\pi} F[u^0 - a^0 + (u^c - a^c) \cos x + (u^s - a^s) \sin x + h(x)] \sin x dx \quad (14.9)$$

$$\tau \frac{da^0}{dt} = Bu^0 - a^0 \quad (14.10)$$

$$\tau \frac{da^c}{dt} = Bu^c - a^c \quad (14.11)$$

$$\tau \frac{da^s}{dt} = Bu^s - a^s \quad (14.12)$$

We note that the number of ODEs above (six) is an immediate consequence of using only a constant and $\cos x$ term in the coupling function $G(x)$. Our coupling function can be thought of as the truncation of the Fourier series of a general 2π -periodic function. Including more terms in this truncation would lead to the system being described by more ODEs, in the obvious way: each new harmonic would result in 4 more ODEs.

For a given $h(x)$ we can find the average speed of the resulting bump by finding the relevant periodic solution of (14.7–14.12). The average speed is then 2π (the size of the domain) divided by the period of this orbit. However, numerically integrating (14.7–14.12) is computationally costly, and since we are only interested in periodic solutions of these ODEs we represent their solutions as truncated Fourier series in time, i.e. we write

$$u^0(t) = u_0^0 + \sum_{i=1}^M \left[u_i^0 \cos(i\omega t) + u_{M+i}^0 \sin(i\omega t) \right] \quad (14.13)$$

$$u^c(t) = u_0^c + \sum_{i=1}^M [u_i^c \cos(i\omega t) + u_{M+i}^c \sin(i\omega t)] \tag{14.14}$$

$$u^s(t) = u_0^s + \sum_{i=1}^M [u_i^s \cos(i\omega t) + u_{M+i}^s \sin(i\omega t)] \tag{14.15}$$

$$a^0(t) = a_0^0 + \sum_{i=1}^M [a_i^0 \cos(i\omega t) + a_{M+i}^0 \sin(i\omega t)] \tag{14.16}$$

$$a^c(t) = a_0^c + \sum_{i=1}^M [a_i^c \cos(i\omega t) + a_{M+i}^c \sin(i\omega t)] \tag{14.17}$$

$$a^s(t) = a_0^s + \sum_{i=1}^M [a_i^s \cos(i\omega t) + a_{M+i}^s \sin(i\omega t)] \tag{14.18}$$

where $\omega = 2\pi/T$ and T is the unknown period of the periodic orbit we wish to find. We have

$$\frac{du^0}{dt} = \sum_{i=1}^M [-i\omega u_i^0 \sin(i\omega t) + i\omega u_{M+i}^0 \cos(i\omega t)]$$

and similarly for the other five functions. We are going to solve (14.7–14.12) by collocation. To do that we impose that the functions given in (14.13–14.18) satisfy the differential equations (14.7–14.12) at $2M+1$ different times in the interval $[0, T]$. Let these times be $t_j = jT/(2M+1)$, $j = 1, \dots, 2M+1$. Then we have

$$\begin{aligned} 0 = & -u^0(t_j) + \sum_{i=1}^M [i\omega u_i^0 \sin(i\omega t_j) - i\omega u_{M+i}^0 \cos(i\omega t_j)] \\ & + 0.09 \int_0^{2\pi} F[u^0(t_j) - a^0(t_j) + (u^c(t_j) - a^c(t_j)) \cos x \\ & + (u^s(t_j) - a^s(t_j)) \sin x + h(x)] dx \end{aligned} \tag{14.19}$$

$$\begin{aligned} 0 = & -u^c(t_j) + \sum_{i=1}^M [i\omega u_i^c \sin(i\omega t_j) - i\omega u_{M+i}^c \cos(i\omega t_j)] \\ & + 0.45 \int_0^{2\pi} F[u^0(t_j) - a^0(t_j) + (u^c(t_j) - a^c(t_j)) \cos x \\ & + (u^s(t_j) - a^s(t_j)) \sin x + h(x)] \cos x dx \end{aligned} \tag{14.20}$$

$$\begin{aligned}
0 = & -u^s(t_j) + \sum_{i=1}^M [i\omega u_i^s \sin(i\omega t_j) - i\omega u_{M+i}^s \cos(i\omega t_j)] \\
& + 0.45 \int_0^{2\pi} F[u^0(t_j) - a^0(t_j) + (u^c(t_j) - a^c(t_j)) \cos x \\
& + (u^s(t_j) - a^s(t_j)) \sin x + h(x)] \sin x \, dx
\end{aligned} \tag{14.21}$$

$$0 = Bu^0(t_j) - a^0(t_j) + \tau \sum_{i=1}^M [i\omega a_i^0 \sin(i\omega t_j) - i\omega a_{M+i}^0 \cos(i\omega t_j)] \tag{14.22}$$

$$0 = Bu^c(t_j) - a^c(t_j) + \tau \sum_{i=1}^M [i\omega a_i^c \sin(i\omega t_j) - i\omega a_{M+i}^c \cos(i\omega t_j)] \tag{14.23}$$

$$0 = Bu^s(t_j) - a^s(t_j) + \tau \sum_{i=1}^M [i\omega a_i^s \sin(i\omega t_j) - i\omega a_{M+i}^s \cos(i\omega t_j)] \tag{14.24}$$

for $j = 1, \dots, 2M + 1$. This gives us $12M + 6$ equations: (14.19–14.24), but there are $12M + 7$ unknowns (T being the last unknown). We also have freedom to choose the origin of time, so to remove this degeneracy and obtain the correct number of equations we add one more (largely arbitrary) condition to fix the phase of the periodic orbit: $u^c(0) = 0$, i.e.

$$u_0^c + \sum_{i=1}^M u_i^c = 0. \tag{14.25}$$

Equations (14.19–14.24) and (14.25) can be solved straightforwardly using Newton's method, where the integral over x is evaluated using the trapezoidal rule.

We now turn to the representation of the frozen noise, $h(x)$. We assume that it is a uniform random field with mean zero and covariance

$$C(x, y) = \frac{\sigma}{2b} \exp \left[-\frac{\pi}{4} \left(\frac{x - y}{b} \right)^2 \right] \tag{14.26}$$

so that σ determines its “strength” and b is the characteristic correlation length. We will represent $h(x)$ by its Karhunen-Loève decomposition [20, 22, 33]. To do this we need to find the eigenpairs of C , $\{\lambda_m, e_m(x)\}_{m=1}^\infty$, defined by

$$\int_{-\infty}^{\infty} C(x, y) e_m(y) \, dy = \lambda_m e_m(x) \tag{14.27}$$

and then order the eigenvalues (which are known to be positive and real): $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$ and normalise the eigenfunctions (which are known to be orthogonal) [33]. The Karhunen-Loève decomposition of h is then

$$h(x) = \sum_{m=1}^{\infty} \sqrt{\lambda_m} e_m(x) \beta_m \tag{14.28}$$

where the β_m are pairwise independent random variables with mean zero taken from the uniform distribution on $[-1, 1]$. We use a uniform random field rather than the more common Gaussian random field, where the β_m are normally-distributed, because we want the random field $h(x)$ to be bounded. The reason for this is that as the amplitude of $h(x)$ is increased, the moving bump seen in Fig. 14.1a can become “pinned” by the heterogeneity [18, 37]. This type of pinned solution is far in phase space from the original moving bump, and cannot be regarded as a small perturbation from it, due to the nonlinear nature of the problem, so it is not appropriate to consider such a solution. This is also the reason that we will only consider sufficiently small values of σ below. Similar reasoning is used when a random field is constrained by physical reasons to be strictly positive, for example [43].

To find the eigenpairs of C consider the function $\cos (my)$, where $m \in \mathbb{N}^+$. This is periodic on the domain $[0, 2\pi]$ and we have

$$2b \int_0^{2\pi} C(x, y) \cos (my) dy = \sigma \int_0^{2\pi} \exp \left[-\frac{\pi}{4} \left(\frac{x-y}{b} \right)^2 \right] \cos (my) dy \tag{14.29}$$

$$= \sigma \int_{x-2\pi}^x \exp \left[-\frac{\pi}{4} \left(\frac{z}{b} \right)^2 \right] \cos (m(x-z)) dz \tag{14.30}$$

Now if b is small relative to the domain size (2π), we can approximate this integral by the infinite one:

$$2b \int_0^{2\pi} C(x, y) \cos (my) dy \approx \sigma \int_{-\infty}^{\infty} \exp \left[-\frac{\pi}{4} \left(\frac{z}{b} \right)^2 \right] \cos (m(x-z)) dz \tag{14.31}$$

$$\begin{aligned} &= \sigma \cos (mx) \int_{-\infty}^{\infty} \exp \left[-\frac{\pi}{4} \left(\frac{z}{b} \right)^2 \right] \cos (mz) dz \\ &\quad + \sigma \sin (mx) \int_{-\infty}^{\infty} \exp \left[-\frac{\pi}{4} \left(\frac{z}{b} \right)^2 \right] \sin (mz) dz \end{aligned} \tag{14.32}$$

$$= 2b\sigma \cos (mx) \exp \left[\frac{-(mb)^2}{\pi} \right] \tag{14.33}$$

where we have used the fact that [38]

$$\int_{-\infty}^{\infty} \exp \left[-\frac{\pi}{4} \left(\frac{z}{b} \right)^2 \right] \cos (mz) dz = 2b \exp \left[\frac{-(mb)^2}{\pi} \right] \tag{14.34}$$

and that $\exp[-(\pi/4)(z/b)^2] \sin(mz)$ is an odd function. Thus (keeping in mind the approximations made above) a partial set of eigenvalues and eigenfunctions for C is

$$\lambda_m^{(1)} = \sigma \exp\left[\frac{-(mb)^2}{\pi}\right]; \quad e_m^{(1)}(x) = \frac{\cos(mx)}{\sqrt{\pi}} \quad (14.35)$$

for $m = 1, 2, \dots$. A similar argument shows that the remaining set of eigenvalues and eigenfunctions is

$$\lambda_m^{(2)} = \sigma \exp\left[\frac{-(mb)^2}{\pi}\right]; \quad e_m^{(2)}(x) = \frac{\sin(mx)}{\sqrt{\pi}} \quad (14.36)$$

for $m = 1, 2, \dots$. Eigenpairs for other covariance functions can be found either analytically or numerically [20, 22, 33, 43]. We truncate the series (14.28) to give a finite-dimensional representation of the random field. We write

$$h(x) = \sum_{m=1}^N \beta_m \sqrt{\lambda_m^{(1)}} e_m^{(1)}(x) + \sum_{m=1}^N \beta_{N+m} \sqrt{\lambda_m^{(2)}} e_m^{(2)}(x) \quad (14.37)$$

where the $\beta_1, \dots, \beta_{2N}$ are randomly chosen from the uniform distribution on $[-1, 1]$.

The idea is now that for each realisation of the $\{\beta_m\}$ we can numerically solve (14.19–14.24) and (14.25). One way to regard the solutions of these equations is that they are given by $12M + 7$ variables, $u_0^0, u_1^0, \dots, a_{2M}^s, T$, each of which is a function of the $2N$ variables $\beta_1, \dots, \beta_{2N}$. There are traditionally two different ways to find these functions. The first is stochastic Galerkin [33, 41], where each of the $12M + 7$ variables (for example, T) is expanded in orthogonal polynomials of the $\{\beta_m\}$. This expansion is truncated and then all of the coefficients in this truncated expansion are found by solving a very large set of coupled equations, often exploiting the orthogonality of the polynomials. (The form of the polynomials is determined by the probability density function of the random variables, the $\{\beta_m\}$. [42, 44]) Once the coefficients have been found, any quantity such as the expected value of, say u_0^0 , can be found by integrating over the space of random variables. Unfortunately, modifying code capable of solving (14.19–14.24) and (14.25) to find all coefficients in the expansion just mentioned is non-trivial.

The other common alternative is referred to as stochastic collocation [33, 41], which involves solving (14.19–14.24) and (14.25) at a number of different points in the random parameter space, i.e. using different $\{\beta_m\}$. We then have the value of all variables $u_0^0, u_1^0, \dots, a_{2M}^s, T$ at these different points and can use interpolation to estimate the values of these variables at other points in the random parameter space. If the values of $\{\beta_m\}$ at which (14.19–14.24) and (14.25) are solved are chosen appropriately, the solutions of these equations at these points can be used to estimate, for example, the expected value of u_0^0 very accurately. This method is referred to as “non-intrusive”, as it does not require modification of the code to solve (14.19–14.24) and (14.25), just some decisions about the values of $\{\beta_m\}$ to use, and some




Fig. 14.2 A typical distribution of T values, where 10,000 realisations of the $\beta_1, \dots, \beta_{2N}$ have been used. Parameters: $\sigma = 0.00003$, $b = 1$, $N = 3$, $M = 6$, $B = 0.1$, $\tau = 14$

postprocessing of the results. This method is also trivially parallelisable and is the one we use here.

The main variable we are interested in is $T(\beta_1, \dots, \beta_{2N})$, the period of the periodic solution, an example of which is shown in Fig. 14.1. A typical distribution of T , for 10,000 randomly-chosen $\{\beta_m\}$, is shown in Fig. 14.2. For these parameter values, the period when $\sigma = 0$ is approximately 124.4007, and the presence of the spatial noise always increases the period. To obtain the expected value of T (which we refer to as \bar{T}) we need to average $T(\beta_1, \dots, \beta_{2N})$ over $\beta_1, \dots, \beta_{2N}$. This average period will itself be a function of parameters of interest such as the strength of the random field, σ , and the correlation length b . Thus, knowing the distribution of the variables $\{\beta_m\}$, we want to calculate the $2N$ -dimensional integral

$$\frac{1}{2^{2N}} \int_{-1}^1 \dots \int_{-1}^1 T(\beta_1, \dots, \beta_{2N}) d\beta_1 \dots d\beta_{2N} \tag{14.38}$$

Note that having found \bar{T} it is equally easy to find, for example, the variance of $T(\beta_1, \dots, \beta_{2N})$:

$$V_T = \frac{1}{2^{2N}} \int_{-1}^1 \dots \int_{-1}^1 [T(\beta_1, \dots, \beta_{2N}) - \bar{T}]^2 d\beta_1 \dots d\beta_{2N} \tag{14.39}$$

We will evaluate these integrals in several different ways.

The integrals in (14.19–14.21) are over periodic domains, so the trapezoidal rule which we use converges very quickly as the number of points used increases [40]. We use 275 points in x and assume that this is accurate enough. We also set M , the number

of Fourier modes in time, to be $M = 6$, and do not consider varying this number further. We will investigate varying N , which determines the number of modes used to represent the random field $h(x)$, and the number of points used to approximate the integrals in (14.38). Note that the computational effort to evaluate (14.38) is proportional to M (and to the number of points used to evaluate the integrals in (14.19–14.21)) but grows extremely quickly with N , as this is proportional to the dimension of the space to be integrated over.

14.3 Results

14.3.1 Convergence

We first show some results regarding convergence of three different schemes for evaluating the integral (14.38). To be concrete we take b (the correlation length of the random field) to be 1. We truncate the series (14.37) at $N = 3$. Typical realisations of the $h(x)$ are shown in Fig. 14.3 (top) and the average covariance of 1000 realisations is shown in Fig. 14.3 (bottom). There is significant deviation between the theoretical and actual covariances, and this is mostly due to approximations made in analytically determining the eigenpairs of the covariance operator.

14.3.1.1 Monte Carlo

We wish to approximate the integral in (14.38). Firstly, consider Monte-Carlo integration. In this method we generate ν vectors K^i , $i = 1, \dots, \nu$, each of length $2N$, where each component of each vector is randomly and independently chosen from a uniform distribution on $[-1, 1]$. We then approximate the $2N$ -dimensional integral in (14.38) by the average

$$\frac{1}{\nu} \sum_{i=1}^{\nu} T(K_1^i, \dots, K_{2N}^i) \quad (14.40)$$

This method has the advantage that it is very simple, and will converge to the correct result as $\nu \rightarrow \infty$. Unfortunately, it is well-known that the error converges as $1/\sqrt{\nu}$ [17]. The convergence of this method is demonstrated in Fig. 14.4

14.3.1.2 Gaussian Quadrature: Full Grids

Next we consider using Gaussian quadrature, forming a tensor product of one-dimensional rules [2, 32, 33] in order to approximate the integral in (14.38). One-dimensional Gauss-Legendre quadrature involves approximating the integral




Fig. 14.3 *Top* four typical realisations of $h(x)$ for $b = 1$, given by (14.37). *Bottom* theoretical covariance of $h(x)$, given by (14.26) (dashed), and the average covariance of 1000 independent realisations (solid). Parameters: $\sigma = 0.00003$, $b = 1$, $N = 3$




Fig. 14.4 Error convergence. *MC* Monte-Carlo integration. For large ν the data points lie on a line with slope $-1/2$, as expected. *full* full tensor product. *sparse* sparse tensor product. ν is the number of $2N$ -dimensional vectors needed in the approximation of the integral in (14.38), i.e. the total number of distinct values of $T(\beta_1, \dots, \beta_{2N})$ that are used to approximate (14.38). Parameters: $\sigma = 0.00003, b = 1, N = 3, M = 6, B = 0.1, \tau = 14$

$$\frac{1}{2} \int_{-1}^1 f(x) dx \tag{14.41}$$

for sufficiently smooth functions f by the sum

$$\sum_{j=1}^{\widehat{N}} w_j f(x_j) \tag{14.42}$$

where x_j is the j th root of $P_{\widehat{N}}$, the \widehat{N} th Legendre polynomial (normalised so that $P_{\widehat{N}}(1) = 1$), and the weights w_j are given by

$$w_j = \frac{1}{(1 - x_j^2) \left[P'_{\widehat{N}}(x_j) \right]^2} \tag{14.43}$$

These rules can be used to approximate multi-dimensional integrals where the variable in each direction is uniformly distributed, in the obvious way. Figure 14.5 shows the tensor product in two spatial dimensions, and weights, for $\widehat{N} = 10$.

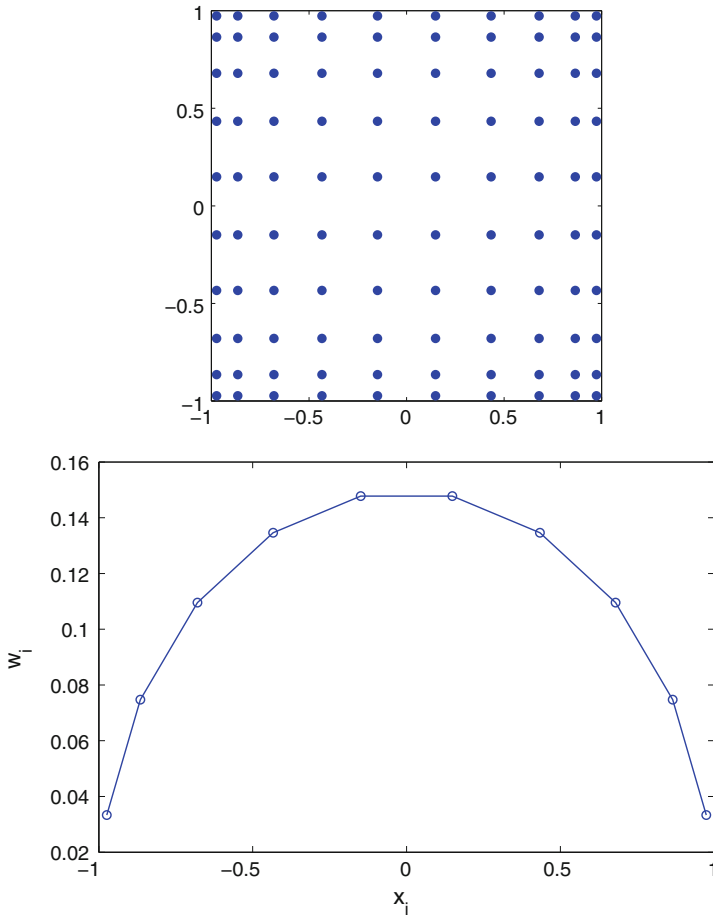


Fig. 14.5 Top tensor product in two dimensions, formed from two one-dimensional Gauss-Legendre rules. Bottom weights w_i as a function of the nodes x_i , for $\widehat{N} = 10$

For fixed \widehat{N} we approximate the integral (14.38) by the multiple sum

$$\sum_{j_1=1}^{\widehat{N}} \sum_{j_2=1}^{\widehat{N}} \cdots \sum_{j_{2N}=1}^{\widehat{N}} w_{j_1} w_{j_2} \cdots w_{j_{2N}} T(x_{j_1}, \dots, x_{j_{2N}}) \tag{14.44}$$

There are a total of $(\widehat{N})^{2N}$ terms in this multiple sum, which grows rapidly as a function of \widehat{N} for moderate to large N —this is the curse of dimensionality. Results using this method are shown in Fig. 14.4, where $\nu = (\widehat{N})^{2N}$. We see rapid convergence, as expected from a spectral method such as this [40]. However, the curse of

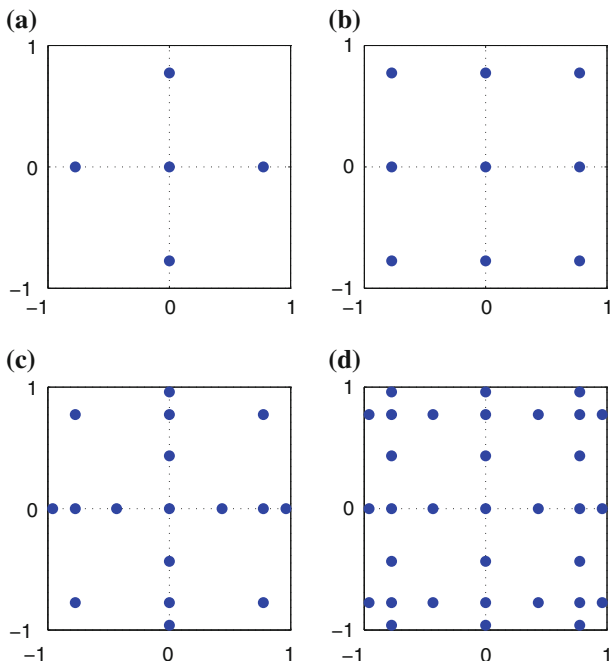


Fig. 14.6 Sparse tensor products in two spatial dimensions for (a)–(d): levels 2–5. The number of points is, respectively, 5, 9, 17 and 33

dimensionality makes this method infeasible for many problems. For example, if $N = 5$ and $\hat{N} = 5$, i.e. we use just five points in each of 10 random dimensions, we have $\nu \approx 10^7$.

14.3.1.3 Gaussian Quadrature: Sparse Grids

The third method we consider involves the use of sparse tensor grids [3, 19, 21, 32]. Tensor products are still formed, as in Fig. 14.5, but many of the points are then discarded, as they do not contribute significantly to the evaluation of the integral. For a specific spatial dimension, different “levels” of grids, and thus accuracies, are constructed. An example is shown in Fig. 14.6 for two spatial dimensions. (We use the code associated with [21], available at <http://www.sparse-grids.de/>) We do not present the general theory here but instead refer the reader to references above. For a given level of accuracy, sparse grids use fewer points than full tensor grids, and the advantage of using sparse grids as opposed to full increases as the dimension of the space to be integrated over increases. Figure 14.4 shows results from using sparse grids. We see that this method is the most accurate of the three considered, converging more rapidly than the full tensor product. We expect this advantage to increase as N , the number of dimensions integrated over, increases.




Fig. 14.7 Mean period, \bar{T} (top) and standard deviation, ($s \equiv \sqrt{V_T}$) (bottom) as a function of random field strength σ . Parameters: $b = 1$, $N = 6$

14.3.2 Varying Parameters

Having compared three common schemes for approximating the integral (14.38), we now use the most accurate one (sparse tensor products) to investigate the effects of varying parameters in the model. We first consider varying the “strength” of the random field, σ . To obtain specific results we keep the correlation length b at $b = 1$ and set $N = 6$. We set the sparse grid level to be 5, which means using a total of 11,073 points in the approximation of (14.38). The results are shown in Fig. 14.7, and we see that both the mean and standard deviation of the distribution of periods increases almost linearly with σ .

In Fig. 14.8 we vary the correlation length b . Because the coefficients of the random field $h(x)$ decay more slowly as b is decreased, we need to keep a large number of terms in the truncation (14.37) to accurately represent the field $h(x)$, i.e. we need to integrate over a high dimensional space. We set $N = 50$, which, for the level we choose, gives 20,001 points in the approximation of (14.38). We see




Fig. 14.8 Mean period, \bar{T} (top) and standard deviation, $s \equiv \sqrt{V_T}$ (bottom) as a function of random field correlation length b . Parameters: $\sigma = 3 \times 10^{-6}$, $N = 50$

that for these parameter values and the truncations used, the standard deviation of the distribution of periods decreases as b is increased, while the mean period shows a non-monotonic dependence on b .

Note that for this type of high-dimensional integration, full tensor grids are impossible to use. Even using sparse tensor grids, as above, is problematic, as the number of points used still grows very rapidly with the level used, and a large amount of time is spent actually calculating the grid points before they are used. However Monte Carlo methods are still feasible, as are other extensions of Monte Carlo methods such as Quasi-Monte Carlo (QMC) [12, 17, 26, 34]. QMC methods are similar to Monte Carlo in that the integrand is evaluated at many points and then averaged, but in QMC methods the points are not randomly chosen, but rather chosen in some “optimal” way. Many variations exist, and rather than go into details here we show in Fig. 14.9 a comparison between Monte Carlo and one particular QMC algorithm. We see that QMC does better than Monte Carlo, at least for these parameter values. The convergence rate for the error in the Monte Carlo method is known to scale as




Fig. 14.9 Error in calculating the mean period as a function of the number of integrand evaluations for Monte Carlo and Quasi-Monte Carlo methods (rule lattice-38005-1024-1048576.5000 from <http://web.maths.unsw.edu.au/~fkuo/lattice/index.html>). The *dashed line* has slope -0.50 and the *solid* has slope -0.68 . Parameters: $b = 0.1, \sigma = 10^{-6}, N = 50$

$1/\sqrt{\nu}$, and “randomised” QMC methods can be used to obtain error estimates for these types of method [34].

Figure 14.10 shows the same calculation as in the top panel of Fig. 14.8, but with QMC. The results are essentially identical, and if errorbars were plotted in Fig. 14.10 they would be smaller than the markers shown.

14.4 Other Forms of Heterogeneity

In the model (14.1 and 14.2) we included the random field inside the nonlinear function F , thinking of it as a spatial perturbation of the firing threshold. We now show how several other forms of heterogeneity could be dealt with using the ideas presented here.

14.4.1 Modulated Connectivity

Suppose that (14.1) was replaced by

$$\frac{\partial u(x, t)}{\partial t} = -u(x, t) + \int_0^{2\pi} G(x - y)[1 + h(y)]F[u(y, t) - a(y, t)]dy \quad (14.45)$$




Fig. 14.10 Mean period, \bar{T} as a function of random field correlation length b , calculated using the same QMC algorithm as in Fig. 14.9. Errorbars are smaller than the markers shown. Parameters: $\sigma = 3 \times 10^{-6}$, $N = 50$

as in [6, 14], where we can think of the new connectivity, $G(x - y)[1 + h(y)]$, as no longer being a function of $x - y$ only. The above analysis would go through, with (14.10–14.12) being the same, but (14.7–14.9) being replaced by

$$\frac{du^0}{dt} = -u^0 + 0.09 \int_0^{2\pi} [1 + h(x)]F[u^0 - a^0 + (u^c - a^c) \cos x + (u^s - a^s) \sin x] dx \tag{14.46}$$

$$\frac{du^c}{dt} = -u^c + 0.45 \int_0^{2\pi} [1 + h(x)]F[u^0 - a^0 + (u^c - a^c) \cos x + (u^s - a^s) \sin x] \cos x dx \tag{14.47}$$

$$\frac{du^s}{dt} = -u^s + 0.45 \int_0^{2\pi} [1 + h(x)]F[u^0 - a^0 + (u^c - a^c) \cos x + (u^s - a^s) \sin x] \sin x dx \tag{14.48}$$

respectively, with a corresponding modification of (14.19–14.21). Parametrising $h(x)$ as in (14.37) we can find $T(\beta_1, \dots, \beta_{2N})$ by solving this new set of equations as before, and this form of heterogeneity introduces no new complexity. Results are shown in Fig. 14.11. We see that as above, both the mean period and its standard deviation increase as σ is increased.




Fig. 14.11 Modulated connectivity. Mean period, \bar{T} (top) and standard deviation, $s = \sqrt{\bar{V}_T}$ (bottom) as a function of random field strength σ . Parameters: $b = 1, N = 6$

14.4.2 Modulated Drive

Suppose instead that (14.1) was replaced by

$$\frac{\partial u(x, t)}{\partial t} = -u(x, t) + \int_0^{2\pi} G(x - y)F[u(y, t) - a(y, t)]dy + h(x) \quad (14.49)$$

as originally proposed by Amari [1]. Using $h(x)$ as in (14.37) and writing

$$u(x, t) = U^0(t) + \sum_{i=1}^{\infty} [U_i^c(t) \cos(ix) + U_i^s(t) \sin(ix)] \quad (14.50)$$

and

$$a(x, t) = A^0(t) + \sum_{i=1}^{\infty} [A_i^c(t) \cos(ix) + A_i^s(t) \sin(ix)] \tag{14.51}$$

we see that all $U_i^c(t)$, $U_i^s(t)$, $A_i^c(t)$ and $A_i^s(t)$ will decay to zero if $i > N$, so we can truncate (14.50) and (14.51) at $i = N$. For $2 \leq i \leq N$ we have (after transients)

$$0 = -U_i^c + \beta_i \sqrt{\frac{\lambda_i^{(1)}}{\pi}} \quad \text{and} \quad 0 = -U_i^s + \beta_{N+i} \sqrt{\frac{\lambda_i^{(2)}}{\pi}} \tag{14.52}$$

and

$$0 = BU_i^c - A_i^c \quad \text{and} \quad 0 = BU_i^s - A_i^s \tag{14.53}$$

which can all be trivially solved. Now U^0 , U_1^c , U_1^s , A^0 , A_1^c and A_1^s satisfy

$$\frac{dU^0}{dt} = -U^0 + 0.09 \int_0^{2\pi} F[u(x, t) - a(x, t)] dx \tag{14.54}$$

$$\frac{dU_1^c}{dt} = -U_1^c + 0.45 \int_0^{2\pi} F[u(x, t) - a(x, t)] \cos x dx + \beta_1 \sqrt{\frac{\lambda_1^{(1)}}{\pi}} \tag{14.55}$$

$$\frac{dU_1^s}{dt} = -U_1^s + 0.45 \int_0^{2\pi} F[u(x, t) - a(x, t)] \sin x dx + \beta_{N+1} \sqrt{\frac{\lambda_1^{(2)}}{\pi}} \tag{14.56}$$

$$\tau \frac{dA^0}{dt} = BU^0 - A^0 \tag{14.57}$$

$$\tau \frac{dA_1^c}{dt} = BU_1^c - A_1^c \tag{14.58}$$

$$\tau \frac{dA_1^s}{dt} = BU_1^s - A_1^s \tag{14.59}$$

This set of equations is no more complex than (14.7–14.12) and can be solved the same way. The results of varying the random field strength σ are shown in Fig. 14.12. Comparing with Figs. 14.7 and 14.11 we see qualitatively the same behaviour: increasing the heterogeneity both slows the bump and increases the width of the distribution of periods.




Fig. 14.12 Modulated Drive. Mean period, \bar{T} (top) and standard deviation, $s = \sqrt{V_T}$ (bottom) as a function of random field strength σ . Parameters: $b = 1, N = 6$

14.5 Conclusion

In this chapter we have used ideas from the field of Uncertainty Quantification to investigate the effects of spatial heterogeneity on the speed of a moving “bump” of activity in a neural field model. Neural field models are intrinsically infinite-dimensional, as is the spatially-extended “frozen noise” that we included in the model. In order to make computational progress we need to represent these processes in a finite-dimensional way. The form of the coupling function G (constant plus cosine) allowed us to exactly write the neural field dynamics (after transients) in the form of six coupled ODEs (14.7–14.12), thus making the spatial part of the dynamics finite-dimensional. Any other spatially-periodic coupling function could be represented arbitrarily well by a finite number of similar spatial modes via a Fourier series, resulting in a similar set of ODEs. These ODEs can be approximately solved in any number of ways, but collocation, as used here, is very efficient. It should be emphasised that the type of solution we were interested in, namely a moving bump,

meant that we were interested in periodic solutions of these ODEs. If we were interested in, for example, a moving front [5, 10], we would look for different sorts of solutions. A specific example of such a front moving over a heterogeneous domain is given in Sect. 4.1 of [15]. And as is standard, the noise process is approximated in a finite-dimensional way by truncating the Karhunen-Loève decomposition (14.28).

Finding the expected value of a quantity (in this case, the period of a periodic orbit) in a system with stochastic or uncertain parameters is equivalent to averaging over a multi-dimensional space. For a small to moderate dimensions such an integral can be performed using full or sparse tensor product grids [32], but for high-dimensional integrals techniques such as Quasi-Monte Carlo [17] must be used. We have demonstrated each of these methods and found several interesting results. For the parameters studied, adding spatial noise to the system always slows the moving bump. Also, varying the spatial scale of the noise shows a nonmonotonic response of the bump's speed (Figs. 14.8 and 14.10).

The results presented here are computationally intensive, and the field of uncertainty quantification can only benefit from both continuing increases in computational power and continued theoretical advances.

Conflict of Interest

The author declares that they have no conflict of interest.

References

1. Amari, S.: Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol. Cybern.* **27**(2), 77–87 (1977)
2. Babuška, I., Nobile, F., Tempone, R.: A stochastic collocation method for elliptic partial differential equations with random input data. *SIAM Rev.* **52**(2), 317–355 (2010)
3. Barthelmann, V., Novak, E., Ritter, K.: High dimensional polynomial interpolation on sparse grids. *Adv. Comput. Math.* **12**(4), 273–288 (2000)
4. Ben-Yishai, R., Bar-Or, R.L., Sompolinsky, H.: Theory of orientation tuning in visual cortex. *Proc. Nat. Acad. Sci.* **92**(9), 3844–3848 (1995)
5. Brackley, C.A., Turner, M.S.: Random fluctuations of the firing rate function in a continuum neural field model. *Phys. Rev. E* **75**(4), 041913 (2007)
6. Bressloff, P.C.: Traveling fronts and wave propagation failure in an inhomogeneous neural network. *Phys. D: Nonlinear Phenom.* **155**(1), 83–100 (2001)
7. Bressloff, P.C.: From invasion to extinction in heterogeneous neural fields. *J. Math. Neurosci.* **2**, 6 (2012)
8. Bressloff, P.C.: Spatiotemporal dynamics of continuum neural fields. *J. Phys. A: Math. Theor.* **45**(3), 033001 (2012)
9. Bressloff, P.C., Folias, S.E.: Front bifurcations in an excitatory neural network. *SIAM J. Appl. Math.* **65**(1), 131–151 (2004)
10. Bressloff, P.C., Webber, M.A.: Front propagation in stochastic neural fields. *SIAM J. Appl. Dyn. Syst.* **11**(2), 708–740 (2012)

11. Bressloff, P.C., Wilkerson, J.: Traveling pulses in a stochastic neural field model of direction selectivity. *Front. Comput. Neurosci.* **6**(90), 14 (2012)
12. Cools, R., Kuo, F.Y., Nuyens, D.: Constructing embedded lattice rules for multivariate integration. *SIAM J. Sci. Comput.* **28**(6), 2162–2188 (2006)
13. Coombes, S.: Waves, bumps, and patterns in neural field theories. *Biol. Cybern.* **93**(2), 91–108 (2005)
14. Coombes, S., Laing, C.R.: Pulsating fronts in periodically modulated neural field models. *Phys. Rev. E* **83**(1), 011912 (2011)
15. Coombes, S., Schmidt, H., Laing, C.R., Svanstedt, N., Wyller, J.A.: Waves in random neural media. *Discrete Contin. Dyn. Syst.* **32**, 2951–2970 (2012)
16. Coombes, S., Venkov, N.A., Shiau, L., Bojak, I., Liley, D.T.J., Laing, C.R.: Modeling electrocortical activity through improved local approximations of integral neural field equations. *Phys. Rev. E* **76**(5), 051901 (2007)
17. Dick, J., Kuo, F.Y., Sloan, I.H.: High-dimensional integration: The quasi-Monte Carlo way. *Acta Numerica* **22**, 133–288, 5 (2013)
18. Folias, S.E.: Nonlinear analysis of breathing pulses in a synaptically coupled neural network. *SIAM J. Appl. Dyn. Syst.* **10**(2), 744–787 (2011)
19. Gerstner, T., Griebel, M.: Numerical integration using sparse grids. *Numer. Algorithms* **18**(3), 209–232 (1998)
20. Ghanem, R., Spanos, P.D.: *Stochastic Finite Elements: A Spectral Approach*. Springer, New York (1991)
21. Heiss, F., Winschel, V.: Likelihood approximation by numerical integration on sparse grids. *J. Econometrics* **144**(1), 62–80 (2008)
22. Huang, S.P., Quek, S.T., Phoon, K.K.: Convergence study of the truncated Karhunen-Loeve expansion for simulation of stochastic processes. *Int. J. Numer. Methods Eng.* **52**(9), 1029–1043 (2001)
23. Kilpatrick, Z.P., Ermentrout, B.: Wandering bumps in stochastic neural fields. *SIAM J. Appl. Dyn. Syst.* **12**(1), 61–94 (2013)
24. Kilpatrick, Z.P., Folias, S.E., Bressloff, P.C.: Traveling pulses and wave propagation failure in inhomogeneous neural media. *SIAM J. Appl. Dyn. Syst.* **7**(1), 161–185 (2008)
25. Kuehn, Christian, Riedler, Martin: Large deviations for nonlocal stochastic neural fields. *J. Math. Neurosci.* **4**(1), 1 (2014)
26. Kuo, F.Y., Sloan, I.H.: Lifting the Curse of Dimensionality. *Notices of the AMS*, **52**(11):1320–1328 (2005)
27. Laing, C., Coombes, S.: The importance of different timings of excitatory and inhibitory pathways in neural field models. *Netw.: Comput. Neural Syst.* **17**(2), 151–172 (2006)
28. Laing, C.R., Longtin, A.: Noise-induced stabilization of bumps in systems with long-range spatial coupling. *Phys. D: Nonlinear Phenom.* **160**(3), 149–172 (2001)
29. Laing, C.R., Troy, W.C.: PDE methods for nonlocal models. *SIAM J. Appl. Dyn. Syst.* **2**(3), 487–516 (2003)
30. Laing, C.R., Frewen, T.A., Kevrekidis, I.G.: Coarse-grained dynamics of an activity bump in a neural field model. *Nonlinearity* **20**(9), 2127 (2007)
31. Laing, C.R., Troy, W.C., Gutkin, B., Ermentrout, G.B.: Multiple bumps in a neuronal model of working memory. *SIAM J. Appl. Math.* **63**, 62 (2002)
32. Laing, C., Zou, Y., Smith, B., Kevrekidis, I.: Managing heterogeneity in the study of neural oscillator dynamics. *J. Math. Neurosci.* **2**(1), 5 (2012)
33. Le Maître, O.P., Knio, O.M.: *Spectral Methods for Uncertainty Quantification*. Springer, New York (2010)
34. Mike, G., Kuo, F.Y., Sloan, I.H., Waterhouse, B.J.: Quasi-Monte Carlo for finance applications. *ANZIAM J.* **50**, C308–C323 (2008)
35. Owen, M.R., Laing, C.R., Coombes, S.: Bumps and rings in a two-dimensional neural field: splitting and rotational instabilities. *New J. Phys.* **9**, 378 (2007)
36. Pinto, D.J., Ermentrout, G.B.: Spatially structured activity in synaptically coupled neuronal networks: I. traveling fronts and pulses. *SIAM J. Appl. Math.* **62**(1), 206–225 (2001)

37. Rankin, J., Meso, A.I., Masson, G.S., Faugeras, O., Kornprobst, P.: Bifurcation study of a neural field competition model with an application to perceptual switching in motion integration. *J. Comput. Neurosci.* **36**(2), 193–213 (2014)
38. Shardlow, Tony: Numerical simulation of stochastic PDEs for excitable media. *J. Comput. Appl. Math.* **175**(2), 429–446 (2005)
39. Steyn-Ross, M.L., Steyn-Ross, D.A., Sleigh, J. W., Whiting, D.R.: Theoretical predictions for spatial covariance of the electroencephalographic signal during the anesthetic-induced phase transition: Increased correlation length and emergence of spatial self-organization. *Phys. Rev. E* **68**, 021902 (2003)
40. Trefethen, L.N.: *Spectral methods in MATLAB*. Society for Industrial Mathematics, vol. 10 (2000)
41. Xiu, D.: Fast numerical methods for stochastic computations: a review. *Commun. Comput. phys.* **5**(2–4), 242–272 (2009)
42. Xiu, D.: *Numerical methods for stochastic computations: a spectral method approach*. Princeton University Press (2010)
43. Xiu, D., Hesthaven, J.S.: High-order collocation methods for differential equations with random inputs. *SIAM J. Sci. Comput.* **27**(3), 1118–1139 (2005)
44. Xiu, D., Karniadakis G.E.: The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM J. Sci. Comput.* **24**(2), 619–644 (2002)
45. Xiu, D., Kevrekidis, I.G., Ghanem, R.: An equation-free, multiscale approach to uncertainty quantification. *Comput. Sci. Eng.* **7**(3), 16–23 (2005)

Chapter 15

In-Silico Models of Trabecular Bone: A Sensitivity Analysis Perspective

Marlène Mengoni, Sebastien Sikora, Vinciane d'Otreppe,
Ruth Karen Wilcox and Alison Claire Jones

Abstract This chapter provides an overview from a sensitivity analysis perspective of computational mechanical modeling of trabecular bone, where models are generated from Computed Tomography images. Specifically, the effect of model development choices on the model results is systematically reviewed and analyzed for both micro-Finite Element and continuum-Finite Element models. Particular emphasis is placed on the image processing effects (thresholding, down-sampling, image to material properties relationships), the mesh-related aspects (mesh size, element type), and the computational representation of the boundary conditions. Typical issues are highlighted and recommendations are proposed with respect to various model applications, including global stiffness/strength and local failure stress/strain behavior.

Keywords Image-based FE models · Trabecular bone · Boundary conditions · Image processing · Numerical convergence

15.1 Introduction

Biological processes such as bone and soft tissue remodeling are triggered or influenced by the local mechanical environment through the cells' mechanotransduction [9, 57]. An accurate model of that environment is therefore important to capture the coupled mechano-biological response of biological tissues. In particular, this chapter will focus on the structural modelling of trabecular bone, which is a driver in applications such as bone adaptation and repair [8, 41]. There is wide-spread interest in replicating the behavior of trabecular bone within a computational environment.

M. Mengoni (✉) · S. Sikora · R.K. Wilcox · A.C. Jones
Institute of Medical and Biological Engineering, School of Mechanical Engineering,
University of Leeds, Leeds, UK
e-mail: m.mengoni@leeds.ac.uk

M. Mengoni · V. d'Otreppe
Department of Aerospace and Mechanical Engineering, Non-Linear
Computational Mechanics (LTAS-MN²L), University of Liege, Liège, Belgium




Fig. 15.1 Trabecular core of a deer antler (2.11 mm cubic specimen): **a** 3D visualization obtained from micro-CT images, **b** hexahedral micro-FE mesh (195,862 elements), and **c** hexahedral continuum-FE mesh (27 elements)

The improvement of bone strength estimation is a key clinical driver, which would enhance treatment of patients whose bones are weakened, such as those with osteoporosis. In parallel, researchers are developing detailed theoretical models of bone structure, material properties and macroscopic behavior in order to better understand its fracture mechanics and regeneration mechanisms.

The development of these patient-specific, highly detailed computer models of trabecular bone has been made possible by high-resolution, bench-top imaging systems such as micro-Computed Tomography (micro-CT) scanners. The three-dimensional images generated by a micro-CT scanner can be processed to extract geometric details at a macroscopic and microscopic level as well as maps of how the material properties vary throughout the bone.

This chapter reviews two of the dominant approaches to image-based modeling of trabecular bone. These are micro-Finite Element (micro-FE) models, which explicitly represent the micro-structure of the bone, and continuum-Finite Element models, which use a continuous, inhomogeneous material property field to implicitly represent that micro-structure (see Fig. 15.1).

Any theoretical representation of a physical system is built upon a series of assumptions about the behavior of that system. Finite element models of trabecular bone by necessity use boundary conditions and loading regimes which are either pure assumption or approximations of experimental conditions [42]. In addition the geometry, and in some cases the material properties, are derived from image data using theoretical relationships, calibrated conversion values, or user-controlled processes.

This chapter is specifically concerned with the sensitivity of key trabecular bone measures to the variation caused by modeling assumptions and decisions made during the derivation of geometries and material properties. Analysis of these sensitivities is a fundamental part of the model development process [18]. Data from the literature is reviewed and specific cases from the authors' work are used to provide more detailed examples.

15.1.1 Micro-Finite Element Models

Trabecular bone models which fall into the category of micro-FE explicitly include individual trabecular struts resolved as part of the three-dimensional model geometry. The element size within the mesh is small enough for several elements to span the thickness of a typical trabecular strut. The marrow and blood within the bone are generally excluded from the finite element mesh.

Due to the large number of elements and complex geometries, micro-FE models of trabecular bone typically require a large amount of computing resource and are often analyzed on high performance computers. With the advances in imaging techniques as well as computational power, micro-FE models have emerged since the 1990s. The high computational cost however generally restricts the size of samples which are processed, a typical area of bone represented is $(5\text{--}10\text{ mm})^3$ [50].

In some cases the trabecular bone within a micro-FE model may be assigned homogeneous material property constants across all elements [55]. In other cases the material properties are assigned on an element-by-element basis, using information from a micro-CT image [5, 25].

In order to construct a subject-specific micro-FE model of trabecular bone a three-dimensional image must be available with a resolution small enough to capture the trabecular geometry. Although bench-top micro-CT scanners can generate sufficiently high-resolution images, they are limited to relatively small in vitro specimens: a whole human vertebra can currently be scanned but generally not a complete human femur. In vivo imaging is possible through the use of high resolution peripheral Quantitative Computed Tomography (HR-pQCT) which makes it possible to capture trabecular level images of small peripheral human joints.

15.1.2 Continuum Finite Element Models

Continuum finite element models of trabecular bone use an element size which is too large for individual trabeculae to be resolved. A single element will typically cover an area which is large enough to contain several trabecular struts and the marrow space between them. Since fewer elements are needed, whole bone models can be analyzed at relatively low computational cost.

The trabecular structure within each element can be represented by a separate material property definition, creating a continuous but inhomogeneous map of properties throughout the bone. The sophistication and accuracy of the element-specific material models depends on both the source image and the modeling approach. In many cases the material properties are isotropic [22]. However, information on trabecular directionality has been used to derive orthotropic material properties for each element [37, 58].

15.1.3 Trabecular Bone Modeling Applications

Both the micro- and the continuum-FE modeling methods are capable of replicating apparent load-displacement behavior of trabecular bone, equivalent to the measurements obtained from a materials testing machine. In the case of continuum-FE, stiffness and strength of whole bones can be modeled [22, 27, 32, 44]. However, the use of a relatively low mesh density, which does not capture individual trabecular struts, limits the possible outputs to those at a macro level. Since micro-FE models are capable of capturing deformation of individual trabecular members, their outputs can include stress, strain and failure initiation within the microstructure.

The development of image-based computer modeling of bone has largely been driven by the need to assess the integrity of bone in patients with a suspected fracture risk. This may be to diagnose bone weakening conditions such as osteoporosis or metastatic involvement, or to measure the effectiveness of on-going treatment [3, 24, 48]. Models designed for this purpose generally use continuum material properties due to the low resolution of *in vivo* scanning. The apparent stiffness and strength measurements taken from these models have been compared to traditional DEXA scans in terms of reliability in predicting fracture [7].

The continuum-FE method has also been used in pre-clinical research setting to compare the effect of treatment across a set of specimens [54]. High resolution micro-CT source images may be available in the laboratory environment and can enhance the accuracy of continuum-FE models while the computational cost of the model solution remains low.

Where the high resolution imaging was available, micro-FE modeling has been developed in order to analyze micro-mechanics and detailed damage mechanics of trabecular bone [30, 36]. The ability to model deformation at the level of individual trabecular struts has led to the use of micro-FE as the mechanical driver for bone remodeling prediction [41] where the local strain/strain energy field drives the remodeling algorithm.

Initial development of the micro-FE method used images of *in vitro* specimens. With the development of peripheral quantitative computed tomography (pQCT), micro-FE has now been applied to *in vivo* studies [6, 49, 50].

15.1.4 Sensitivity Analysis

Sensitivity analysis is the process of establishing how sensitive the outputs of a model are to various inputs or settings [1]. This is done by varying the input within a range considered reasonable or realistic, and measuring the effect on the model outputs and on any conclusions drawn from model comparisons. A sensitivity test will give information on how precise the value of a parameter should be but will not show how accurate it is.

For the trabecular bone models discussed in this chapter, the assumptions which could introduce errors into the solution can be grouped as follows.

1. Digital representation: the accuracy with which the imaging modality represents the real bone.
2. Geometry and mesh generation: the accuracy with which the source image is segmented and the bone geometry is represented by the mesh.
3. Mesh quality: the numerical effect of element size, element shape and order of integration.
4. Material model: the accuracy with which the material model represents the material mechanical behavior.
5. Boundary conditions and constraints: the effect of assumptions made about behavior at the model boundaries.
6. Loading regime: accuracy of representation of the load in the simulated scenario.

Both the geometry (2) and the material model (4) are central to the development of a specimen-specific model and can be dependent on the source image.

In contrast, the boundary conditions (5) and loads (6) are usually independent of the source image and aim to replicate some in vivo or in vitro scenario. This chapter will review the effect of boundary condition choices but not that of loading cases.

The finite element solution depends in the mesh quality (3). A mesh sensitivity analysis is one of the fundamental verification steps highlighted in any course on finite element analysis. The choice of element type in any finite element model is not always clear cut. In solid models, the common types are linear or higher order elements of tetrahedral or hexahedral shape. There is often a balance between using a higher order integration function, which in some cases will converge with fewer elements but have more integration points (and hence a higher computational cost) and a linear option where a larger number of elements may be needed to reach convergence.

Sections 15.2 and 15.3 give some details of the methodology employed for micro-FE and continuum-FE respectively, along with evidence of sensitivity of key outputs to the assumptions detailed above. Although the image resolution is mentioned in the context of image segmentation, digital representation (1) is otherwise neglected. The bone material model (4) is assumed throughout the chapter to follow an isotropic Hookean elastic behavior. The effect of using a, possibly more accurate, nonlinear and/or anisotropic material model is not considered in this work.

15.2 Micro-FE Models of Trabecular Bone

There are two main methods to create micro-FE meshes from images. The first method requires a triangulation of the surface that first needs to be extracted from the images [13, 35]. The triangulated surface can then be filled with tetrahedral elements. The second one, referred to as the voxel-based method, creates the elements by converting the images voxels into hexahedral elements [47, 51]. This second

method can be used either to directly create hexahedra from voxels or to create hexahedral meshes that are mass-compensated [47]. Indeed, direct conversion of voxels to hexahedra can lead to loss of connectivity. These disconnected elements are not an active part of the model and are thus disregarded by any computation method. This induces a loss of bone mass compared to the actual bone mass of the specimen. The mass-compensated method accounts for this loss of connectivity and adds mass by artificially thickening the remaining bone trabeculae.

The voxel-based method is straightforward; however, if not smoothed, it produces jagged surfaces and edges, known as a staircase artifact, and thus is not accurate at the boundaries.

The meshes produced by a surface extraction method are smoother than voxel-based meshes; however, tetrahedral elements do not perform as well as hexahedral elements of the same order from a computational point of view [10, 39, 46].

When used in a small strain analysis with uniform elastic properties, the voxel-based meshes (using 8-noded elements) are particularly efficient as all elements are the same (same fixed orientation and shape), allowing an elementary tangent operator to represent the entire linear system [40, 52]. In this case, models with millions of degrees-of-freedom can be resolved on standard desktop computers.

15.2.1 Sensitivity to Imaging and Material Property Assignment

Before 3D CT images are converted into micro-FE models, they are usually binarized by thresholding, segmenting marrow and bone. The threshold level and method influence the results of micro-FE models as they influence the mesh and lose information on the partial volume voxels (voxels representing both bone and marrow). Another method of producing the micro-FE models is to directly convert the grey level into equivalent mechanical properties, thus accounting in the model for both bone and marrow (and partial volume voxels). The greyvalue to mechanical properties relationship thus influences the micro-FE results.

15.2.1.1 Sensitivity to the Threshold

The image threshold levels and the threshold method are two of the key points affecting the model behavior. They determine whether a voxel is represented or not as a solid element in the model. The threshold method thus has a direct influence on the bone volume modeled, as well as on the trabecular thickness and trabecular connection. Changes in these three parameters affect the apparent behavior of a trabecular model, since a high bone mass sample will be stiffer than a low bone mass one of the same size. Equally, local behavior will be affected. For example thin trabeculae are more likely to fracture and unconnected trabeculae do not participate in the weight bearing of the bone sample.

Table 15.1 Deviation of structural parameters for a 0.5% increase in threshold level; data from [20]

Deviation in	Low volume fraction samples (<0.15) (%)	High volume fraction samples (>0.2) (%)
BV/TV	5	2
Tb.Th	<3	<3
Computed apparent stiffness	9	3

Early on in the introduction of micro-FE models, the effect of the threshold method was analyzed. By comparing the behavior of human trabecular bone models at different image resolutions [47], it was shown that the use of a direct voxel conversion produces good results for high resolution models but not for lower resolution ones (with however a better apparent behavior representation than the local behavior one). The use of a mass compensated method produces more accurate results on an apparent point of view, however it compensates the loss of connectivity by a thickening of the remaining structure and thus significantly changes the micro architecture, and the local behavior.

The threshold level used for trabecular bone is also significant. Manual thresholding by different users leads to different results. Even though there is a low inter-user variability (0.5% difference in threshold value) to produce a “judged as optimal” threshold [20], the change in structural parameters extracted within that variability can be significant (see Table 15.1). It has been shown [4] that visual thresholding usually under-estimates the bone volume. Due mainly to the difference in the thickness of the trabeculae, the sensitivity of a visual threshold is lower for high volume fraction samples [4]. It is likely that the error is systematic per user [20]. This systematic error would therefore not be an issue for a comparative study performed by a single user as differences between groups could be detected anyway. It cannot however be blindly used to extract absolute quantitative mechanical parameters. Choosing a threshold value that accounts for an experimentally measured BV/TV (e.g. measured with Archimedes’ principle) would reduce the errors associated to manual thresholding. However, such an experimental value can prove difficult to measure. Indeed, the specimen whose density is measured using Archimedes’ principle needs to be completely immersed into distilled water, or another submersion liquid of known density, and degassed to remove all trapped air [16]. The measured density thus depends on the reliability of the degassing phase which is not easy to assess.

It should be pointed out that these different studies [4, 20, 47] did not account for geometrical nonlinearities that could occur even at low apparent strains. Their conclusions over the representation of mechanical parameters are thus valid only under a small strain hypothesis.

15.2.1.2 Sensitivity to the Relationship Used

Most micro-FE models of trabecular bone consider a homogeneous tissue-level bone modulus. However, the tissue modulus is dependent on the mineralization and thus can vary both within a trabecula and between struts [29]. Using a density-dependent modulus can thus account for the mineral content but also for the partial volume effects at the trabecular surfaces. This partial volume effect is caused by the error in capturing the surface of the trabeculae. Depending on the threshold value, voxels representing a mix of bone tissue and air or marrow can be considered as being 100% bone or 100% air/marrow. The material properties of those areas are thus either over-estimated or under-estimated. Using material properties function of a local greyscale rather than homogeneous values after thresholding can thus help to reduce the sensitivity of the model to the threshold. With a non-linear relationship, Homminga et al. [25] showed that using a greyscale-based modulus value for each voxel instead of a homogeneous value reduces the mean deviation and the range of deviation of the computed apparent elastic modulus from the experimental apparent elastic modulus. Bourne et al. [5], proposed a linear relationship between the X-ray attenuation and tissue modulus. This relationship assumes a modulus of 20 GPa for a tissue of 1.1 g/cc. Using different slopes for the linear relationship, they found a slope of 1.4 most precisely predicted experimental modulus. They demonstrated the apparent elastic modulus value for a homogeneous 20 GPa model was significantly greater than for all types of inhomogeneous models. Following the same principle, Harrison et al. [21] calibrated their linear relationship with micro-indentation tests. Finally, a more complex relationship accounting for the mineral content of the bone was proposed by Bourne et al. [5]. They defined a theoretical relationship relating micro-CT mineral density to tissue density and elastic modulus. The derivation uses prior knowledge of the bone constituents' volume fractions and individual constituent densities to calculate the volume and mass of each constituent within a voxel.

15.2.2 Sensitivity to the Finite Element Mesh

To perform a finite element analysis on processed images, the geometry they represent needs to be discretized into a finite element mesh. A finite number of geometrically simple elements (such as hexahedra and tetrahedra) is used to represent the potentially very complex geometry represented in the images. The mesh built on the processed images is the next source of errors to which the results are very sensitive. The element size needs to be appropriate so that the geometrical discretization is accurate. However, a series of numerical errors occurs depending on the element size, type, and integration method. The sensitivity of the results to these types of errors is overviewed here.

15.2.2.1 Sensitivity to the Finite Element Mesh Size

Early work on the influence of element size [31] showed that computing an apparent stiffness with a linear model and a voxel-based mesh produced results which were very sensitive to the mesh resolution. However, the image resolution those meshes were built from (given the trabeculae are properly resolved) did not seem to play an important role. Low density samples seemed more sensitive to the mesh resolution than high density ones. Convergence studies for a fully linear model and a voxel-based mesh [19] have shown that the element size should be less than one fourth of the trabecular thickness. No convergence studies for meshes built on triangulated surfaces of trabecular structures were however found in the literature. This gap is partially covered hereafter.

Micro-CT images of the trabecular core of a bone antler [13] were used to analyze several sensitivity aspects of micro-FE models. It consisted of a 2.11 mm cubic specimen (BV/TV of 10.88 %), imaged at a cubic voxel size of 8.64 μm (see Fig. 15.1a). Ten triangulation surface meshes obtained with a surface reconstruction algorithm [13] were constructed at different resolutions (producing from 65,000 to 120,000 triangles). Linear tetrahedral meshes built on those surface triangulations (i.e. meshes with 130,000–290,000 elements) were used to analyze the effect of mesh size on a finite strain model of compression tests. The material at trabecular level was considered as following an isotropic Hookean elastic behavior, described with a Young's modulus of 7.78 GPa [14], and a Poisson's ratio of 0.3. The performance of those meshes was evaluated by comparing the computed apparent stiffness (computed with a linear regression at 0.2 % compression), and the force level for 5 % compression tests in each direction. The finite strains micro-FE models were solved using the non-linear object-oriented implicit software Metafor (LTAS-MN²L, University of Liège, Belgium).

The coarser meshes show places (highlighted in red on Fig. 15.2) where only one element spans across the trabecular thickness. The convergence study shows (Fig. 15.3) that both the apparent stiffness and the maximal force decrease when increasing the number of elements. The apparent stiffness decreases by 6.2 % from the coarser mesh to the finer ones, while the maximal force decreases by 10.4 %.

15.2.2.2 Sensitivity to Type of Finite Element

The performance of a model is not only sensitive to the mesh resolution but also to the shape of each element. Early studies on the subject [47] showed that for a fully linear elastic model, there were no significant differences between the performances of a mass-compensated linear hexahedral mesh and a tetrahedral one. As soon as geometrical nonlinearities are included in a model, those conclusions may no longer be valid. To analyze the element shape influence on a model performance, the bone antler sample introduced previously was meshed using three different algorithms. The first one used a surface-reconstruction algorithm as presented in Fig. 15.2 (we use here the finest mesh from that study), thus producing a smooth tetrahedral mesh; the





Fig. 15.2 Cut through antler meshes. *Left* 130,000 elements mesh, *Right* 290,000 elements mesh

Fig. 15.3 Convergence study—normalized values for compression tests in three directions (*color dots*): stiffness (*dashed line*) and force (*plain line*) (*color online*)



second one used a direct voxel conversion algorithm, thus producing a jagged hexahedral mesh (see Fig. 15.1b); the third one used a direct voxel conversion followed by a topology-preserving smoothing algorithm [12], producing a smooth hexahedral mesh. Two meshes were produced using this last method, the first one involved one smoothing iteration, the second one two smoothing iterations. The performance of these four meshes was evaluated comparing the computed apparent stiffness (computed with a linear regression at 0.2% compression), the maximal force level, and the deformed micro-structure for 10% compression tests in each direction, using the same material model as earlier.

Even though the smoothing algorithm preserves the initial topology as accurately as possible, some shrinkage is inevitable. The bone volume represented by each mesh is slightly different (Table 15.2). Smooth hexahedral meshes represent a smaller bone volume than the voxel-based mesh, the tetrahedral mesh represents however a similar bone volume. The computed apparent stiffness slightly decreases with smoothing in the hexahedral meshes. The computed force (Fig. 15.4) is equivalent for the three

Table 15.2 Mesh type dependent parameters

Relative difference (in % of the voxel mesh value)	Smooth hexahedral mesh (1 smoothing iteration)	Smooth hexahedral mesh (2 smoothing iterations)	Tetrahedral mesh
Volume	-4.32	-7.89	+0.75
Stiffness	-3.71 (± 1.72)	-5.55 (± 0.52)	+39.63 (± 7.05)

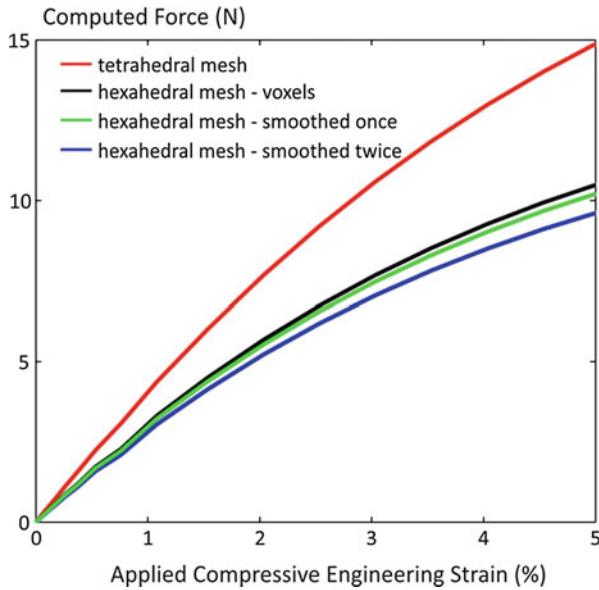


Fig. 15.4 Example of force versus applied engineering strain function of the element type

hexahedral meshes, slightly decreasing as the smoothing increases. The tetrahedral mesh requires however a higher force to be applied, especially at large strains.

Finally, the deformation pattern between the hexahedral meshes shows only slight differences while that of the tetrahedral mesh is completely different (Fig. 15.5). In particular, direction of trabecular bending can be opposite.

It is thus clear that the choice of element shape for a given mesh resolution influences not only the local behavior of the model but also its apparent behavior. A thorough comparison with experimental data on both the apparent and local level is needed to fully evaluate the best choice of element type. It should finally be pointed out that the conclusions addressed here are only valid on a sample with low BV/TV and considered as following an isotropic Hookean elastic behavior. Differences in results for the force level or the deformation pattern may be different for other types of material behavior more appropriate to model bone trabeculae at large strains.

For a given element shape and size, the chosen integration method will affect the integration results. Geometrical discretization errors lead generally to softening




Fig. 15.5 von Mises equivalent stress fields and values of the applied force at 10% applied engineering strains

(due to an underestimation of the volume at the threshold and the meshing phases). For hexahedral elements in a linear integration (8 integration points), there can be a stiffening effect due to shear locking if elements happen to be submitted to pure bending. In a finite element solution process the stress field, and all other secondary fields, are computed accurately at the integration points, and reducing their number reduces the accuracy of the stress field. A reduced integration (one integration point) avoids shear-locking but degrades the computed stress field. Quadratic integration (27 integration points) also avoids shear locking and does not degrade the secondary fields (as the number of integration points increases). It thus can be used even at low resolutions for accurate stress fields. The main disadvantages of a quadratic integration are the increase in computational cost and the increase in sensitivity in inaccurate geometry (such as staircase artifacts). The opposite behavior between geometrical discretization softening and integration stiffening explains [13] why 8-noded hexahedra, using a sufficient resolution, are accurate concerning the computation of global apparent values in small strains analysis. For a quadratic integration however, as discretization errors are not compensated by integration errors, the computed apparent values are less accurate than 8-noded hexahedra (even though both results are strongly correlated).

15.2.2.3 Discussion

A number of studies have been performed on the analysis of hexahedral mesh performances representing trabecular bone microstructure. However voxel-based meshes, while straightforward to build, do not represent the trabecular surface accurately as they produce jagged edges. Tetrahedral meshes allow the representation of smooth surfaces more easily. There are few convergence studies on the performance of tetrahedral meshes representing trabecular microstructure. The current work illustrated

that the mesh has to be fine enough in order for more than one element to span across the trabecular thickness. However, drawing a definite conclusion on the number of elements needed over the trabecular thickness is difficult. Indeed as all elements do not have the same size, measuring the quality of the mesh by the ratio of the mean edge length to the mean trabecular thickness is not representative of the mesh quality. For the presented meshes, that ratio indeed varies from 3.04 to 3.47 when increasing the number of elements. Looking at that ratio only would thus not highlight the differences observed between meshes. Only a local inspection of the mesh can help determine whether or not the mesh is fine enough.

Another method for avoiding jagged surfaces is to smooth out voxel-based meshes, thus obtaining smooth hexahedral meshes. This smoothing operation has to be done with as little volume loss as possible. The performance of smooth hexahedral meshes was compared to that of a voxel-based mesh and a tetrahedral mesh, showing differences in computed force and stiffness. The stiffness and force decrease of the hexahedral meshes can be explained by the proportional volume loss. The tetrahedral mesh shows higher apparent stiffness and force that cannot be explained by the small volume increase. The difference is most probably due to the numerical stiffness of the linear tetrahedron. Indeed, the linear tetrahedron (1 integration point) is known to be stiff while the hexahedra are here integrated on 8 nodes with selective reduced integration to reduce shear-locking effects. The increased apparent stiffness behavior of the tetrahedral model is therefore most likely to be a numerical artifact. This numerical stiffness is less present in second (or higher) order tetrahedra. However, due to their simplicity and robustness, elements with linear shape functions are often preferred for non-linear problems, particularly when these involve large strains, frictional contact or material nonlinearities.

15.2.3 Sensitivity to Boundary Conditions

This section discusses the sensitivity of micro-FE models to the representation of boundary conditions. Boundary conditions represent experimental loading and support conditions of the modeled specimens. When qualitatively comparing the performance of several models, the applied boundary conditions might not be of importance to extract differences or similitudes between different groups as long as they are applied in the same way for each group. When quantitatively comparing models and experimental data, the accuracy of the boundary condition representation can be of great importance.

Most experimental tests of trabecular structure are compression tests of cylindrical samples. Representing the experimental setup in details can be considered. The interaction between the bone sample and the experimental setup (the sample extremities can be embedded into end-caps) is however often unknown. The setup is thus often simplified into a fixed end, on top of which lays the sample, and a moving one, compressing it. The applied boundary conditions can allow either for the material to move in the plane perpendicular to the loading direction (free boundary condition)

or not (constrained boundary conditions). Several studies have shown the influence of model choice on the results. The models are highly sensitive to the representation of the gripped end by constrained or free boundary conditions. Apparent parameters such as the apparent stiffness can show variations up to 40% [31], with a constrained numerical setup stiffer than a free one. Note however that a compression experiment with or without end-caps can show comparable deviations [28]. Using micro-FE models coupled to an optimization method to compute the hard tissue modulus gives however results which are better correlated [26] when the samples are tested with end-caps and modeled with constrained BC's, than tested without end-caps and modeled with frictionless contact conditions or free moving boundaries. The difference may be due to the fact that modeling a free moving boundary with frictionless behavior is not fully representative of the experimental conditions as purely frictionless behavior is not obtained. Approximating the boundary conditions usually leads to a higher tissue modulus than representing the actual experimental setup [4]. The method used to represent constrained or free boundary conditions can still be achieved in different ways using for instance either nodal constraints or contact conditions.

To analyze the effect of boundary representations, the bone antler sample introduced previously was compressed using eight different boundary setups (Fig. 15.6), all representing a fixed bottom surface and a moving upper one. Four of those setups had constrained boundary conditions at each end (Fig. 15.6a–d) and four had free boundary conditions at each end (Fig. 15.6e–h). The constrained condition on the fixed surface was represented either as a node constraint (cases a and b), for which the surface nodes were pinned in 3D, or as a friction contact with a rigid plane (cases c and d) using a friction coefficient of $\mu = 0.8$. The constrained condition on the




Fig. 15.6 Different types of constrained (*top row*) and free (*bottom row*) boundary conditions for a vertical displacement. Constraints in *grey* are applied to contact planes; constraints in *black* are applied to surface nodes


Table 15.3 Boundary conditions type dependent parameters (values in parentheses are the standard deviations)

Relative difference in stiffness (in % of the “no contact” stiffness)	Contact on moving side	Contact on pinned side	Contact on both sides
Constrained type BC’s	51.3 (±19.4)	76.9 (±23.5)	86.8 (±15.4)
Free type BC’s	26.8 (±3.5)	34.5 (±10.5)	51.3 (±9.4)

moving surface was similarly represented either as a node constraint (cases a and c) or as a contact condition (cases b and d). Similarly the free condition on the fixed surface was represented either as a node constraint (cases e and f), for which the nodes were pinned in the directions perpendicular to the compression, or as a frictionless contact condition (cases g and h). The free condition on the moving surface was either a node constraint (cases e and g) or a frictionless contact condition (cases f and h). For the case where two frictionless contact conditions were used (case h) the central node of the bottom surface was pinned to avoid rigid body motion.

As previously a compression of 5% was applied in a large strains framework and the performance of each model was assessed comparing the computed apparent stiffness and the maximal force reached. Exactly as constrained BC’s are stiffer than free ones, contact BC’s are stiffer than node constraints, whether constrained or free (Table 15.3). The deviation in the force between the different free representations stay proportional with the level of compression while the constrained representation shows increasing deviation with compression (Fig. 15.7).

Fig. 15.7 Example of force versus engineering strain function of the BC type (plain lines free BC’s, dashed lines constrained BC’s)



The ability to move laterally while applying a compression allows the virtual setup to be more compliant. The computed apparent stiffness and force decreases with the increasing flexibility for the sample to move or expand laterally.

15.3 Continuum Level Models of Trabecular Bone


Finite element models of trabecular bone at the continuum level have existed for many decades and have been used to investigate a range of clinical situations. Whilst micro-FE models have become more prevalent for the simulation of small regions of trabecular bone, the computational cost can be prohibitive, and whole bones and joints are still routinely simulated at the continuum level. Advances in imaging technologies such as CT and micro-CT have enabled more information to be derived for the generation of such models, including both the geometry and the spatially varying material properties.

15.3.1 *Sensitivity to Imaging and Material Property Assignment*

Continuum-level finite element models of bones with inhomogeneous material properties based on the underlying bone density have become widely adopted, and have been shown to provide better agreement with experimental data than those using uniform properties [44]. The elastic modulus of elements representing the trabecular bone regions within these models are often assigned on an element-by-element basis. Two approaches are commonly used to derive the elastic modulus values, as illustrated in Fig. 15.8. In the first (the ‘greyscale’ approach), an average greyscale is calculated from the voxels within the element volume, and the elastic modulus is calculated as a function of this value.

In the second (the ‘segmentation’ or ‘BV/TV’ approach), the underlying image is first segmented in order to calculate the bone volume fraction (BV/TV) and the modulus is then calculated as a function of the BV/TV. The segmentation approach allows extracting further information on the microstructure. In particular, information about the anisotropic organization of trabeculae can be computed from the segmented images. In that case, fabric tensor based orthotropic material properties can be derived. Accounting for anisotropy showed it can improve the correlation of bone morphology to bone strength for several anatomical sites [32, 37].

Fig. 15.8 Example of two different approaches for the derivation of mechanical properties from the micro-CT image data. For each element within the model, the elastic modulus is calculated based on either the mean greyscale (greyscale approach) or the BV/TV (segmentation approach) of that region



15.3.1.1 Sensitivity to the Methodology Used

In order to examine the sensitivity to the approach used, models from two previous specimen-specific studies [45, 54] were examined using both methods. In total, ten vertebral bodies (four human and six porcine) were imaged using micro-CT at a cubic voxel size of 0.074 mm, and tested under axial compression. From the micro-CT scans, two models were built of each specimen with the same element size of 1 mm. In one model, the greyscale approach was used, with the modulus linearly related to the mean greyscale of the voxels within the element. In the other, the segmentation approach was used, by first segmenting the images with a species-specific threshold (i.e. different for the human and porcine specimens). The elastic modulus was then linearly related to BV/TV. In both cases, the relationship between the image data and elastic modulus was optimized until the average error between the predicted stiffness of the models and the experimental values was minimized. The resulting predictions were compared to the experimental values and are shown in Fig. 15.9. The absolute average errors in stiffness for the two sets of models compared to the experimental results were then calculated. It was found that these errors were very similar for the two methods (6.5% for BV/TV method and 8.3% for greyscale method). Since the experimental error is likely to be of a similar order of magnitude to these errors, the results suggest that there is no advantage in using one method over the other, providing that the parameters used have been optimized.

15.3.1.2 Sensitivity to the Threshold for the Segmentation Method

As discussed in Sect. 15.2.1.1, the threshold selected to segment an image into trabecular bone and trabecular space will affect the thickness of the trabeculae and it has been shown that the choice of threshold can have a considerable effect on the

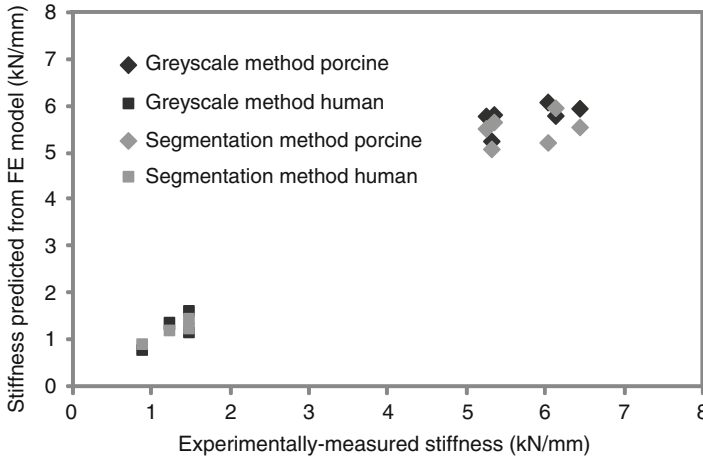


Fig. 15.9 Agreement between FE model predictions and experimental stiffness values for ten porcine and human vertebrae modeled using element-specific material properties based on BV/TV and mean greyscale data derived from the micro-CT image data of the specimens




Fig. 15.10 Predicted stiffness of four FE models of human vertebra [54] generated using the segmentation method to assign element-specific elastic modulus values. For each specimen, the predicted stiffness when two different threshold values were used to calculate the BV/TV is shown

calculated BV/TV values [38]. This then has a knock-on effect on the resulting FE model stiffness, as is illustrated in Fig. 15.10 for the four models of cadaveric vertebrae described above [54]. Here, two threshold values were selected to represent extremes of the range likely to be picked ‘by eye’, and it was found that the predicted stiffness varied by a mean of over 30%.

15.3.1.3 Sensitivity to the Relationship Used


Whichever method is used to determine the elastic modulus or other material properties from the underlying image data, the relationship between the greyscale and the property will have an effect on the final model. A number of different relationships have been used for both the BV/TV and the greyscale methods. Many of these relationships originate from density-modulus equations derived from previous experimental tests on trabecular bone specimens. The number of available equations in the literature is large, and several studies have investigated the sensitivity of the model outputs to the equation used, in many cases making direct comparisons with experimental tests in order to determine the most suitable relationship (e.g. [2, 11, 17]). These studies have generally found that the FE model predictions of apparent stiffness and local strain are highly sensitive to the equation adopted, which is unsurprising when the range of different equations in the literature is considered. However, there is little consensus across the studies on a single ‘optimum’ equation and, even within studies, different equations appear to fit different individual specimens better (e.g. [2, 17]). Many of the equations used are based on a power-law relationship between a measure of bone density (ρ) (ash density, apparent density, BV/TV etc.) and elastic modulus (E):

$$E = a\rho^b$$

The relationship derived by Morgan et al. [34] with a relatively low power ($b = 1.49$) appears to commonly be amongst the closest when the resulting FE models are compared to experimental data [11, 17]. In the study by Cong et al. [11], the authors determined optimum values of a and b to best fit the stiffness predictions of the FE models of femora to results obtained experimentally. They found an even lower power ($b = 1.16$) obtained the best results from a power-law equation. In another study [54], it was found that the value of the power had little effect on the performance of an individual model of a spinal vertebra, providing the associated constant, a , was optimized, as can be seen in Fig. 15.11.

The effect of the power will depend on the spread of the greyscale values within the underlying images. Extremely bright and dark regions in the image, which may

Fig. 15.11 Load–displacement curves for a vertebral model generated using different relationships. Adapted from [54]



be caused by artifacts, will dominate the model behavior with higher power terms by causing regions of overly large or small modulus values. The low sensitivity seen here makes a linear relationship a reasonable choice for this study. However the sensitivity may be affected by bone density range within the specimen set.

15.3.1.4 Discussion

A number of different methods have been used to both extract information from three dimensional image data sets and to derive finite element mechanical properties from this information. Two common methods for extraction, described here as the ‘greyscale’ method and the ‘segmentation’ method appear to yield models with relatively similar levels of accuracy. Both have their advantages and disadvantages. The segmentation method is sensitive to threshold, which in turn is sensitive to the scanner and settings used. Use of phantoms and/or automatic segmentation methods could allow scans from different imaging systems to be used, but as yet, little work has been undertaken to develop a robust framework for this process. The greyscale method is dependent on both the scanner and its settings, as well as the material within the marrow space. The use of phantoms is common place to calibrate scanners and relate the greyscale to the bone density. However, the greyscale of the trabecular space will be very different for dry bone specimens compared to those where the marrow is intact. So, whilst phantoms may get around some of the problems of using different scanners, it is more difficult to take account of different materials within the trabecular space.

The model predictions have been shown to be sensitive to the relationships used to assign the properties from this image information. The literature is awash with different equations and there is no clear consensus on an optimum choice. For applications with a limited density range, it may be possible to use a simple linear relationship, provided that the terms in the equation can be tuned for the specific species and type of bone that is used.

15.3.2 Sensitivity to the Finite Element Mesh

15.3.2.1 Sensitivity to the Finite Element Mesh Size

Many factors affect the convergence behavior with respect to the elements size of a finite element model, and the inhomogeneous properties, often coupled with a complex geometry, add particular challenges to the analysis of bone. To isolate the effects of material properties from those of the geometry, a study was undertaken where six rectangular cores of bone were extracted from continuum-level models of whole vertebrae and examined in isolation [27]. As the mesh density was altered (Fig. 15.12), models where the bone properties were inhomogeneous, and based on the underlying image greyscale, displayed less consistent convergence than those




Fig. 15.12 A typical cuboid ($12 \times 12 \times 32 \text{ mm}^3$) of trabecular bone meshed with three different element sizes (Model 4 = 4 mm^3 , Model 2 = 2 mm^3 , Model 1 = 1 mm^3). Adapted from [27]

Table 15.4 Difference in predicted stiffness between homogenous and greyscale-based FE models of trabecular bone cores generated with different element sizes (Model 4 = 4 mm, Model 2 = 2 mm, Model 1 = 1 mm)

	Homogenous properties	Greyscale-based properties
(Model 4-Model 2)/Model 2 (%)	6.3	5.5 (± 4.3)
(Model 2-Model 1)/Model 1 (%)	0.030	3.8 (± 2.5)

From [27]

where the properties were homogenous (Table 15.4). This is not surprising, since within the inhomogeneous greyscale-based models element size will alter not only the number of degrees of freedom, but also the distribution of material properties.

Zhao [56] decoupled these two effects by generating models of synthetic trabecular bone specimens based on micro-CT images in two different ways. In the first method, the images were down-sampled to different resolutions using a method incorporating partial volume effects. Then finite element models were generated with elements of the same size as the down-sampled images resolutions. In the second method, the images were down-sampled to the coarsest level, and then finite element models created from the images with varying mesh sizes, such that there were different numbers of elements but the materials properties of each were based on the same underlying greyscale grid (Fig. 15.13). Where this latter method was used, there was a rapid convergence that remained as the element size decreased. However with the first method, there was some evidence of convergence at larger element sizes (where the element was much larger than the trabecular bone structure), but as the image

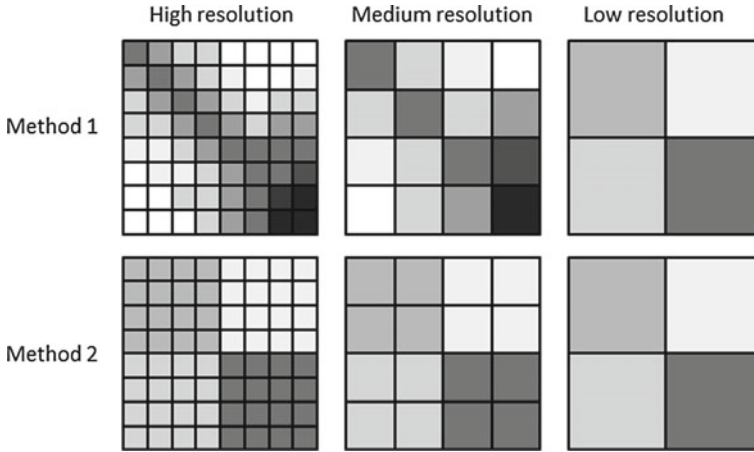


Fig. 15.13 Example in 2D of two methods of generating FE models from images at different element sizes. In Method 1, the underlying image is down-sampled, and the elements are assigned properties based on the image greyscale at that down-sampled level. In Method 2, the images are first down-sampled to the lowest resolution, then meshes of different sizes are used, such that the properties are always based on the same down-sampled image. Adapted from [56]

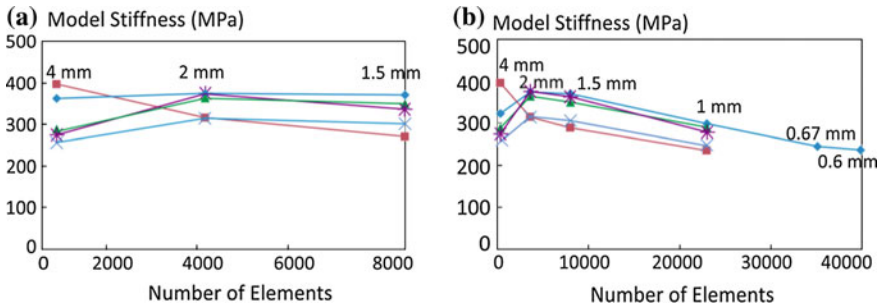


Fig. 15.14 Results of mesh convergence tests using **a** Method 1 and **b** Method 2 for models generated from images of a synthetic trabecular-like structure. Each line depicts the results of a different image data set. In Method 2, as the element size is decreased, the elements become a similar size to the trabeculae (~1 mm in this case) and eventually a second convergence occurs once the model becomes a micro-FE model. From [56]

resolution neared that of the trabecular structure itself, there was instability, since the elements were beginning to represent either trabecular space, or trabecular bone, rather than an average of the two (Fig. 15.14).

In a study including both isotropic and orthotropic models of cylindrical samples of deer antlers and other cellular materials, Mengoni et al. [33] decoupled the mesh size effects from the region size from which the material properties are extracted. The samples were meshed at a given fine resolution while the image-based material properties were extracted on regions of different size, assigning the same set of

properties to the underlying cluster of mesh elements. This method showed evidence of convergence for decreasing size of element clusters, before reaching the element-by-element level. For a given mesh size, while assigning material properties on an element-by-element basis may be necessary for isotropic models where only the modulus is evaluated from the image, for orthotropic models where fabric is also evaluated, it is less necessary and a cluster approach may be appropriate here because more information is extracted from the images.

15.3.2.2 Sensitivity to Type of Finite Element

There are relatively few studies that have specifically analyzed the element type in relation to continuum bone models. Ramos and Simoes [39] compared element type and order for simplified and realistic geometry femur models with homogeneous material properties. For the realistic model, the results for models with hexahedral elements converged with a lower number of elements than for the tetrahedral models, but there were some (<10%) differences between the first and second order converged results. For the tetrahedral elements, there was little difference in the results between the first and second order element types for the same number of elements. Their general conclusion was that the type of element “did not evidence significant differences”, however this will be very dependent on the nature of the problem and level of accuracy required. In the case of bone models with inhomogeneous, image-based material properties a change in element size will affect multiple aspects of the model (Sect. 15.3.2.1) and convergence testing is not straight forward.

15.3.3 Sensitivity to Boundary Conditions

The sensitivity of ovine trabecular bone models to the representation of boundary conditions replicating an experimental test were examined by Sikora [43]. Experimental tests were undertaken in which trabecular cores (approximately 10 mm in diameter and 20 mm in length) were extracted from ovine vertebrae, and set in delrin endcaps using a small quantity of polymethylmethacrylate (PMMA) cement, as shown in Fig. 15.15a. The specimens were imaged using an HR-pQCT (XtremeCT, Scanco Medical, Switzerland) with a voxel size of 0.041 mm and converted to finite element models with a 1 mm mesh size using proprietary software (ScanIP v4.2, Simpleware Ltd., UK). Four different methods of representing the boundary conditions on the bone were investigated. For each case, four specimens were investigated and the mean difference in predicted stiffness between the case representing the full experimental set up (Fig. 15.15a) and the different simplifications (Fig. 15.15b–d) were calculated.

The results are presented in Fig. 15.16. It can be seen that large relative differences on the stiffness occur where the boundary conditions do not provide the lateral




Fig. 15.15 Schematics of FE models of a trabecular bone compression test representing **a** the experimental set-up including the endcaps and PMMA cement, **b** a simplified version without the PMMA cement, **c** only the bone and **d** the bone with additional boundary conditions preventing lateral displacement (shown as *grey lines*). In all cases, the models were tied to rigid plates on the *top* and *bottom* surfaces. Adapted from [43]

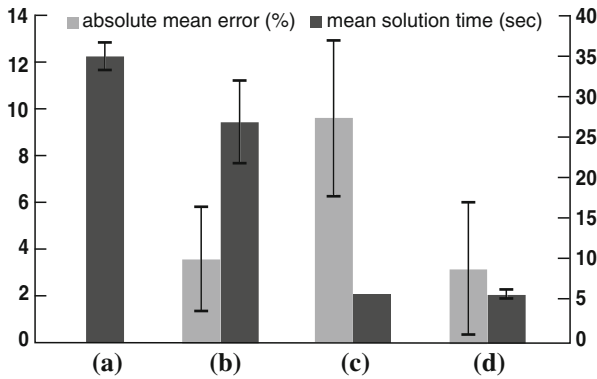


Fig. 15.16 Results for a boundary condition sensitivity study [43] undertaken on four models under the four different boundary conditions shown in Fig. 15.15. The absolute mean stiffness differences (*error bars* show standard deviations) are shown for cases (b)–(d) compared to the reference case (a). Mean solution times are also shown

constraint at the ends (case c). Similar results were also found by Zhao [56] using synthetic bone specimens. Here, an investigation was also undertaken to examine

the effects of different boundary conditions on a model similar to that shown in Fig. 15.15c.

The representation of whole bone specimens is just as sensitive to loading conditions, as was illustrated by Jones and Wilcox [27]. Here, specimen-specific models of spinal vertebrae were constructed to replicate an experimental compression test where the upper endplate was loaded via a steel ball, allowing it to tilt relative to the lower endplate. The predictions of stiffness from the FE models was found to be highly sensitive to the position at which the load was applied, with even deviations of 2 mm in location causing nearly 20 % change in the predicted stiffness. Such results illustrate not only the importance of boundary conditions in an FE model, but that this is also an issue with an experimental test, although it is often less apparent in the laboratory because of the variation between specimens.

From these results, it is clear that boundary conditions are important and changes in their implementation can cause substantial differences in model predictions. Even known conditions in a laboratory test can be represented in different ways, and simplifications can lead to substantial errors. The endcapped trabecular bone specimen should be very ‘easy’ to model and the situation becomes worse if the ends of the specimen are free to move laterally, as Zhao [56] showed, since the coefficient of friction between the bone and the loading platen is usually unknown. Where there are more unconstrained degrees of freedom, such as in the whole vertebra model, then the location of the constraints can play a major role and the replication of an experimental test becomes increasingly more difficult.

Boundary conditions are often applied without much justification, and these examples demonstrate that there is a need for thorough sensitivity tests since subtle changes in their application can lead to quite different results.

15.4 Discussion

This chapter provided evidence on the sensitivity of trabecular bone model outputs to the assumptions made during their construction, for micro-FE and continuum-FE models.

The level of sensitivity to several aspects, such as image segmentation and meshing, depended on the density of the specimen considered and the amount of deformation modeled. Model results were less sensitive to the assumptions made with the presence of higher density bone within the source specimen or with a low strain assumption during virtual testing. This is an important consideration when applying established methods to a new site in the body or new disease state.

15.4.1 Segmentation

Micro-FE predictions are sensitive to the threshold value used to segment the bone from the source image as this process affects both the total bone volume and can affect the connectivity within. Therefore changes to the threshold value will affect both macro behavior and that at an individual trabecular level. This sensitivity is more pronounced for low density bone specimens. The segmentation process also affects outcomes in the case of continuum-FE models which are based on BV/TV values. It is unclear whether one modeling approach provides a lower sensitivity to the image segmentation process as the available tests are not easily comparable (Sects. 15.2.1.1 and 15.3.2.1). In the absence of an ideal threshold value, the segmentation of the micro-structure should be as consistent as possible between specimens. The best chance of this consistency is through the use of phantoms, user training and automation where possible.

15.4.2 Meshing

The choice of optimum element type for micro-FE models is dependent on the intended application. For example, a direct conversion from image voxels to linear hexahedral elements may be a good choice where a linear analysis is sufficient, computational cost must be controlled and only macro level behavior is of interest. The representation of bending behavior of individual trabecular struts remains a challenge with each element shape and type of integration delivering different benefits and drawbacks (Sect. 15.2.2.3). Ultimately the choice for a particular project will likely be driven by the level of accuracy required and the computational resources available.

Regardless of what other sensitivity analyses are undertaken, a mesh convergence test is recommended for any new model or different application of an existing model [23]. However, the choice of element size for continuum-FE models is somewhat arbitrary. This chapter has discussed how the underlying structure, represented by the element-specific material properties, is captured at different resolutions depending on the element size. The image to material property conversion formula is likely dependent on the element size and therefore can be calibrated for a particular choice. The effect of element size is therefore corrected for during that calibration process. For lower resolution source images (such as traditional hospital grade QCT) the image resolution may provide a logical cap for the mesh resolution.

The choice of element integration for the representation of macro behavior using continuum-FE models requires the consideration of similar factors to the representation of micro behavior using micro-FE models. For example, care should be taken when using linear elements with standard integration if there is large deformation or bending, and a sensitivity test is always a useful check.

15.4.3 Image-to-Material Properties Relationship

Applying suitable material properties to finite element models of bone is challenging and there is currently no consensus on the optimal method. Although the use of source images provides a method of applying a realistic distribution of those properties, that distribution is reliant on the accuracy of the imaging modality and the choice of conversion formula. Relationships between image greyscale, density and material constants are sometimes derived from in vitro material tests, sometimes from theory and sometimes calibrated in silico against experimental data. Regardless the resulting material distribution is dependent on specimen preparation and scanner settings, underlining the importance of calibration phantoms during the image capture. In addition it is possible for any image to contain artifacts which will skew the greyscale distribution. The higher the power law in either micro- or continuum-FE methods, the greater the effect of those artifacts.

As well as allowing the analysis of local trabecular behavior, the micro-FE method can produce useful results with average homogeneous material properties. In contrast the continuum-FE method requires the calibration of a conversion formula in order to create the material property map. This calibration requires the use of image phantoms and knowledge of the bone type, density, and specimen-preparation. The use of segmented BV/TV values rather than average greyscale is one way to eliminate some of the variability between cases.

15.4.4 Boundary Conditions

Sensitivity to boundary conditions is an issue for both micro- and continuum-FE models of bone. The choice of boundary conditions has been shown to have one the most significant effects on apparent stiffness values of all of the aspects studied in this chapter. It is easier to match model boundary conditions to experimental boundary conditions where the latter are more constrained. Matching frictional properties at free boundaries is challenging as these are often unknown. It is necessary to take care with the choice of boundary constraints: contact conditions, even when frictionless, can constrain movement more than simple nodal constraints.

15.4.5 Looking Forward

Both micro- and continuum-FE modeling techniques have advantages which should secure them a place in the virtual representation of trabecular bone for the foreseeable future. Continuum-FE can capture the inhomogeneity of the micro-structure sufficiently well to generate stiffness and strength predictions at a macro level, while keeping computational cost low enough to make whole bone models possible. In con-

trast the computational cost of micro-FE is high making it currently more suitable for the analysis of local behavior in small samples.

Information from sensitivity testing is a crucial aspect to be considered alongside experimental data on tissue properties and comparative validation studies, which together provide the necessary confidence in model predictions. The majority of the sensitivity tests reported for trabecular bone models consider overall stiffness as the output of interest. This trend is reflected in the studies detailed in this chapter. Establishing the accuracy of the overall stiffness prediction is a natural starting point for the development of these mechanical models of bone and is therefore the most well documented. However, advances in model sophistication are allowing the prediction of local behavior and the simulation of bone failure. These developments currently outstrip the availability of relevant sensitivity information. Sensitivity data, which quantifies the effect of key parameters on these alternative modeling outputs, will be an important part of next stage of evolution in this research area.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

1. Anderson, A.E., Ellis, B.J., Weiss, J.A.: Verification, validation and sensitivity studies in computational biomechanics. *Comput. Methods Biomech. Biomed. Eng.* **10**(3), 171–184 (2007)
2. Austman, R.L., Milner, J.S., Holdsworth, D.W., Dunning, C.E.: The effect of the density-modulus relationship selected to apply material properties in a finite element model of long bone. *J. Biomech.* **41**, 3171–3176 (2008)
3. Bessho, M., Ohnishi, I., Matsuyama, J., Matsumoto, T., Imai, K., Nakamura, K.: Prediction of strength and strain of the proximal femur by a CT-based finite element method. *J. Biomech.* **40**, 1745–1753 (2007)
4. Bevill, G., Eswaran, S.K., Farahmand, F., Keaveny, T.M.: The influence of boundary conditions and loading mode on high-resolution finite element-computed trabecular tissue properties. *Bone* **44**, 573–578 (2009)
5. Bourne, B.C., van der Meulen, M.C.: Finite element models predict cancellous apparent modulus when tissue modulus is scaled from specimen CT-attenuation. *J. Biomech.* **37**, 613–621 (2004)
6. Boutroy, S., Van Rietbergen, B., Sornay-Rendu, E., Munoz, F., Bouxsein, M.L., Delmas, P.D.: Finite element analysis based on in vivo HR-pQCT images of the distal radius is associated with wrist fracture in postmenopausal women. *J. Bone Miner. Res.* **23**(3), 392–399 (2008)
7. Buckley, J.M., Loo, K., Motherway, J.: Comparison of quantitative computed tomography-based measures in predicting vertebral compressive strength. *Bone* **40**, 767–774 (2007)
8. Carlier, A., Van Oosterwyck, H., Geris, L.: *In Silico Biology of Bone Regeneration Inside Calcium Phosphate Scaffolds Tissue Engineering*, pp. 31–48. Springer, The Netherlands (2014)
9. Checa, S., Prendergast, P.J., Duda, G.N.: Inter-species investigation of the mechano-regulation of bone healing: comparison of secondary bone healing in sheep and rat. *J. Biomech.* **44**, 1237–1245 (2011)

10. Cifuentes, A., Kalbag, A.: A performance study of tetrahedral and hexahedral elements in 3-d finite element structural analysis. *Finite Elem. Anal. Des.* **12**, 313–318 (1992)
11. Cong, A., Buijs, J.O.D., Dragomir-Daescu, D.: In situ parameter identification of optimal density-elastic modulus relationships in subject-specific finite element models of the proximal femur. *Med. Eng. Phys.* **33**, 164–173 (2011)
12. d’Otreppe, V.: From medical imaging to finite element simulations: a contribution to mesh generation and locking-free formulations for tetrahedra. Ph.D. Thesis, University of Liege (2012)
13. d’Otreppe, V., Boman, R., Ponthot, J.-P.: Generating smooth surface meshes from multi-region medical images. *Int. J. Numer. Methods Biomed. Eng.* **28**, 642–660 (2012)
14. de Bien, C., Mengoni, M., d’Otreppe, V., Freichels, H., Jérôme, C., Ponthot, J.-P., Léonard, A., Toye, D.: Development of a biomechanical model of deer antler cancellous bone based on x-ray microtomographic images. In: *Proceedings of Micro-CT User Meeting*, pp. 137–145, 2012
15. Depalle, B., Chapurlat, R., Walter-Le-Berre, H., Bou-Saïd, B., Follet, H.: Finite element dependence of stress evaluation for human trabecular bone. *J. Mech. Behav. Biomed. Mater.* **18**, 200–212 (2013)
16. Ding, M., Odgaard, A., Hvid, I.: Accuracy of cancellous bone volume fraction measured by micro-CT scanning. *J. Biomech.* **32**, 323–326 (1999)
17. Eberle, S., Göttlinger, M., Augat, P.: An investigation to determine if a single validated density-elasticity relationship can be used for subject specific finite element analyses of human long bones. *Med. Eng. Phys.* **35**, 875–883 (2013)
18. Geris, L., Gomez-Cabrero, D.: An introduction to uncertainty in the development of computational models of biological processes. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
19. Guldberg, R.E., Hollister, S.J., Charras, G.T.: The accuracy of digital image-based finite element models. *Trans. ASME J. Biomech. Eng.* **120**, 289–295 (1998)
20. Hara, T., Tanck, E., Homminga, J., Huiskes, R.: The influence of microcomputed tomography threshold variations on the assessment of structural and mechanical trabecular bone properties. *Bone* **31**, 107–109 (2002)
21. Harrison, N.M., McDonnell, P.F., O’Mahoney, D.C., Kennedy, O.D., O’Brien, F.J., McHugh, P.E.: Heterogeneous linear elastic trabecular bone modelling using micro-CT attenuation data and experimentally measured heterogeneous tissue properties. *J. Biomech.* **41**, 2589–2596 (2008)
22. Helgason, B., Perilli, E., Schileo, E., Taddei, F., Brynjólfsson, S., Viceconti, M.: Mathematical relationships between bone density and mechanical properties: a literature review. *Clin. Biomech.* **23**, 135–146 (2008)
23. Henninger, H.B., Reese, S.P., Anderson, A.E., Weiss, J.A.: Validation of computational models in biomechanics. *Proc. Inst. Mech. Eng. Part H J. Eng. Med.* **224**, 801–812 (2010)
24. Hojjat, S.-P., Beek, M., Akens, M.K., Whyne, C.M.: Can micro-imaging based analysis methods quantify structural integrity of rat vertebrae with and without metastatic involvement? *J. Biomech.* **45**, 2342–2348 (2012)
25. Homminga, J., Huiskes, R., Van Rietbergen, B., Rügsegger, P., Weinans, H.: Introduction and evaluation of a gray-value voxel conversion technique. *J. Biomech.* **34**, 513–517 (2001)
26. Jacobs, C.R., Davis, B.R., Rieger, C.J., Francis, J.J., Saad, M., Fyhrie, D.P.: The impact of boundary conditions and mesh size on the accuracy of cancellous bone tissue modulus determination using large-scale finite-element modeling. *J. Biomech.* **32**, 1159–1164 (1999)
27. Jones, A.C., Wilcox, R.K.: Assessment of factors influencing finite element vertebral model predictions. *Trans. ASME J. Biomech. Eng.* **129**, 898–903 (2007)
28. Keaveny, T.M., Pinilla, T.P., Crawford, R.P., Kopperdahl, D.L., Lou, A.: Systematic and random errors in compression testing of trabecular bone. *J. Orthopaed. Res.* **15**, 101–110 (1997)
29. Keaveny, T.M., Morgan, E.F., Niebur, G.L., Yeh, O.C.: Biomechanics of trabecular bone. *Annu. Rev. Biomed. Eng.* **3**, 307–333 (2001)

30. Kosmopoulos, V., Keller, T.S.: Predicting trabecular bone microdamage initiation and accumulation using a non-linear perfect damage model. *Med. Eng. Phys.* **30**, 725–732 (2008)
31. Ladd, A.J.C., Kinney, J.H.: Numerical errors and uncertainties in finite-element modeling of trabecular bone. *J. Biomech.* **31**, 941–945 (1998)
32. Matsuura, M., Eckstein, F., Lochmueller, E.-M., Zysset, P.K.: The role of fabric in the quasi-static compressive mechanical properties of human trabecular bone from various anatomical locations. *Biomech. Model. Mechanobiol.* **7**(1), 27–42 (2008)
33. Mengoni, M., Voide, R., de Bien, C., Freichels, H., Jérôme, C., Léonard, A., Toye, D., van Lenthe, G.H., Müller, R., Ponthot, J.-P.: A non-linear homogeneous model for bone-like materials under compressive load. *Int. J. Numer. Methods Biomed. Eng.* **28**(2), 334–348 (2012)
34. Morgan, E.F., Bayraktar, H.H., Keaveny, T.M.: Trabecular bone modulus-density relationships depend on anatomic site. *J. Biomech.* **36**, 897–904 (2003)
35. Müller, R., Rügsegger, P.: Three-dimensional finite element modelling of non-invasively assessed trabecular bone structures. *Med. Eng. Phys.* **17**, 126–133 (1995)
36. Niebur, G.L., Feldstein, M.J., Yuen, J.C., Chen, T.J., Keaveny, T.M.: High-resolution finite element models with tissue strength asymmetry accurately predict failure of trabecular bone. *J. Biomech.* **33**, 1575–1583 (2000)
37. Pahr, D.H., Zysset, P.K.: A comparison of enhanced continuum FE with micro FE models of human vertebral bodies. *J. Biomech.* **42**, 455–462 (2009)
38. Parkinson, I.H., Badiei, A., Fazzalari, N.L.: Variation in segmentation of bone from micro-ct imaging: implications for quantitative morphometric analysis. *Australas. Phys. Eng. Sci. Med.* **31**, 160–164 (2008)
39. Ramos, A., Simoes, J.: Tetrahedral versus hexahedral finite elements in numerical modelling of the proximal femur. *Med. Eng. Phys.* **28**, 916–924 (2006)
40. Rügsegger, P., Koller, B., Müller, R.: A microtomographic system for the nondestructive evaluation of bone architecture. *Calcif. Tissue Int.* **58**, 24–29 (1996)
41. Schulte, F.A., Zwahlen, A., Lambers, F.M., Kuhn, G., Ruffoni, D., Betts, D., Webster, D.J., Müller, R.: Strain-adaptive in silico modeling of bone adaptation: a computer simulation validated by in vivo micro-computed tomography data. *Bone* **52**, 485–492 (2013)
42. Schliemann-Bullinger, M., Fey, D., Bastogne, T., Findeisen, R., Scheurich, P., Bullinger, E.: The experimental side of parameter estimation. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
43. Sikora, S.: *Experimental and Computational Study of the Behaviour of Trabecular Bone-Cement Interfaces*. PhD Thesis, University of Leeds, Leeds (2013)
44. Taddei, F., Cristofolini, L., Martelli, S., Gill, H.S., Viceconti, M.: Subject-specific finite element models of long bones: an in vitro evaluation of the overall accuracy. *J. Biomech.* **39**, 2457–2467 (2006)
45. Tarsuslugil, S.M., O’Hara, R.M., Dunne, N.J., Buchanan, F.J., Orr, J.F., Barton, D.C., Wilcox, R.K.: Development of calcium phosphate cement for the augmentation of traumatically fractured porcine specimens using vertebroplasty. *J. Biomech.* **46**, 711–715 (2013)
46. Taylor, R.L., Simo, J.C., Zienkiewicz, O.C., Chan, A.C.H.: The patch test: a condition for assessing FEM convergence. *Int. J. Numer. Methods Eng.* **22**, 39–62 (1986)
47. Ulrich, D., Van Rietbergen, B., Weinans, H., Rügsegger, P.: Finite element analysis of trabecular bone structure: a comparison of image-based meshing techniques. *J. Biomech.* **31**, 1187–1192 (1998)
48. Unnikrishnan, G.U., Morgan, E.F.: A new material mapping procedure for quantitative computed tomography-based continuum finite element analyses of the vertebra. *J. Biomech. Eng.* **133**(7), 071001 (2011)
49. Varga, P., Pahr, D.H., Baumbach, S., Zysset, P.K.: HR-pQCT based FE analysis of the most distal radius section provides an improved prediction of Colles’ fracture load in vitro. *Bone* **47**(5), 982–988 (2010)
50. Van Lenthe, G.H., Müller, R.: Prediction of failure load using micro-finite element analysis models: towards in vivo strength assessment. *Drug Discov. Today Technol.* **3**(2), 221–229 (2006)

51. Van Rietbergen, B., Weinans, H., Huiskes, R., Odgaard, A.: A new method to determine trabecular bone elastic properties and loading using micromechanical finite-element models. *J. Biomech.* **28**, 69–81 (1995)
52. Van Rietbergen, B., Weinans, H., Huiskes, R., Polman, B.J.W.: Computational strategies for iterative solutions of large fem applications employing voxel data. *Int. J. Numer. Methods Eng.* **39**, 2743–2767 (1996)
53. Wagner, D.W., Lindsey, D.P., Beaupre, G.S.: Deriving tissue density and elastic modulus from microCT bone scans. *Bone* **49**, 931–938 (2011)
54. Wijayathunga, V.N., Jones, A.C., Oakland, R.J., Furtado, N.R., Hall, R.M., Wilcox, R.K.: Development of specimen-specific finite element models of human vertebrae for the analysis of vertebroplasty. *Proc. Inst. Mech. Eng. Part H-J. Eng. Med.* **222**, 221–228 (2008)
55. Wolfram, U., Wilke, H.J., Zysset, P.K.: Valid μ finite element models of vertebral trabecular bone can be obtained using tissue properties measured with nanoindentation under wet conditions. *J. Biomech.* **43**, 1731–1737 (2010)
56. Zhao, Y.: Finite Element Modelling of Cement Augmentation and Fixation for Orthopaedic Applications. Ph.D. Thesis, University Of Leeds, Leeds (2010)
57. Zöllner, A.M., Tepole, A.B., Kuhl, E.: On the biomechanics and mechanobiology of growing skin. *J. Theor. Biol.* **297**, 166–175 (2012)
58. Zysset, P.K.: A review of morphology-elasticity relationships in human trabecular bone: theories and experiments. *J. Biomech.* **36**(10), 1469–1485 (2003)

Part V
Model Predictions Under Uncertainty

Chapter 16

Neuroswarm: A Methodology to Explore the Constraints that Function Imposes on Simulation Parameters in Large-Scale Networks of Biological Neurons

David Gomez-Cabrero, Salva Ardid, Maria Cano-Colino,
Jesper Tegnér and Albert Compte

Abstract Candidate mechanisms of brain function can potentially be identified using biologically detailed computational models. A critical question that arises from the construction and analysis of such models is whether a particular set of parameters is unique or whether multiple different solutions exist, each capable of reproducing some relevant phenomenology. Addressing this issue is difficult, and systematic procedures have been proposed only recently, targeting small systems such as single neurons or small neural circuits [16] (Marder and Taylor, *Nat Neurosci* 14:133–138, 2011), [1] (Achard and De Schutter, *PLoS Comput Biol* 2:e94, 2006). However, how to develop a methodology to address the problem of non-uniqueness of parameters in large-scale biological networks is yet to be developed. Here, we describe a computational strategy to explicitly approach this issue on large-scale neural network models, which has been successfully applied to computational models of working memory (WM) and selective attention [2] (Ardid, *J Neurosci Off J Soc Neurosci* 30:2856–2870, 2010), [3] (Cano-Colino et al., *Cereb Cortex* 24:2449–2463, 2014). To illustrate the approach, we show in this chapter how our strategy applies to the problem of identifying different mechanisms underlying visuospatial WM. We use a well-established biological neural circuit model in the literature [6] (Compte et al., *Cereb. Cortex* 10:910–923, 2000) as a reference point, which we then

D. Gomez-Cabrero (✉) · J. Tegnér
Unit of Computational Medicine, Center for Molecular Medicine, Department of Medicine,
Karolinska Institutet, Solna, Sweden
e-mail: david.gomezcabrero@ki.se

D. Gomez-Cabrero · S. Ardid · M. Cano-Colino · A. Compte
Institut d'Investigacions Biomediques August Pi i Sunyer (IDIBAPS),
Barcelona, Spain

S. Ardid
Center for Computational Neuroscience and Neural Technology (CompNet),
Department of Mathematics and Statistics, Boston University, Boston, USA

M. Cano-Colino
Champalimaud Neuroscience Programme, Champalimaud Centre for the Unknown,
Lisbon, Portugal

© Springer International Publishing Switzerland 2016
L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_16

perturb by using the Swarm Optimization Algorithm. This algorithm explores the space of biologically unconstrained parameters in the model under the constraint of preserving a solution defined here as a network in which the activity of model neurons mimics the properties of neurons in the dorsolateral prefrontal cortex (dlPFC) of monkeys performing a visuospatial WM task [7] (Funahashi et al., *J Neurophysiol* 61:331–349, 1989). The results are: (1) identification of a set of model solutions, composed of alternative and, in principle, feasible and sufficient mechanisms generating WM function in a cortical network. In particular, we found that the dynamics of interneurons play a main role in distinguishing among potential circuit candidates. Secondly we uncovered compensatory mechanisms in a subset of the parameters in the model. In essence, the compensatory mechanisms we observe in the different solutions are based on correlations between sets of parameters that shift the local Excitatory/Inhibitory balance in opposite directions. In summary, our approach is able to identify distinct mechanisms underlying a same function, as well as to propose a dynamic solution to the problem of fine-tuning. Our results from the proposed workflow would be strengthened by additional biological experiments aimed to refine the validity of the results.

Keywords Prefrontal cortex · Workflow · Ensemble analysis · Working memory model · Neuroscience · Computational biology

16.1 Introduction

A branch of Computational Biology makes use of mathematical modelling (such as differential ordinary equation systems) to understand better the mechanisms of the biological system of interest. In those cases, models are tools to test and generate hypotheses, to then validate experimentally. But the use of models is not trivial and requires robust methodologies of data analysis, model generation [14, 22], parameter estimation [4, 23], and experimental design [27]. A particularly crucial decision in this process is how complexity and uncertainty are being considered during the modelling [11, 13, 15]. However, the major challenge in using computational models under uncertainty is the generation of relevant hypotheses that are not exclusively dependent on choices during modelling (such as parameter selection). In this chapter, we provide a methodology for robust hypothesis generation in the context of Neuroscience and under parameter uncertainty; relevant work in the book addressing similar challenges are described in [5, 15].

Mechanistic aspects of brain function can be studied with the use of biologically detailed computational models. Those models detail relations and/or interactions between entities through mathematical formulations that depend on a set of parameter values. The first major success in modelling was the classical model of Hodgkin and Huxley ([10] HHM) of the action potential. HHM was developed to provide a mechanistically unified system description by mathematically organizing experimental observations. Interestingly, HHM not only described known facts but also

allowed to generate predictions that only many years later, when technical development allowed, were experimentally validated [9]. The success of the model started what is known nowadays as computational biology. Typically, a general approach for modelling-based studies has three phases: (1) the collection of relevant experimental facts and expert knowledge to be considered, (2) the mathematical description of the system, and (3) the fine-tuning of the parameters to reproduce the “expected” behaviours.

A critical question that arises from that approach is to what extent conclusions depend on particular simulation parameters (see [5, 15]). However, to demonstrate whether a model is unique in reproducing some relevant phenomenology can be a tall order. Interestingly, the questions when more than one solution (parameter set) reproduce expected experimental behaviours (which we will term “solution” parameter sets), uncover many other critical questions such as: (1) are those parameter sets providing different biological mechanisms or predictions? [5], (2) do parameter sets group into families of solutions or do they provide a continuum of solutions? Following these questions [16] proposed the need to investigate in populations of parameter sets in order to discover compensatory mechanisms in neurons or circuits.

In the context of single neurons and small neural networks this problem has been addressed by constructing and analysing databases of models compatible with biological function [1, 20, 21]. Small networks of three interconnected cells could reproduce the rhythmic patterns of activity in the crustacean stomatogastric ganglion for millions of different, disparate parameter combinations [21]. In the analysis of a cerebellar Purkinje cell models authors found 20 different *solution* models (i.e. parameter sets compatible with experimental data) [1]. Furthermore, by studying the parameter landscape created by the good models it was found that the *parameter space of good models* could be defined by a set of “*loosely connected hyperplanes*” [1].

In the present chapter we present a computational strategy to explicitly explore, group and characterize parameter-sets on large-scale neural network models. This strategy is similar to that described in [16], but adapted to deal with the complexity and computational cost of large-scale neural simulations. We applied this strategy to study a specific cognitive function, visuospatial working memory, which can be modelled with a biological neural network [6] that mimics the properties and dynamics of neurons in the dorsolateral prefrontal cortex (PFC) of monkeys engaged in oculomotor delayed response tasks [7]. In such tasks, the monkey is required to retain the location of a visual cue during a delay period between the cue stimulus and the memory-guided saccadic response, and PFC neurons reflect this memorized information through selective persistent activation in the delay period [7].

The typical experimental design is depicted in Fig. 16.1a. The trial starts with a blank screen containing just a central cross on which the monkey fixates its gaze to initiate the trial (“pre-stimulus”). While fixating, a stimulus cue appears briefly in one of eight possible locations equidistant from the fixation point (“stimulus”). After cue presentation a delay period of a few seconds follows during which the monkey needs to remember the location of the previously presented cue (“delay period”). At the end of the delay period, the fixation cross disappears and the monkey makes a

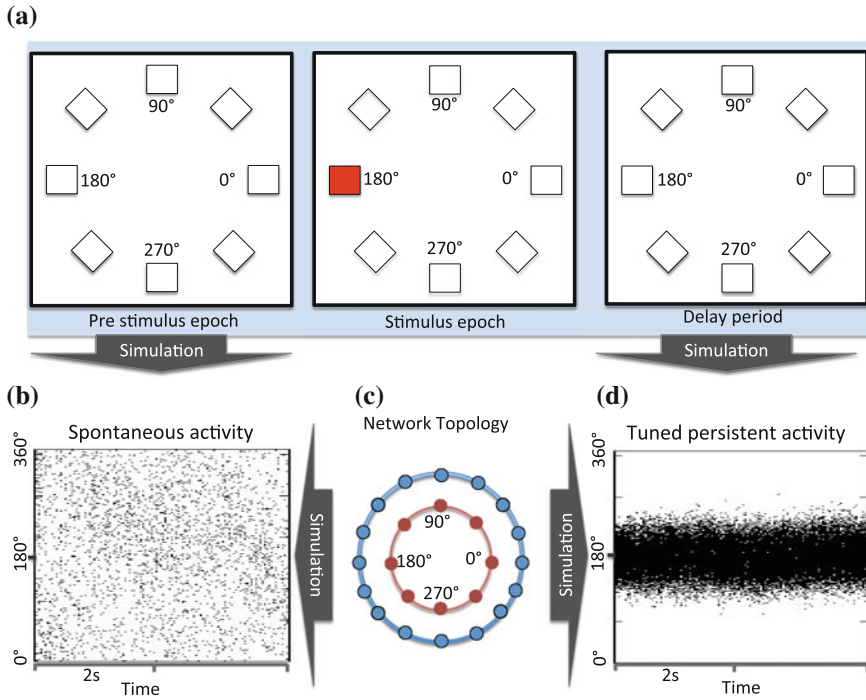


Fig. 16.1 The experiment, the model and the simulation results explained. **a** Three stages of the experiment. In a first stage there is no stimulus and a blank screen is observed. In a second Stage a visual stimulus is briefly flashed in one of 8 possible locations; finally in a third stage the stimulus disappears but it must be remembered during few seconds. The task consists in reporting the location of the briefly flashed stimulus after a delay of a few seconds and it thus require memorising this location. **c** Topology of the neuronal network used to model neural activity during this task. Neurons have a 1-D ring topology, and their position in the ring is associated to the possible location of the stimulus. There are two concentric rings, a ring of pyramidal (excitatory) neurons (in blue) and a ring of interneurons (inhibitory neurons, in red). The position of the two rings denotes the selectivity of excitatory and inhibitory neurons. Neurons are connected according to their relative angular distance on the ring. **b** Spontaneous firing pattern of the neurons over time. A point denotes a given neuron (y-axis) firing in a given time (x-axis). **d** Persistent firing pattern of the neurons associated with the location of the presented stimulus

saccadic eye movement to the location where he remembers the cue was presented. In experiments, PFC neurons show tuned persistent activity in the delay period of this behavioural protocol [7]. This experimental design can also be simulated using a network model of excitatory and inhibitory neurons [6]. By arranging neurons according to spatial selectivity (Fig. 16.1c), connectivity parameters can be tuned so that strong local excitation and strong global feedback inhibition combine to produce neuronal responses in line with experimental data. Figure 16.1b, d shows in rastergrams the activity of model pyramidal neurons in one tuned network (see next section) during the pre-stimulus and delay period epochs, respectively [6]. Notably, activity in the pre-stimulus epoch appears uniform and at low firing rates, and we

term this condition “spontaneous activity”, and network activity in the delay period is tuned and sustained, and we call this “tuned persistent activity”.

The methodology proposed here is similar to the one described in [8], which was used to analyse a computational model of atherosclerosis. We update the necessary steps to apply it in Neuroscience, in particular for the analysis of the PFC model (see also [2, 3]). Interestingly our approach is able to identify: (1) compensatory mechanisms, (2) characteristics of inhibitory neurons firing patterns that allow the grouping of solutions, and (3) significant opposite compensatory mechanisms in some of the groups identified. Therefore, our computational approach that explores the “solution space” identifies relevant mechanisms to be further tested by appropriate experimental designs.

Section 16.2 of the chapter presents a review of the PFC model. Section 16.3 details how feasible parameter sets are searched. The following section, Sect. 16.4, summarizes the integrative analysis of the parameter sets that are considered to be correct based on the experimental data. The final section, Sect. 16.5, presents the conclusions regarding the utility of the proposed approach and summarizes the biological results associated to PFC.

16.2 PFC-Working Memory Model

We used the PFC network model described in [6]. We refer the reader to this publication for a thorough account of the computational model and we provide here only a succinct description. The network contains 1,024 excitatory neurons and 256 inhibitory neurons modelled according to the leaky integrate-and-fire formalism [24]. Model neurons are arranged according to their preferred cue directions in a ring topology as shown in Fig. 16.1c and they are interconnected via conductance-based synapses with dynamics consistent with AMPA, NMDA and GABA_A receptor mediated synaptic transmission in the cortex. Specifically, AMPA and GABA_A synaptic conductances jump instantaneously when a pre-synaptic spike occurs and decay exponentially with time constant 2 ms for AMPA and 10 ms for GABA_A.

NMDA conductances are voltage dependent and their dynamics are defined by a rise time (set to 2 ms), a decay time (set to 100 ms) and a saturation term such that they become insensitive to high presynaptic firing rates.


$$\frac{ds}{dt} = -\frac{1}{\tau_s}s + \alpha_s x(1-s)$$

$$\frac{dx}{dt} = -\frac{1}{\tau_x}x + \sum_i \delta(t - t_i)$$

While the parameters defining these time dynamics are relatively well constrained by experimental data, the strength of the conductances are much more unconstrained

Notation	Description	Lower Bound	Upper Bound	Value of Reference
$G_{EE,AMPA}$ (nS)	Recurrent excitatory-to-excitatory synaptic conductance mediated by AMPAR channels	0	0.98	0
$G_{EE,NMDA}$ (nS)	Recurrent excitatory-to-excitatory synaptic conductance mediated by NMDAR channels	0	1.95	0.76
$G_{EI,AMPA}$ (nS)	Recurrent excitatory-to-inhibitory synaptic conductance mediated by AMPAR channels	0	0.98	0
$G_{EI,NMDA}$ (nS)	Recurrent excitatory-to-inhibitory synaptic conductance mediated by NMDAR channels	0	1.95	0.58
G_{II} (nS)	Recurrent inhibitory-to-excitatory synaptic conductance mediated by GABAAR channels	0	7.81	2.67
G_{II} (nS)	Recurrent inhibitory-to-inhibitory synaptic conductance mediated by GABAAR channels	0	7.81	2.05
σ_{EE} (°)	Excitatory-to-excitatory width of the connectivity footprint	7.2	360	14.4 **
σ_{EI} (°)	Excitatory-to-inhibitory width of the connectivity footprint	7.2	360	Unstructured *
σ_{IE} (°)	Inhibitory-to-excitatory width of the connectivity footprint	7.2	360	Unstructured *
σ_{II} (°)	Inhibitory-to-inhibitory width of the connectivity footprint	7.2	360	Unstructured *
J_{EE}^+	Excitatory-to-excitatory strength of the stronger isodirectional connections	0.8	3	1.62
J_{EI}^+	Excitatory-to-inhibitory strength of the stronger isodirectional connections	0	2	1*
J_{IE}^+	Inhibitory-to-excitatory strength of the stronger isodirectional connections	0	2	1*
J^{II}	Inhibitory-to-inhibitory strength of the stronger isodirectional connections	0	2	1*
$g_{ext,EE}$ (nS)	External excitatory (AMPA-mediated) conductance on excitatory cells	0	6	3.1
$g_{ext,II}$ (nS)	External excitatory (AMPA-mediated) conductance on inhibitory cells	0	6	2.38

(a) Parameter Description



(b) Structured connectivity of the model.

Fig. 16.2 Parameters selected to vary in the PFC-model. **a** Parameter names, description and lower and upper bounds are included. (*) cross- and isodirectional components of these connections were equally strong (i.e. $J^+ = 1$; in this case, the value of σ is irrelevant; see panel **b**). (**) corrected from [6]. **b** Structured connectivity of the model. The synaptic connection strength decreases with the difference in the preferred cues of two neurons, with strong interactions between neighboring neurons and weak interactions between more distant neurons

and need to be tuned to achieve the required function. We impose however a topographical constraint, so that the strength of synaptic conductances is a function of the difference in the preferred cues of the presynaptic and postsynaptic neurons. We specified this function to be a Gaussian, defined by three parameters (see Fig. 16.2b): the width σ , the tuning parameter J^+ , and the overall strength g . It has been shown that if excitatory connections among excitatory neurons are such that synaptic strength decreases with the difference in the preferred cues of two neurons, with strong interactions between neighbouring neurons and weak interactions between more distant neurons, the network has regimes of operation compatible with working memory physiology: tuned persistent delay period activity (Fig. 16.1d) bi-stable with a low-rate, unstructured spontaneous activity (Fig. 16.1b).

For a given set of parameters the model is simulated for 5 s in repeated trials. Some trials are run without any phasic external stimulation to test for the stability of the spontaneous activity (Fig. 16.1b). In other trial simulations (persistent activity trials), a stimulus is applied by transiently injecting current to a subset of neurons after the first second of simulation. Every time a model is run a numerical seed is

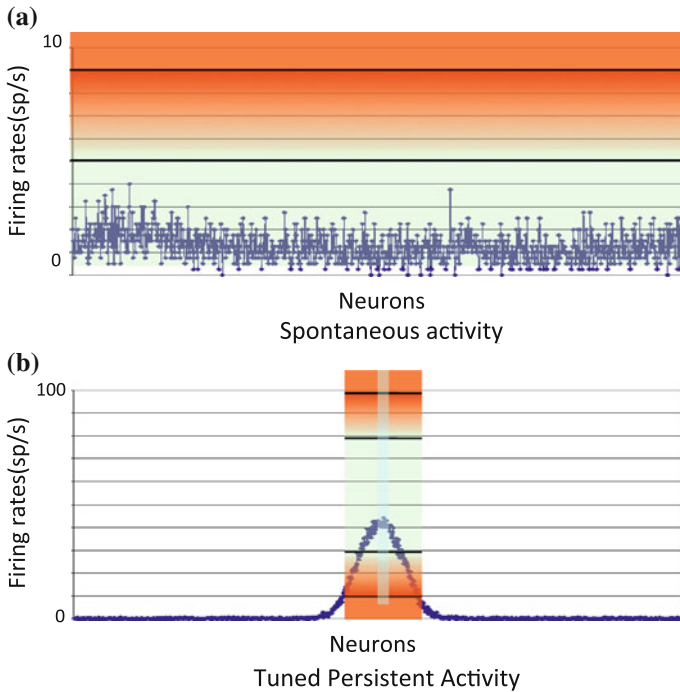


Fig. 16.3 Evaluation functions in plots. **a** Expected firing rate for spontaneous activity. The maximum firing rate is limited to 8 sp/s, but penalized if more than 5 sp/s. Also there is a penalty for the fraction of pyramidal neurons that are silent. **b** Expected firing rate pattern for persistent activity. The maximum firing rate is limited to 100 sp/s, and penalized if it is more than 80 sp/s. The minimum firing rate is 10 sp/s, and penalized if it is less than 30 sp/s

randomly selected that controls the timing of nonspecific external Poisson spiking activity that depolarizes the network neurons and set a general random background activity in the network. The output of the model for a given random seed is the timing of the spiking events for all the neurons included in the model (*rastergram*); examples of rastergrams with the desired patterns of activity are presented in Fig. 16.1 for persistent activity (Fig. 16.1d) and for spontaneous activity trials (Fig. 16.1b). The stationary patterns in these rastergrams can be summarized with *firing rates* (measures in *spikes per second*) computed in a final 4s window of the simulation. Firing rate plots are shown in Fig. 16.3. For the same parameter set the rastergram and firing rates may differ between different trial simulations depending on the initial numerical seeds selected.

16.3 Neuroswarm: A Tool to Explore the Parameter Space

While the network solution of Compte et al. [6] (Fig. 16.1b, d) had qualitative features consistent with experimental data, some aspects of model function did not match quantitatively: the model displayed a large gap between persistent activity to preferred and non-preferred stimuli (>20 sp/s) while experimentally this is a narrower gap (<10 sp/s); the model required a very precise symmetry in the translationally invariant connectivity; connectivity parameters required a significant degree of fine tuning; and neuronal activity in the persistent state was more stable in the model than observed experimentally. These discrepancies could mark fundamental flaws in the model or some of them could be specific of the dynamical regime ensuing from our particular choice of parameters and could be alleviated in a different parameter regime. Here, we will design a protocol to address this question by exploring network behaviour in very different parameter configurations. The model has over 100 parameters, and we selected 16 that we considered to be unconstrained by experimental data and relevant candidates to regulate the expected behaviour of the model. The name of the parameters, their ranges and their lower and upper bounds are detailed in Fig. 16.2a; Fig. 16.2a includes also the values originally considered in [6].

A likely hypothesis is that there is no unique but several combination of these parameters that are able to produce network activity as in Fig. 16.1b, d, in qualitative agreement with experimental data [7]. To test this hypothesis we need to find different parameter sets that are in agreement with the observed experimental results. To this end, we need to define two major elements: (1) a method to evaluate the correspondence of network function with the expected dynamics (Fig. 16.1b, d) for each parameter set, and (2) a method to explore the parameter space to search efficiently for parameter sets that match the required function optimally. In addition, and considering the computational costs implicated in the exploration of this high-dimensional parameter space, we designed a heuristic approach to extend the search by exploring the linear relations between good candidate solutions. We termed our implementation of this procedure “*Neuroswarm*” because it used a Particle Swarm algorithm to search for optimal solutions, but it can be readily extended to other optimization procedures. In the following we describe each of the steps involved.

16.3.1 Evaluating a Parameter Set

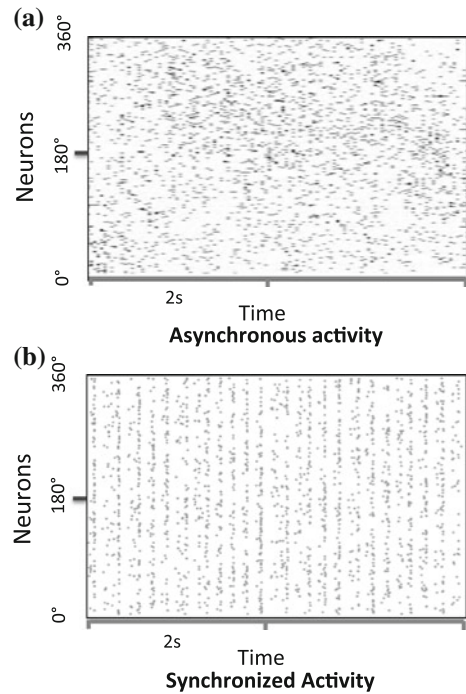
We defined a set of network activity properties expected from a “feasible” parameter set, and we specified each such property in a mathematical fitness function of the simulated network activity. Thus, after each simulation we could compute one fitness value for each of the expected network activity properties so that we could then compute a single evaluation value for each parameter set as a weighted sum of all these fitness values. The overall evaluation function is a cornerstone in the searching process as it defines the landscape and therefore it is largely associated to the

difficulty of the search. Evaluations are designed to have linear components to allow the optimization over gradients. Fitness functions were computed from the spiking times of only excitatory neurons in each simulation where for every parameter set we ran five simulations with different initial random seeds. The five fitness functions that we employed were the following:

- FF1. *Maximum firing rate in spontaneous activity*: a maximum firing rate of 8 sp/s in *spontaneous activity*. A linear fitness function taking values from 0 to 1 starts penalizing maximum firing rates of 5 sp/s (penalization of 0) and penalizes with a maximum value of 1 firing rates of 8 sp/s and above. We ran 5 different network simulations with different random seeds and took the largest penalization as fitness value here. See Fig. 16.3a.
- FF2. *Percentage of silent neurons in spontaneous activity*: the amount of excitatory neurons being silent (firing rate of 0 sp/s) during *spontaneous activity* is penalized. The fitness function takes the fraction of silent excitatory neurons. The maximum fraction computed from 5 simulations is taken as the FF2 evaluation. See Fig. 16.3a.
- FF3. *Maximum and minimum firing rate in persistent activity*: for excitatory cells targeted by the stimulus in *persistent activity* trials, the preferred peak firing rate is in the range 30 sp/s to 80 sp/s. A linear fitness function starts penalizing maximum firing rates of 80 sp/s (penalization of 0) up to 100 sp/s (penalization of 1). Similarly, the linear fitness function starts penalizing minimum firing rates of 30 sp/s (penalization of 0) up to 10 sp/s (penalization of 1). See Fig. 16.3b. The penalization is computed for 5 simulations and the maximum is selected.
- FF4. *Asynchronous activity*: parameter sets that generate extremely synchronized activity patterns during *spontaneous* and/or *persistent activity* are also penalized. For 5 simulations of each type we compute the average binary Pearson correlation between all pairs of neurons computed over windows of 4 ms. The maximum among the 5 values is selected as FF4. In Fig. 16.4 we provide examples of what we consider asynchronous activity (upper panel) and synchronized activity (lower panel).
- FF5. *Homogeneity*: we run five simulations with different random seed for each of the two conditions of interest: *spontaneous activity* (Fig. 16.1b) and *persistent activity* (Fig. 16.1d), and we penalize networks that did not provide stable results across these five simulations in each case. The penalization is computed by subtracting to 1 the average of the p-values computed from comparing the firing rate distributions across pairs of simulations by the Kolmogorov-Smirnov test; a value close to 0 denoted highly correlated distributions.

Finally, each fitness function had an associated weight to compute the total fitness evaluation value (TFEV) as its weighted average. We found that it was necessary to assign the highest weights to FF3 (5×10^3) and FF2 (10^3) fitness functions described above in order to find satisfactory solutions; the weights assigned to other fitness function were 10^2 , 10^2 and 10^0 , respectively for FF1, FF4 and FF5. Thus, for each parameter set, 5 simulations were run in each condition (spontaneous activity and persistent activity), one fitness value was obtained by evaluating the ensuing spiking activity for each of 5 fitness function and this was all combined in one single

Fig. 16.4 Asynchronous versus Synchronized neuron activity. Rastergrams showing examples of what we consider **a** asynchronous and **b** synchronized activity



evaluation value (TFEV) that characterized how well activity for this parameter set matched the required working memory function depicted in Fig. 16.1. The objective was then to find those parameter sets that yielded the lowest TFEV.

16.3.2 Exploring the Parameter Space

We sampled the parameter space in order to find parameter sets which yielded simulations with stable spontaneous activity and tuned persistent activity as illustrated in Fig. 16.1 (we denote these sets as *feasible parameter sets*, *FPS*). We assessed network function for each parameter set using the fitness functions defined above. Because of non-linearities both in the model and in the parameter set evaluation, no exact algorithm can be used to find the parameters that minimize the evaluation value. We used instead a heuristic algorithm to search for FPS: the Particle Swarm Optimization Algorithm (PSO; [12]). As a black-box optimization algorithm, the PSO can operate with any fitness function and it was originally designed to search optimally in a hyperspace of real numbers [12]. More detailed description of the algorithm is provided in [4] elsewhere in this volume. In brief, each model instantiation defined by a particular choice of the values of 16 parameters (Fig. 16.2a) represents one particle in a 16-dimensional space in the PSO. We run simulations

for each network model (particle) and compute the TFEV value for that particle. By working simultaneously with a large number of particles, PSO establishes parameter updates that move all particles towards the particle with the best TFEV value, while also attracting each particle towards its previous best evaluation. This optimization process has been shown to be effective in several applications [19] and the authors of this chapter have shown its usefulness in several cases in Neuroscience [2, 3].

Due to the high computational demands (in our equipment each network simulation ran for approximately 15 min and each iteration of the PSO algorithm required 500 simulations, for near 50 iterations in total), we executed all simulations in parallel in a computational facility with 200 available CPUs and with a modified parallelised code (Grid-SuperScalar technology at Barcelona Supercomputing Center—Centro Nacional de Supercomputacion for which we adapted the code).

To sample among the FPS we ran PSO 31 times, in each run we considered 50 particles, with a maximum of 50 iterations. An earlier stop was considered if the method was not finding a better-evaluated solution for more than 15 iterations. To define the working FPS we (1) first selected for each PSO run the best evaluated parameter set; then (2) we discarded 2 parameter sets because they were poorly evaluated (PSO did not find good solutions). Finally (3) non-filtered best solutions were used to define thresholds for the fitness functions (considering the maximum for each fitness function). Those thresholds (defined as 110% of the maximum values) were used to define the selection criteria that defined FPS from all the parameter sets evaluated in all PSO runs; by doing this we were able to recover more than one high-quality parameter set per PSO run. In the analysis, **FPS** denotes the original 29 parameter sets while **rFPS** denotes FPS extended with those high-quality recovered parameter sets.

16.3.3 Increasing the Set of Solutions

Once we had FPS we sought to explore the relation between them by simulating networks (parameter sets) that interpolated between them. We designed an exploratory greedy search of the parameter space. For each pair of the best solutions of the 31 PSO we investigated 9 equidistant points in the linear path between them in parameter space. We call these linearly interpolated parameter sets **Llps**. Figure 16.5 shows an example of the methodology, C1, C2 and C3 are the best solutions, the lines represent the shortest path (line) in 16 dimensions between each pair; each cross represent a new parameter set to be evaluated, a (un)filled cross denotes a (un)feasible model. By using this procedure we can evaluate if different solutions form part of a large continuous region of solutions or else if they are separated by regions without solutions, suggesting that they could constitute qualitatively different solutions.




Fig. 16.5 Extending the solution space. All the linear paths between every pair of PSO’s best solutions (FPS) are evaluated. For each linear path 9 equidistant points are evaluated. We observed that no single path contained all parameter sets evaluation as good (see “*Truncated straight line*”). In addition, a few possible new clusters of solutions were identified in some paths (see “*Possible new cluster*”)

16.3.4 Testing the Allowed Range of Parameters

One concern in this whole procedure is that one needs to define beforehand an allowed range for each of the parameters explored (see Fig. 16.2a), and this may limit the capacity of the search algorithm to find the best solutions. We performed one analysis to test this, and we iterated the procedure if we found that some of the ranges needed to be expanded. We plotted the density functions of the FPS (Fig. 16.6, black continuous), FPS extended with new solutions by linear paths (Lips, Fig. 16.6 black discontinuous), and the set of all parameter sets explored (Fig. 16.6, grey). By analysing FPS density plots we can confirm that no parameter had a large number of FPSs clustered at one of the imposed range limits, so that parameter boundaries did not seem to be a limiting factor to find FPS. Interestingly, densities are in many cases centred in the point equidistant between the lower and upper bound; we consider two explanations for this observation: (a) it may reflect a bias generated by the range definition (centred values are explored more often in random trajectories) and (b) this may denote parameters that have no large effects on the network described by the




Fig. 16.6 Ranges of parameter in good solutions. For each parameter we evaluated the density of the values observed in all parameter sets evaluated (*grey*), FPS (*black, continuous line*), and rFPS (*black, discontinuous line*)

parameter set. Importantly, despite the benefits of extending FPS by linear paths, we consider that **LIPs** reflect biases and have to be considered carefully when statistical analysis in parameter space are conducted.

16.4 Results

We applied the procedure described above and we found 93 different networks that could produce spatial working memory function as described in Fig. 16.1; 29 were identified in FPS and the rest were included through linear path extension (LIPs, Fig. 16.5). We then analysed these solutions with the aim of identifying relevant characteristics of the solutions and parameters associated to them.

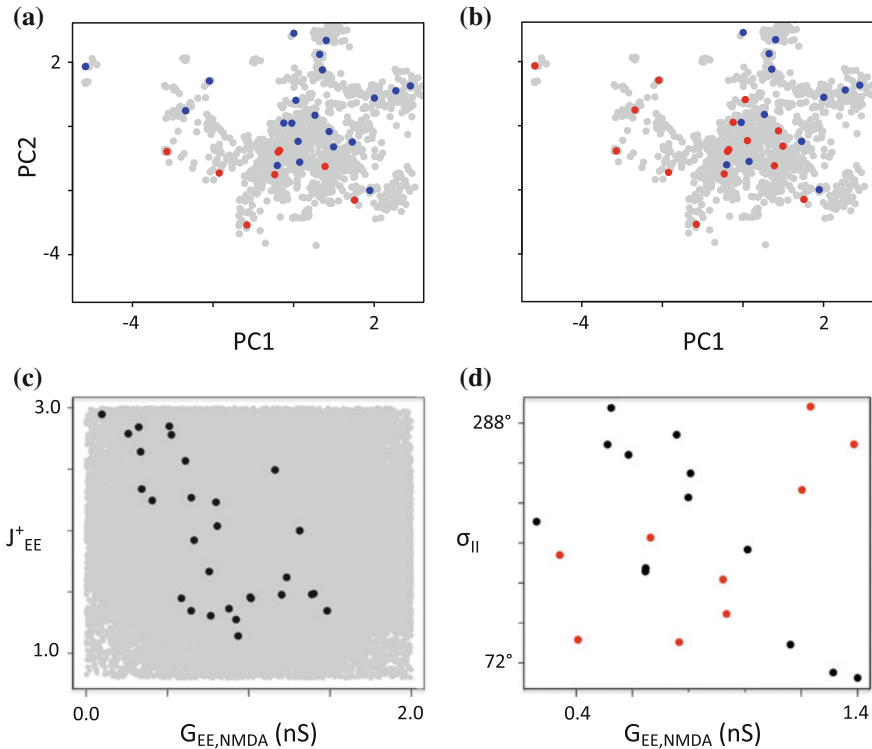


Fig. 16.7 Identification of solution grouping and compensatory mechanisms. **a** Shows the FPS's PCA with all parameter sets evaluated plotted in *gray* in the plane of the 2 first principal components. *Red dots* are FPS with two peaks in the firing rates of inhibitory neurons (see Fig. 16.7a, b) and *blue dots* are FPS with one single peak in inhibitory neuron activity (*blue*, see Fig. 16.7c, d). **b** Similar to **(a)** but *blue dots* mark FPS where inhibitory neurons had activity centred around the stimulus during persistent activity (Fig. 16.7a, d) and *red dots* mark FPS for peak inhibitory activity opposite to stimulus location (Fig. 16.7b, c). **c** Identified Compensatory mechanism in FPS between $G_{EE,NMDA}$ and J_{EE}^+ . *Black dots* represent solutions in FPS, while *grey dots* represented discarded parameter sets evaluated during the parameter search. **d** Opposite compensatory mechanism between $G_{EE,NMDA}$ and σ_{II} : solutions with inhibitory activity centred around the stimulus had a negative correlation between $G_{EE,NMDA}$ and σ_{II} ($\rho = -0.66$, *black dots*), while this correlation was positive for solutions with inhibitory activity maximal 180° away from the stimulus ($\rho = 0.69$, *red dots*)

16.4.1 2-D Representation of FPS to Detect Structure in the Solutions

We applied Principal Component Analysis (PCA) to the parameter sets explored by our search algorithm (Fig. 16.7a). PCA analysis (see [4]) allowed dimensionality reduction from the 16 parameters to the 2 first principal components (which accounted for 33% of the variance). This bi-dimensional representation allowed us to identify that many of the explored parameter sets clustered around the best solutions (FPS,

Table 16.1 Principal component analysis of FPS

Parameter	PC1	PC2	PC3
$G_{EE,AMPA}$	0.31	-0.27	0.33
$G_{EE,NMDA}$	0.35	-0.25	-0.35
$G_{EI,AMPA}$	0.31	-0.02	-0.15
$G_{EI,NMDA}$	0.03	0.16	-0.56
G_{IE}	0.26	-0.37	0.02
G_{II}	-0.23	-0.06	-0.40
σ_{EE}	-0.23	-0.05	-0.23
σ_{EI}	-0.02	0.20	0.05
σ_{IE}	0.07	0.10	-0.10
σ_{II}	0.08	-0.08	0.27
J_{EE}^+	-0.23	0.35	0.29
J_{EI}^+	0.42	0.16	0.04
J_{IE}^+	0.32	0.33	-0.21
J^{+II}	-0.19	-0.44	-0.02
$g_{ext,E}$	0.16	-0.27	-0.01
$g_{ext,I}$	-0.33	-0.35	-0.08

shown in red or blue). We may consider this initial clustering to be a consequence of the searching algorithm; however, from the observations during LIps computation (Fig. 16.5) we observed that many linear combinations of FPSs were not considered with quality enough to be part of the FPS; therefore we may conclude that FPS is non-convex (but we cannot conclude anything about parameter connectivity in the topological sense [18]).

The loadings of the 3 first PCA components (which accounted for 46% of the variance) can be found in Table 16.1. Careful inspection of these loadings could help identify what parameters of our simulation were most informative in distinguishing between the different FPS. We found that connectivity strengths, especially among excitatory neurons, and the tuning strengths of all neuron connectivities were discriminating factors between solutions, while the width of the connectivities did not seem to differentiate them significantly.

16.4.2 Clustering Solutions in the Parameter Space

Next we aimed to identify if solutions that are close in parameter space were sharing specific characteristics in their network activity. As a first exploratory approach we considered the solutions as they were mapped on the 2-D PCA reduction. We investigated the rastergrams and firing rates of excitatory and inhibitory neurons. FPS are clearly separated into subgroups if we inspect the activity of inhibitory




Fig. 16.8 Characterization of inhibitory rasters. Four types of inhibitory rastergrams are shown. In each case, 5 simulations initiated with different random seeds are shown, both for tuned persistent activity (PA 1–5 in the *upper panels*) and for spontaneous activity trials (SA 1–5 in the *lower panels*). These four examples illustrate: **a** single peak of activity centred at 180° ; **b** single peak of activity not centred at 180° ; **c** two peaks of activity, not centred at 180° ; and **d** two peaks of activity centred at 180°

neurons. We considered two characterizations of inhibitory rasters: (1) CvsNC: centred around the presented stimulus (180° , Fig. 16.7a, d) or not centred (Fig. 16.8b, c); and (2) 1vs2: 1 peak (Fig. 16.8a, b) or 2 peaks (Fig. 16.8c, d). In both characterizations there is a separation between groups in the PCA, see Fig. 16.7a, b.

We identified with a Kruskal-Wallis test that J_{EI}^+ was significantly different (after multiple testing correction) in comparison CvsNC, while J_{II}^+ changed significantly in the 1vs2 comparison. Interestingly, J_{EI}^+ and J_{II}^+ were the most relevant parameters in the first and second component of the PCA, respectively.

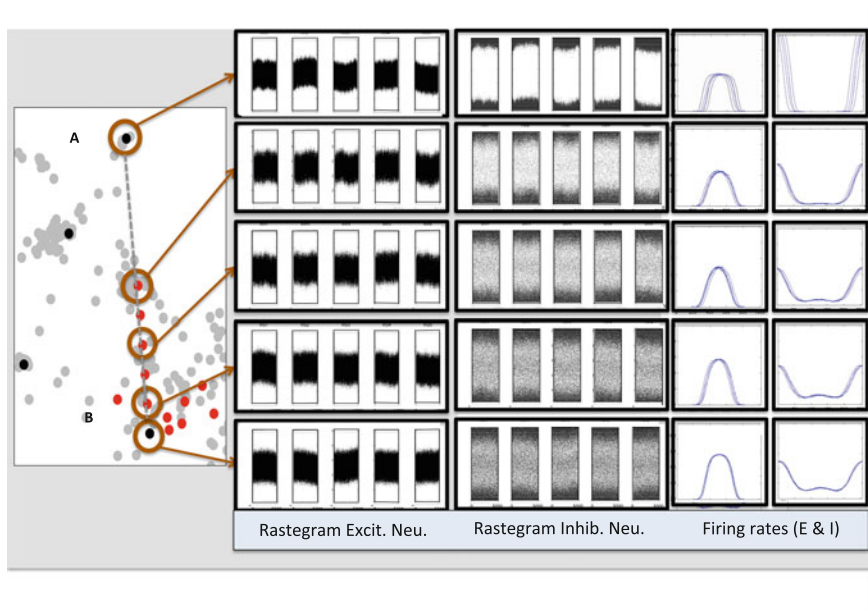


Fig. 16.9 PFC model solutions are linearly unconnected, but new solutions are Inferred: example of a novel cluster in the Linear Path between two FPS-B. The linear path (*grey dashed line*) between two best solutions (**A** and **B**) from PSO is inspected at 9 interpolating points. The *red points* in the *line* denote new feasible solutions. For **A**, **B** and 3 new solutions the rastergram and firing rates of excitatory (*left*) and inhibitory (*right*) are shown. Rastergrams are provided for 5 simulations with different seeds. Average firing rates over the delay period for excitatory and inhibitory neurons are plotted in the *rightmost* panels, each simulation with a different line to show diffusion of activity in different trials

16.4.3 LIps Identifies Transition Between Solution Types

In LIps we extended FPS by investigating linear combinations of FPS parameter sets (Fig. 16.5). Through this methodology we observed that a straight line of feasible solutions connected no pair of solutions; therefore we validate the non-convexity of FPS.

In LIps we identified novel clusters of high-quality parameter sets; importantly, we found that in all cases the novel cluster firing rates are similar to those of one of the pair of solutions that define the linear combination, possibly suggesting a non-linear connection between them. Figure 16.9 shows an example of a novel cluster. We observe that in this case the characteristics of the behaviour of the novel cluster are similar to those found in Solution B. However, we observe a gradual transformation that reduces the high bump diffusion observed in solution A to the non-existent diffusion of solution B. We conclude that the study of the linear paths helps in the

identification of those parameters associated to specific characteristics. Interestingly, the major differences between Solution A and Solution B are not occurring in pyramidal neurons but in inhibitory neurons.

16.4.4 Studying Compensatory Mechanisms

We were interested in observing whether there were compensatory mechanisms between the parameters that may help to alleviate the fine-tuning problem biologically. We measured this by computing Spearman non-parametric correlation between the values of the parameter sets. We considered only parameter sets in FPS, without including the linear combinations of those, to prevent introducing an artificial bias. Table 16.2 shows all correlations. Among them, we selected those that were statistically significant, strong absolute correlations ($|\rho| > 0.5$). We found three compensatory mechanisms of interest: (1) $G_{EE,NMDA}$ versus J_{EE}^+ ($\rho = -0.63$, Fig. 16.8c); (2) J_{EI}^+ versus J_{IE}^+ ($\rho = 0.62$, expected positive correlation); and (3) J_{EI}^+ versus $g_{ext,I}$ ($\rho = -0.53$). These compensatory mechanisms can be interpreted as addressing critical principles of network operation that sustain the required working memory function. Thus, the first of these compensations, relating negatively the strength of $G_{EE,NMDA}$ and the tuning of excitatory connections J_{EE}^+ , seems to work to keep a sufficient level of excitation locally: if the strength of excitatory synapses is strong enough, local potentiation relative to other synapses is not required and may actually lead to unreasonably high firing rates. The second compensatory mechanism (positive relation between the tuning of excitatory-to-inhibitory and inhibitory-to-excitatory connections) suggests a control of the spatial specificity in the local excitation-inhibition loop. Finally, the third compensation addresses the local excitation-inhibition balance: if inhibitory cells have high spontaneous activity (as a result of increased $g_{ext,I}$), local excitation from excitatory cells should be reduced. These particular requirements for the operation of working memory networks have not been scrutinized in previous studies and suggest new avenues to study the conditions for working memory operation in such networks.

Finally, we also searched for compensatory mechanisms that were opposite in different sub-groups as defined by 1vs2 and/or CvsNC. We found one such opposite compensatory mechanism relating $G_{EE,NMDA}$ and σ in the CvsNC classification (Fig. 16.7d). Thus, for solutions in which inhibitory neurons presented a stable firing pattern with maximal activity centred around the stimulus location (i.e. congruent with maximal activity in excitatory neurons), these two parameters correlated negatively, while they correlated positively for those solutions where inhibitory neurons had activity profiles peaking at the same location as excitatory neurons (Fig. 16.7d). This relationship is intriguing and suggests a different mechanistic link between within population interactions (inhibition to inhibitory neurons, and excitation to excitatory neurons) for each of these two network function organization (CvsNC). This remains to be explored in depth in a future study.

Table 16.2 Correlation structure between parameters in FPS

	$G_{EE,AMPA}$	$G_{EE,NMDA}$	$G_{EL,AMPA}$	$G_{EL,NMDA}$	G_{IE}	G_{II}	σ_{EE}	σ_{EI}	σ_{IE}	σ_{II}	J_{EE}^+	J_{EI}^+	J_{IE}^+	J^{+II}	$g_{exl,E}$	$g_{exl,I}$
$G_{EE,AMPA}$	1.00	0.21	0.18	-0.42	0.45	-0.40	-0.23	-0.06	0.06	0.02	-0.36	0.29	0.03	0.12	0.14	-0.13
$G_{EE,NMDA}$	0.21	1.00	0.19	0.41	0.45	-0.03	0.01	-0.02	-0.10	0.02	-0.63	0.40	0.39	0.13	0.23	-0.11
$G_{EL,AMPA}$	0.18	0.19	1.00	0.20	0.21	0.12	-0.12	-0.17	0.20	-0.05	-0.09	0.05	0.21	-0.34	0.46	-0.19
$G_{EL,NMDA}$	-0.42	0.41	0.20	1.00	-0.33	0.37	0.04	-0.09	0.07	-0.14	-0.15	0.04	0.22	-0.08	-0.02	0.04
G_{IE}	0.45	0.45	0.21	-0.33	1.00	0.18	-0.17	0.22	-0.12	0.02	-0.44	0.15	0.01	0.15	0.27	0.02
G_{II}	-0.40	-0.03	0.12	0.37	0.18	1.00	0.03	0.07	-0.10	-0.31	-0.25	-0.40	-0.21	-0.02	-0.26	0.22
σ_{EE}	-0.23	0.01	-0.12	0.04	-0.17	0.03	1.00	0.08	-0.09	-0.24	-0.16	-0.32	0.00	0.27	-0.12	0.16
σ_{EI}	-0.06	-0.02	-0.17	-0.09	0.22	0.07	0.08	1.00	-0.01	-0.08	0.01	0.08	0.06	-0.20	-0.17	-0.07
σ_{IE}	0.06	-0.10	0.20	0.07	-0.12	-0.10	-0.09	-0.01	1.00	-0.32	0.07	-0.04	0.19	-0.15	0.35	-0.16
σ_{II}	0.02	0.02	-0.05	-0.14	0.02	-0.31	-0.24	-0.08	-0.32	1.00	0.07	0.03	-0.10	-0.10	0.19	0.25
J_{EE}^+	-0.36	-0.63	-0.09	-0.15	-0.44	-0.25	-0.16	0.01	0.07	0.07	1.00	-0.05	-0.11	-0.13	0.12	0.00
J_{EI}^+	0.29	0.40	0.05	0.04	0.15	-0.40	-0.32	0.08	-0.04	0.03	-0.05	1.00	0.62	-0.11	0.03	-0.53
J^{+IE}	0.03	0.39	0.21	0.22	0.01	-0.21	0.00	0.06	0.19	-0.10	-0.11	0.62	1.00	-0.43	-0.04	-0.44
J^{+II}	0.12	0.13	-0.34	-0.08	0.15	-0.02	0.27	-0.20	-0.15	-0.10	-0.13	-0.11	-0.43	1.00	0.23	0.45
$g_{exl,E}$	0.14	0.23	0.46	-0.02	0.27	-0.26	-0.12	-0.17	0.35	0.19	0.12	0.03	-0.04	0.23	1.00	0.17
$g_{exl,I}$	-0.13	-0.11	-0.19	0.04	0.02	0.22	0.16	-0.07	-0.16	0.25	0.00	-0.53	-0.44	0.45	0.17	1.00

16.5 Conclusions

We show here that large-scale biological simulations of neural networks for specific cognitive functions can be evaluated for generality using an optimization procedure in their high-dimensional parameter space. Typically, these simulations are very unconstrained and generality has been tested using mathematical simplifications in mean field formulations. While this approach is indeed very general, the initial assumptions on the simplifications to use may constrain the validity of the results to a subset of possible solutions. Our computational workflow approach can classify what parameters are more critical for the identified behaviour, and what compensatory or synergistic associations between parameters are imposed by the required behaviours. These relationships can guide simplifications for further mathematical analysis.

Our study provides two major conclusions arguing for the exploration of multiple high-quality parameter sets (or solutions), which support and extend those shown in Marder et al. [17]. First, to consider a single solution (such as a single set of parameters fitting the expected data) provides limited insights on a given model: *are we sure that the conclusions observed in a single solution (parameter set) are true for all feasible parameter sets?* Secondly, the characterization of the set of feasible parameter sets provides a deeper understanding of the model because it can (a) characterize and enumerate the set of hypotheses that cannot be rejected based on the present experimental and theoretical understanding of the phenomenon; (b) identify specific experiments that can be most informative in distinguishing between these pending alternative scenarios; and (c) provide insights about what parameters of the models are critical, and could be used as targets for specific manipulations in subsequent simulations.

We have made use of the approach proposed in the present chapter in other computational works [2, 3]; where this parameter exploration procedure was used to confirm that a specific property observed in one model was general to the class of possible models constrained by experimental and behavioural results. See also [5] for other parameter exploration-based approaches to generate hypothesis.

References

1. Achard, P., De Schutter, E.: Complex parameter landscape for a complex neuron model. *PLoS Comput. Biol.* **2**(7), e94 (2006)
2. Ardid, S., Wang, X.J., Gomez-Cabrero, D., Compte, A.: Reconciling coherent oscillation with modulation of irregular spiking activity in selective attention: gamma-range synchronization between sensory and executive cortical areas. *J. Neurosci. Off. J. Soc. Neurosci.* **30**(8), 2856–2870 (2010)
3. Cano-Colino, M., Almeida, R., Gomez-Cabrero, D., Artigas, F., Compte, A.: Serotonin regulates performance nonmonotonically in a spatial working memory network. *Cereb. Cortex (New York, N.Y.:1991)* **24**(9), 2449–2463 (2014). doi:[10.1093/cercor/bht096](https://doi.org/10.1093/cercor/bht096)
4. Cedersund, G., Samuelsson, O., Ball, G., Tegnér, J., Gomez-Cabrero, D.: Optimization in biology parameter estimation and the associated optimization problem. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)

5. Cedersund, G.: Prediction uncertainty estimation despite unidentifiability: an overview of recent developments. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
6. Compte, A., Brunel, N., Goldman-Rakic, P.S., Wang, X.J.: Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex* **10**(9), 910–923 (2000)
7. Funahashi, S., Bruce, C.J., Goldman-Rakic, P.S.: Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.* **61**(2), 331–349 (1989)
8. Gomez-Cabrero, D., Compte, A., Tegnér, J.: Workflow for generating competing hypothesis from models with parameter uncertainty. *Interf. Focus* **1**(3), 438–449 (2011)
9. Hodgkin, A.L.: Chance and Design in Electrophysiology: An informal account of certain experiments on nerve carried out between 1934 and 1952. *J. Physiol.* **263**(1), 1–21 (1976)
10. Hodgkin, A.L., Huxley, A.F.: Currents carried by sodium and potassium ions through the membrane of the giant axon of *Loligo*. *J. Physiol.* **116**, 449–472 (1952)
11. Hug, S., Schmidl, D., Bo Li, W., Greiter, M.B., Theis, F.J.: Bayesian model selection methods and their application to biological ODE systems. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
12. Kennedy, J., Eberhart, R.: Particle swarm optimization. In: *IEEE International Conference on Neural Networks Proceedings, Vols. 1–6, 1942–1948* (1995)
13. Kirk, P., Silk, D., Stumpf, M.P.H.: Reverse engineering under uncertainty. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
14. Lejon, A., Samaey, G.: Stochastic modeling and simulation methods for biological processes: overview. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
15. Mannakee, B.K., Ragsdale, A.P., Transtrum, M.K., Gutenkunst, R.N.: Sloppiness and the geometry of parameter space. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
16. Marder, E., Taylor, A.L.: Multiple models to capture the variability in biological neurons and networks. *Nat. Neurosci.* **14**(2), 133–138 (2011)
17. Marder, E., Tobin, A.E., Grashow, R.: How tightly tuned are network parameters? Insight from computational and experimental studies in small rhythmic motor networks. *Progr. Brain Res.* **165**, 193–200 (2007)
18. Munkres, J.R.: *Topology*, 2nd edn. Prentice Hall, New Jersey (2000)
19. Poli, R.: Analysis of the publications on the applications of particle swarm optimisation. *J. Artif. Evol. Appl.* **2008**, 1–10 (2008). doi:[10.1155/2008/685175](https://doi.org/10.1155/2008/685175)
20. Prinz, A.A., Billimoria, C.P., Marder, E.: Alternative to hand-tuning conductance based models: construction and analysis of databases of model neurons. *J. Neurophysiol.* **90**(6), 3998–4015 (2003)
21. Prinz, A.A., Bucher, D., Marder, E.: Similar network activity from disparate circuit parameters. *Nat. Neurosci.* **7**(12), 1345–1352 (2004)
22. Sunnåker, M., Stelling, J.: Model extension and model selection. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
23. Tucker, W.: Interval methods. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
24. Tuckwell, H.: *Introduction to Theoretical Neurobiology* (2 vols.). Cambridge Studies in Mathematical Biology. Cambridge University Press, Cambridge (1988)
25. Van Geit, W., Achard, P., De Schutter, E.: Neurofitter: a parameter tuning package for a wide range of electrophysiological neuron models. *Front. Neuroinform.* **1**, 1 (2007)
26. Van Geit, W., De Schutter, E., Achard, P.: Automated neuron model optimization techniques: a review. *Biol. Cybern.* **99**(4–5), 241–251 (2008)
27. Van Schepdael, A., Carlier, A., Geris, L.: Sensitivity analysis by design of experiments. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)

Chapter 17

Prediction Uncertainty Estimation Despite Unidentifiability: An Overview of Recent Developments

Gunnar Cedersund

Abstract One of the most important properties of a mathematical model is the ability to make predictions: to predict that which has not yet been measured. Such predictions can sometimes be obtained from a simple simulation, but that requires that the parameters in the model are known from before. In biology, the parameters are usually both not known from before and not identifiable, i.e. the parameter values cannot be determined uniquely from available data. In such cases of unidentifiability, the space of acceptable parameters is large, often infinite in certain directions. For such large spaces, sampling-based approaches that try to characterize the entire space have difficulties. Recently, a new type of alternative approaches that circumvent this characterization problem has been proposed: where one only searches those directions in the space of acceptable parameters that are relevant for the uncertainty of a particular prediction. In this review chapter, these recently proposed methods are compared and contrasted, both regarding theoretical properties, and regarding user experience. The focus is on methods from the field of systems biology, but also methods from biostatistics, pharmacodynamics, and biochemometrics are discussed. The hope is that this review will increase the usefulness and understanding of already proposed methods, and thereby help foster a tradition where predictions only are deemed interesting if their uncertainties have been determined.

Keywords Prediction uncertainty · Systems biology · Core predictions · Prediction profile likelihood · Ordinary differential equations · Bayesian methods · Cluster Newton · Neutral parameters

G. Cedersund (✉)

Department of Clinical and Experimental Medicine, IMT, Linköping University,

58185 Linköping, Sweden

e-mail: gunnar.cedersund@liu.se; gunnar.cedersund@gmail.com; +46(0)702-512323

G. Cedersund

Department of Biomedical Engineering, IMT, Linköping University,

58185 Linköping, Sweden

© Springer International Publishing Switzerland 2016

L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,

Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,

DOI 10.1007/978-3-319-21296-8_17

17.1 Introduction

Mathematical modelling has been a part of modern science already since its early foundation. In those early days, different hypotheses concerning the movements of the planets were incorporated in mathematical models, and the ability of these hypotheses to describe available data and to produce reliable predictions were used as the basis for choosing between the hypotheses. Those two components—rejections and predictions—are still the two corner stones of mathematical modelling as it now at last is getting more widely used also within biology (Fig. 17.1a). In this chapter we will deal with the latter of these two components, predictions. More specifically, we

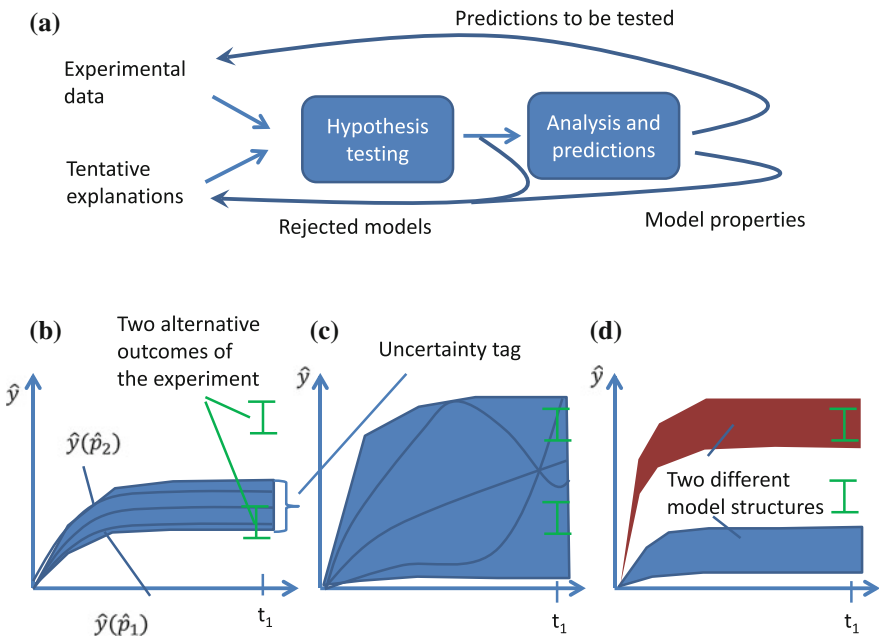


Fig. 17.1 **a** The two main steps of model-based data analysis: hypothesis testing and prediction analysis. The first step has been dealt with in previous chapters (e.g. [26]), and leads to rejections or tentative acceptances of models. The second step is dealt with in this chapter, and identified predictions feeds back to the design of new experiments, or in general to new knowledge regarding the original hypotheses. **b–d** If you know the uncertainty of the prediction, you can in before-hand guarantee that the experiment will give you something, independently of the outcome. In **(b)**, the experiment is done to test the model. Only if the experimental data (*green*) lies within the prediction's uncertainty tag (*blue area*) will the model be accepted. In **(c)**, the experiment is done to further determine the parameters: the more uncertain the prediction, the more the space of the acceptable parameters will be constrained by the data. In **(d)**, the experiment is done to distinguish between the two models: independently of the experimental outcome, at the most one of the models will remain non-rejected after the new data point has been collected. The main result reviewed in this chapter regards how to go beyond characterization of all individual parameter trajectories (such as $\hat{y}(\hat{p}_1)$ and $\hat{y}(\hat{p}_2)$ in **(b)**) to methods that only looks for those parameters that give extreme predictions

will make an overview of recent advances in methods that allow us to produce useful predictions also in biology: predictions that take into account the known uncertainties in data and prior knowledge, and convert these to predictions with uncertainty tags (Fig. 17.1b). Importantly, these new methods work and may identify well-determined predictions also in the case of unidentifiable parameters.

Several of the herein presented methods are only a few years old, and it is therefore important to emphasise how important a development these methods are. This importance comes from the fact that the uncertainties in biology and medicine are so much bigger than corresponding uncertainties within physics, and that the classical methods for generating predictions with uncertainty tags therefore no longer are useful. These uncertainties are of at least two types: structural and parametrical. The structural uncertainties are bigger in biology, because there are more unknown details that make up each hypothesis. In physics, traditional hypotheses have been about the laws of nature, and e.g. whether these follow relativistic or Newtonian mechanics. In contrast, in biology, a hypothesis is often concerning the structure of a part of a biological network, and most meaningful hypotheses regarding such a network are in themselves containing many structural uncertainties. For instance, a hypothesis may be that a certain type of feedback generates a specific behaviour, such as an overshoot [4]. This hypothesis has many potential implementations, and is thus corresponding to many sets of equations, e.g. corresponding to different assumptions regarding the kinetic laws describing the involved reactions. The second type of uncertainty concerns the values of the parameters. In other words, even if one has made all structural decisions, e.g., concerning whether to use mass action or Michaelis-Menten rate law expressions, the values of the kinetic parameters that are to be used in those rate expression are still unknown. This again stands in stark contrast to the situation in physics, where most parameters are natural constants, which can be determined once and for all, and which already have been calculated prior to the modelling.

This second problem, that of uncertain parameters, is often referred to as unidentifiability, and it has been extensively studied in the literature. Identifiability is often considered to exist in two types: structural and practical. Structural identifiability has to do with the structure of the equations, and asks the question whether the parameters in principle can be obtained from those equations, if measurement uncertainty and insufficient excitation not would be an issue. Although much beautiful theory has been produced for structural identifiability examinations, e.g. based on differential geometry [18, 25], this question is not too important for biological applications, where measurement uncertainty and insufficient excitation often are major problems. Therefore, the second type, practical unidentifiability analysis, which takes into account the specific data set at hand, is much more relevant. Much early analysis of practical identifiability has been concerned with the eigenvalues of the Hessian (i.e. the second-derivative) of the cost function. If the predicted output is linearly dependent on the parameters, this Hessian can be proportional to the inverse of the covariance matrix of the parameter uncertainties [17]. Unfortunately, for most biological problems, the assumption of linearity is not fulfilled, and the Hessian then only gives a local measure of parameters, at a probably non-unique point in the parameter space. The result of such a Hessian-based analysis is therefore hard to use.

For this reason, methods that make global assessments are more interest. There are attempts to produce global versions of the Hessian approach (e.g. [24]), but these do not solve the fundamental problem of linearity. A more promising approach, which also is intuitive and easy to implement, is the method based on the profile of the likelihood, simply called profile likelihood (PLH). This approach has been known for decades [21], but has not been used within the systems biology community until recently [23].

Several of the chapters in this book [11] have dealt with various ways to identify and handle uncertainty (e.g. [14]), but this is the only one that exclusively deals with predictions. Predictions are important to characterise, with uncertainty tags, for a wide variety of reasons. First and most importantly, predictions are what feeds back to the experiments, and therefore closes the experimental/theory loop (Fig. 17.1a), which produces an ever-evolving understanding of the system. Here it should be stated that both well-determined predictions (herein called core predictions) and poorly determined predictions (herein called suggestions or beach statements) are of interest: core predictions allow you to test the quality of the model (Fig. 17.1b), and beach statements are important because measuring them will produce a big improvement in the overall well-determination of the parameters in the model (Fig. 17.1c). Note that an important special case of the first usage occurs when you have two different core predictions from two different models (Fig. 17.1d); then the experiment ensures that at least one of the model structures will be rejected, independently of the outcome of the experiment. Second, even those predictions that cannot, or will not, directly be tested experimentally are interesting because they do provide a deeper understanding of the studied model: what are the key properties that have to be fulfilled in this system in order for it to produce the data, how is the model actually producing the observed behaviour? Third, reliable predictions are useful in a medical context. If some specific prediction of the model has been shown to be useful as clinical markers, e.g. to correctly diagnose a patient, it is important that that prediction is obtained with a correct degree of certainty. Fourth, well-determined predictions in a model means that it is useful for interacting with other models, in a supermodel. Say for instance that the first model describes the internal dynamics of a specific organ, and that its input-output profile is well-determined from experimental data. Then this organ model can be used in a hierarchical model, incorporating several organs and their cross-talk, even though some of the internal predictions in the model may be uncertain [22]. Finally, note that prediction uncertainty is not the same as parameter uncertainty: all parameters may be undetermined in the model, even though many well-determined predictions exists. For all these reasons, methods that deals with prediction uncertainty, and that can handle the in biology common situation of unidentifiability, are of utmost importance.

All in all, to find predictions with uncertainty is therefore an important subject, and it is therefore interesting to look at some of the recent developments that have been done within the field. The rest of the chapter is structured as follows. First some basic notations regarding ODEs are introduced. Second, the most straightforward and simple approaches are introduced, which simply collects all found parameters that seem acceptable. Third, the recent developments regarding methods based on profiling the likelihood and modified optimization problems are introduced, and the

different variants are contrasted and discussed. These new methods are then also compared to three other important related approaches: Bayesian methods, cluster Newton, and neutral parameters. The chapter ends with a summary, and with a discussion of what the price is of not bothering with prediction uncertainty.

17.2 Basic Notations

The following presentation is done in the framework of ordinary differential equations (ODEs), but most concepts and methods hold equally well for all systems for which there exists a predictor, and the ability to form a likelihood function. Now follow quite a few notations. An ODE-based model is henceforth described by the following notations:

$$\dot{x} = f(x, p_x, u) \quad (17.1a)$$

$$x(0) = x_0 \quad (17.1b)$$

$$\hat{y} = g(x, p_x, p_y, u) \quad (17.1c)$$

where x are the states, usually describing the concentration or amount of various substances; where \dot{x} represents the derivative of the states with respect to time, and where f is the non-linear smooth function used to calculate these derivatives; where p_x are the parameters used to calculate f , usually corresponding to kinetic rate constants; where u is the input, which may depend on time, and which is usually known and controlled by the experimentalist; where $x(0)$ contains the values of the states at time $t = 0$, and where these values are described by the parameters x_0 ; where \hat{y} are the simulated model outputs corresponding to the measured experimental signals; where g is a smooth nonlinear function; and where p_y are parameters only appearing in the measurement equations. Note that x , u and y may depend on t , but that the notation is dropped unless the time-dependence needs to be especially stated, as in Eq. (17.1b). Note that at this stage there are three types of parameters, with potentially unknown values, and that all unknown parameters are collected in the parameter p , i.e. currently

$$p = (p_x, x_0, p_y) \quad (17.2)$$

Finally, the term parameter requires some further comments, since its usage opens for two interpretations. In some cases, a parameter means a point in the parameter space; in such cases the terms parameter point or parameter value will typically be used. Similarly, to refer to an individual element p_i in the parameter vector p , the term individual parameter will typically be used.

The three equations (17.1) form a model structure, which is denoted \mathcal{M} , and this model structure is turned into a specific model, denoted $\mathcal{M}(p)$, if the parameters are set to specific values. Now, assume that the N measurements are collected in a set Z ,

$$Z = \{y(t_i)\}_{i=1}^N \quad (17.3)$$

where $y(t_i)$, describes the measurement vector at time t_i . Unless otherwise stated, the null hypothesis is that the considered model structure is true, i.e. that

$$y(t_i) = \widehat{y}(t_i, p^0) + v \quad \forall i = 1 \dots N \quad (17.4)$$

where p^0 denotes the “true” parameter values (which exists if the null hypothesis is true), and where v is a noisy signal that follows some distribution denoted D . In this chapter, the term ‘explanation’ is used to describe a hypothesis that not only should produce correct predictions mathematically, but also do this using mechanisms that make sense biologically.

17.3 Simple Approaches to Identify Predictions with Uncertainty

The underlying reason why predictions are uncertain is that there are many parameter points that all correspond to an acceptable agreement with available data and prior knowledge, and therefore must be considered as non-rejected. The formal decision of whether a parameter is to be considered as rejected or not is taken using a cost function, denoted $V(p)$, and a cut-off, denoted δ (aspects of choosing $V(p)$ and δ are discussed in section “The cut-off level” below). In other words, if the set of non-rejected parameters is denoted \mathcal{A} , it can formally be defined as

$$\mathcal{A} := \{p : V(p) \leq \delta\} \quad (17.5)$$

where “:=” means that the left hand side is defined by the right hand side. With this set \mathcal{A} formally defined, a basic principle of identifying parameter uncertainties presents itself: simply find the parameter points within \mathcal{A} that gives a maximal and minimal value of the considered prediction. This is the principle of frequentist approaches, and it is contrasted to the Bayesian approaches later in this chapter. In the most simple approaches, the search for maximal and minimal values of the considered prediction within \mathcal{A} is simplified to the search for maximal and minimal values within an approximation of \mathcal{A} (Step 1 in Fig. 17.2). This approximation can be obtained by traversing \mathcal{A} in different ways, and then saving the encountered parameter points. In [4], this traversing was done using a modified optimization approach, which uses multiple simplexes instead of one, to more fully cover the parameter space. Various alternatives to ordinary optimization and parameter estimation approaches that can be used for this are reviewed e.g. in [8]. Another approach sometimes used is based on the conventional PLH [23]. As the profiles are traversed to search for parameter uncertainties, the parameter points are saved, and then analysed with respect to the considered prediction. However, the limitation of these approaches is that there is no guarantee that the most extreme predictions have been found. In contrast, such a

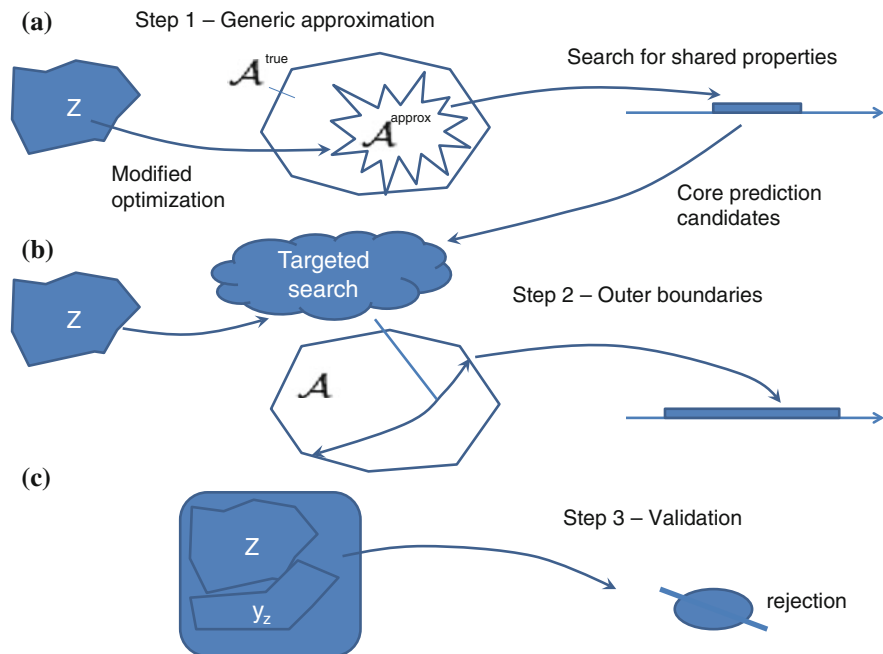


Fig. 17.2 The three main steps involved in prediction uncertainty assessments. **a** *Step 1* is simply to make a point-wise approximation of the true space of acceptable parameters, \mathcal{A} . In Bayesian approaches, this is the only step, and then the space is approximated so that the parameter points are most dense where the most likely parameters are. **b** *Step 2* lies at the heart of the most recent developments, which primarily are in focus in this chapter. One has then picked a specific prediction, probably because *Step 1* judged it to be well-determined, and then does a targeted approximation for that particular prediction. The most important things are to find the max and min values that the prediction can obtain, while still remaining in \mathcal{A} . **c** *Step 3* is in the core prediction approach simply a validation that the prediction rejection really holds, if the measured value lies outside the prediction uncertainty. In PPL, this step also involves assessments of profiles. Figure adopted from [6]

guarantee is obtained, at least in principle, by the newly developed approaches which modify the PLH to instead deal with prediction uncertainty, using various modified optimization problems.

17.4 PLH-Inspired Methods Based on Modified Optimization Problems

The now presented approach has recently been presented in different forms, based on different theoretical traditions, but ending up in similar but not identical methods. The following presentation is an attempt to bring these approaches together into a joint description, where similarities and differences more easily can be highlighted.

17.4.1 The Idea

Consider a specific prediction denoted z . This prediction can be anything that can be calculated using the given model structure and a specific parameter point, p . In other words, z must be a function that depends solely on the parameters, $z = z(p)$, but it can otherwise be anything of interest: e.g. a value of a state at a specific time-point, a value of a reaction rate at a specific time point, a single parameter value, or some other arbitrary function of these values, but that still depends on a given time point. The property z can sometimes also be considered as an entire time-series, but here we have encountered the first difference: that the property can only be a time-series in one of the settings, that based on core predictions.

In fact, the basic idea behind the core-prediction based approach presented in [7] was presented in the more general sense of z being a time-series. Nevertheless, when the full 3-step implementation of the core prediction approach was presented for specific examples in [6], the complications of studying a time-series was not considered. The inclusion of time-series are mainly concerned with the issue of dependencies and statistical interpretations of the results, and this whole issue is further discussed in a separate section below, entitled “Dependencies and the handling of predictions of time-series”. For now, we consider the more specific case of z being a single scalar valued function of parameters and states at a given time point, and consider the general idea as presented in [6]. This idea is given by the following equation

$$\Delta_z^{\max} = z(p_{\max}) \text{ where} \quad (17.6a)$$

$$p_{\max} = \arg \max_p \{z\} \text{ subject to } V(p) < \delta \quad (17.6b)$$

$$\Delta_z^{\min} = z(p_{\min}) \text{ where} \quad (17.6c)$$

$$p_{\min} = \arg \min_p \{z\} \text{ subject to } V(p) < \delta \quad (17.6d)$$

where Δ_z^{\max} is the upper boundary of the prediction, and Δ_z^{\min} is the lower boundary. In words, the idea behind this approach is thus simply the one already stated: to find the maximal and minimal values of z that exists within \mathcal{A} .

The basic idea behind the other approach is based on a theory known as prediction profile likelihood (PPL). The concept of PPL was first mentioned in 1956 [10], and the concept of prediction profile likelihood was first introduced in 1979 [20]. Since then, quite a few various alternatives to this approach have been introduced, as reviewed e.g. in [2]. A first version that is relevant to this discussion was uploaded on the public article database arXiv in [15], and then as a standard paper in [16]. The basic idea behind PPL is to extend the data series Z with a new data point, z^* , which is a measurement of the property z . Often the uncertainty of z^* is zero, and then the maximum likelihood approach simply boils down to maximizing the likelihood, or minimizing the cost $V(p)$, while fulfilling $z(p) = z^*$. This additional data point, z^* ,

then enters as an additional parameter, and one can fix this parameter to different values, while optimizing over the others, as in a traditional PLH approach.

Let us now see how these two ideas may be implemented in practice.

17.4.2 The Stepping

The core-prediction approach is formally a standard, albeit difficult, constrained optimization problem, and may therefore in principle be implemented using any method that can deal with such problems. However, since many of the most commonly used implementations of optimization algorithms cannot deal with such difficult constraints, the following simple implementation was suggested in [6]:

1. Initiate $w > 0$, $f_{\text{incr}} > 1$
2. $\hat{p} = \arg_p \min V_C(p)$, where $V_C(p) = V(p) + w \cdot C(z(p))$
3. if $V(p) > \delta$, return $z(\hat{p})$; else $w = w \cdot f_{\text{incr}}$, go to 2
4. *Output*: the maximum or minimum value that z can take while still agreeing with data

where $C(z)$ is chosen as a function that grows with z if you are seeking Δ_z^{\min} (e.g. $C(z) = z$ or $C(z) = \log(z)$) and where $C(z)$ is chosen as a diminishing function if you are seeking $\Delta_{\text{prop}}^{\max}$ (e.g. $C(z) = 1/z$ or $C(z) = -z$).

The central idea behind this little algorithm is that you increase the contribution from the maximisation (or minimisation) of the new cost function $C(\cdot)$ compared with the component from the old cost function $V(p)$, until the maximisation (or minimisation) cannot proceed without violating $V(p) < \delta$; the optimised parameter values where you stop would then in theory and for sufficiently small step-sizes, f_{incr} , lie at the border of the core prediction uncertainty, i.e. at Δ_z^{\max} (or Δ_z^{\min}). In other words, if there is an upper boundary for the prediction z , it is because that property is observable (i.e. can be identified from data [6]), i.e. because there is a tradeoff for the model between fitting to the data and maximizing that value.

In the other approach, based on PPL, one formulates a similar but not identical constrained optimization problem as in (17.6)

$$\hat{p} = \arg \min_p V(p) \quad \text{where } z(p) = z^* \quad (17.7)$$

This constrained optimization problem corresponds to the case where the added data point, z^* , has zero uncertainty. In practice they solve this equation by optimizing

$$\hat{p} = \arg \min_p V_C(p) \quad \text{where } V_C(p) = V(p) + \eta(z - z^*) \quad (17.8)$$

where η is a Lagrange multiplier, which is increased until the resulting optimization yields a sufficient fulfillment of the constraint $z(p) = z^*$. Once that fulfillment is

achieved, z^* is increased (or decreased) until the maximum (or minimum) border, where $V(p) = \delta$, has been reached.

There are some subtle but crucial differences between these two approaches to the stepping. First, for the core predictions approach, the stepping is done with the Lagrange multiplier w , which should lead to corresponding steps in the property $z(p)$. In contrast, in PPL, the stepping is done with z^* directly, and steps in their Lagrange multiplier η are only taken to make $z(p) - z^*$ closer to zero. On the plus side for the core prediction approach, this difference means that the core prediction approach only needs to perform one constrained optimization problems, whereas the PPL approach solves many such problems. Since constrained optimizations are highly difficult, this may be a big advantage, and something that opens up for usage of more advanced global methods for the single constrained optimization problem (17.6). Nevertheless, the difficulty of the many constrained optimizations in PPL is relieved by the fact that they are similar, and that the end result of the previous optimization, for a slightly different value of z^* , probably is a good start guess to the new optimization, with the new value of z^* . However, in the experience of the author, the most important difference between the two approaches to stepping is another: that the core prediction approach seems to be more inviting to jumps between different qualitative behaviours, which means that incremental changes in w not always lead to incremental changes in $z(p)$. For this reason, if one wants highly resolved profiles, the PPL approach to stepping is probably to be preferred compared to the stepping done in the simple algorithm solving Eq. (17.6). Note, however, that these methods are rapidly evolving.

17.4.3 The Cut-Off Level

The cut-off level is another issue where there are similarities and differences between the two approaches. In short, using the core prediction approach, one can use any cut-off, which opens up for more pragmatic choices, but in PPL they use the same theory as for the parametric PPL, which provides for a stronger theoretical underpinning, but also makes the options fewer.

In other words, for the core prediction approach, one can use any choice of cost function, and any reasoning behind the choice of δ . Nevertheless, the standard choice for $V(p)$ is the traditional least square

$$V(p) = \sum_{t=0}^N \sum_i \frac{(y_i(t) - \widehat{y}_i(t, p))^2}{\sigma_i(t)} + \text{additional terms} \quad (17.9)$$

where i sums over the different measurement signals, σ is the standard deviation of the measurement noise, and the additional terms are optional. If the additional terms are included, one no longer has the traditional least square, and its theoretically sound properties. Nevertheless, often such additional terms are anyway warranted

and useful to include, e.g. to aid the search, or to require subjectively important qualitative behaviours. The most common way to choose a cut-off between acceptable and rejected costs is probably to use the above $V(p)$ without additional terms, and then use the inverse of the cumulative chi-square distribution, where the degrees of freedom equals the number of data points. This is often referred to as the chi-square test. More advanced approaches compensate for the number of identifiable parameters, use empirical distributions from bootstrapping, or other principles altogether, such as model discrimination or residual correlation-based tests. These cut-off methods, and more, are reviewed in the same paper that published the first version of the core prediction approach, [7].

For the PPL cut-off, the value is closely associated to the standard PLH approach. This means two things. First, the cost function must be equal to the likelihood function, which is maximized as described in earlier chapters. This maximum likelihood approach is the same as minimizing the chi-square cost above (Eq. (17.9) without the additional terms), under the assumption of additive and normally distributed measurement noise. Second, the cut-off is chosen as a certain addition, δ_{PPL} to the cost at the optimal parameters, $V(\hat{p})$. This addition, δ_{PPL} , is equal to the inverse of the cumulative chi-square distribution for 1 degree of freedom.

17.4.4 Validation of the Obtained Boundaries

Both of the new approaches have a variation step, or a voluntary final step, which is called something with the word validation in it, but again the two versions are only related and not exactly the same.

For the core prediction approach, the validation step is simply a check that the prior steps have not led to a misconception, and that the corresponding rejection really holds true. In other words, Steps 1 and 2 (Fig. 17.2) have led to a prediction with outer boundaries, and data values outside of these boundaries will probably lead to a rejection. However, even if the experimentally measured data point would lie outside of the obtained uncertainty, that is not a guarantee that the combined dataset, including the original data and the new data point, would lead to a rejection of the model, since that depends on the uncertainty of the new data point, and on how far away from the rejection boundary the model was based on the original data set. In other words, the question of whether the predicted rejection leads to an ultimate rejection needs to be tested, or validated. This validation step, Step 3, thus validates the rejection. This step was introduced already in [4], which also introduced the concept of core predictions, and then the step followed directly on Step 1, which made the validation even more important. Note that also here, the validation of a predicted rejection is made using standard rejection methods, such as the chi-square test, and that no new theory is needed.

For PPL, the validation step is more a variation of the original PPL. In this variation, the added data point z^* has an uncertainty, and is thus added on an equal footing to the other data points. In other words, no constrained optimization is needed to

ensure that $z(p) = z^*$. One instead simply optimizes over the normal parameters, obtains a cost, modifies z^* , does a new optimization, and repeats until an unacceptable agreement with the data is obtained. Note that this method of course has the limitation that it depends on what the uncertainty of z^* is assumed to be, and that the assumed uncertainty may differ from the actual uncertainty of a subsequently collected actual data point.

17.4.5 Dependencies and the Handling of Predictions of Time-Series

A final important difference between the two approaches is that already hinted at: how they relate to predictions of entire time-series.

In the core prediction approach, one can in principle expand the present framework to account for timeseries using only a minor modification, but one need to be careful about the interpretation of the results. The key modification that needs to be done is to introduce a distance measure between two time-series, so that the time-series can be projected down to a scalar, which can be added as a constraint. A standard way to introduce such a distance measure is to take the difference between the two time-series at each time point, and integrate over time. In other words, for two time-series denoted d_1 and d_2 their distance, denoted $D(d_1, d_2)$, is given by

$$D(d_1, d_2) := \int d_1(t) - d_2(t) dt \quad (17.10)$$

With these definitions in place, consider a specific predicted time-series $z = d(\hat{p})$. The max and min values of the uncertainty region around this time series is given by the same algorithm as before, where $C(p)$ is a growing function of $D(d(p), z)$ for finding Δ_z^{\min} , and a diminishing function for finding Δ_z^{\max} . Note that this is a generalization of the original algorithm: if the time-series is collapsed to a single time-point, the time-series approach finds the same boundaries as the original algorithm. Note also that this time-series formulation was the first way this core prediction approach was introduced, and that [6] therefore only presents an implementation and testing of that approach, using a particular solution to the general constrained optimization problem. We will now turn to the PPL viewpoint of this generalization, and then discuss the differences in interpretations and possibilities.

When considering time-series in the PPL framework, statistics enters the picture in another way, which makes things more complicated. In PPL, the statistical measure enters in the calculation of the term δ_{PPL} , which for a single additional data point z^* is calculated as the inverse of the cumulative chi-square distribution, with one degree of freedom, and at the desired significance level. In other words, if the significance level is 0.05, the cut-off is $\delta = V(\hat{p}) + \sim 3.8$, and the prediction should lie within the specified boundaries in 95% of the cases. In [15], they tested this claim, by generating many artificial datasets, with different noise realizations, and saw that at

least for the therein tested example, the fraction of successful prediction uncertainties converges to a reasonable vicinity of the fraction predicted by the significance level. This theory comes from the standard PLH approach, and depends on the assumption that the individual parameters, p_i , are uncorrelated, and that the new additional data point is considered as such an individual parameter. This uncorrelation assumption would have problems if one would switch from a single data point to an entire time-series, comprising many or an infinite number of data points. This is the case because if a datapoint at a specific time-point is predicted to be high, neighbouring datapoints will also be predicted to be high. In other words, for time-series, the statistical basis for PPL breaks down, and they can no longer predict the cut-off using that theory. That problem does not disappear in the validation version of PPL.

Let us now compare to the statistical viewpoint of the core prediction approach. In that approach, the cut-off is decided *a priori*, using any statistical measure, such as the chi-square value. Then, with this specification of \mathcal{A} in place, the remaining exercise is to find the extreme predictions, the max and the min, that lie within this space, and this is purely an optimization problem. In this viewpoint, there is no difference between looking at the most extreme timeseries or the most extreme time-point, because no assumptions of uncorrelation appears, as it does in the PPL approach. However, one needs to keep that in mind when interpreting a predicted time-series: it is the extreme time-series considered as a whole that is predicted. In other words, other acceptable time-series may contain values at specific time-points that are more extreme (higher or lower) than the most extreme time-series. This problem is to some extent illustrated in Fig. 17.1c. Finally, note that if one wants to use the prediction for rejections based on new data-points at various points in this predicted time-series, the validation step in the core prediction approach ensures that the rejection holds statistically: Step 3 is a standard rejection formulation, even though more than one data point has been added. Therefore, no rejections will erroneously be done because of this described problem with correlations.

17.5 Related Approaches

17.5.1 Bayesian Approaches

The most important and state-of-the-art alternative to those presented herein are probably those developed in a Bayesian setting (reviewed in e.g. [14]). These theories are in practice implemented using Markov Chain Monte Carlo (MCMC) simulations [3, 12, 26, 29]. In this setting, one does not exclusively characterise \mathcal{A} and ignores the rest of the parameter space, but one instead gives all parameter points a probability $P(p|Z)$, which says how likely the parameter points are, given the available data. These probabilities can then be used to calculate the corresponding probability of general model properties $P(z|Z)$, which can then be combined with a cut-off to obtain an uncertainty of the prediction, with a corresponding significance. However, despite

the conceptual appeal of this approach, and the existence of rather general methods, there are important limitations. For instance, the method requires the specification of a prior distribution of the parameter values, even if one has no prior knowledge on these distributions; the choice of this prior will determine the outcome. The shortcoming of this assumption in terms of achieving conclusive statements is discussed in the next section. The most important limitation, however, is that these methods only converge if the system is identifiable, or only mildly unidentifiable; such limitations do not apply to the herein presented approach. In cases of non-identifiability it appears that an important difference between the two approaches is revealed. In frequentist approaches such as those presented herein one considers the value of the cost function (or likelihood), whereas in MCMC approaches one looks at the density, and in cases of unidentifiability these do not necessarily coincide. To sort out what is the correct approach, likelihood or density, is an important task for future research comparing and combining frequentist and Bayesian approaches. Finally, just as Step 1 in the frequentist approaches, MCMC approaches suffer the problem of the curse of dimensionality, since they try to approximate a multi-dimensional distribution using points. This problem is circumvented using PPL or Step 2 in the core prediction approach. In fact, Step 2 is per definition better than MCMC at finding extreme points, since you can always use the outcome of MCMC as a start guess for the constrained optimization, which per definition will find something equally good or better.

17.5.2 Other Related Approaches and the Relation to Other Fields

A relatively extensive overview of related methods—such as interval analysis based methods [13, 28], sloppy modelling [5, 19], etc.—is available in [6]. Nevertheless, there are some important subsequent developments, and developments in other fields, which are not mentioned therein. The most important such method is PPL, and it has been described in detail above. One important neighboring research field is known as pharmacokinetics. This is a field that has done modelling of biological systems since the 70s, and they have a relatively well-developed methodological toolbox. For instance, as mentioned above, the method known as profile likelihood (for parameter uncertainty), has been used in the pharmacokinetics community for several decades, but has only recently been discovered in the systems biology community. Similarly, the problem of unidentifiability, and the need to consider multiple parameter sets, has been described also within the pharmacokinetics community, and methods have been proposed to characterize such sets. One such method is cluster Newton [1], where a cluster of parameter sets are considered together, used to approximate a surface, and where this surface then is used in the optimization. However, although cluster Newton has shown some strengths, e.g. regarding speed compared to certain other approaches, cluster Newton should primarily be considered as a new alternative to

Step 1 (Fig. 17.2), and there does not yet seem to exist a correspondence to Step 2 in the pharmacokinetics community. Another important neighboring community is known as chemometrics. Also they have now acknowledged the need to identify the well-determined predictions, and to account for the effect of unidentifiability. In the chemometrics community, the space of acceptable parameters is referred to as “neutral parameters”, and there are recent papers (e.g. [27]), where they analyse this space to find the well-determined properties (the core predictions). A recent and related method to identify the well-determined properties has been published also in the systems biology community [9]. This method starts by the identification of subsets of parameters for which the so-called Hessian matrix has full rank, except for one parameter. Within this subset, one can then numerically determine the interrelations among the parameters, from ordinary profile-likelihood plots. Finally, fits of simple analytical expressions to these numerically determined interrelations can semi-automatically give the expressions for the core predictions.

In summary, the field of handling unidentifiability is rapidly gaining in attention, not only in the field of systems biology but in several other related fields. However, there does not exist a correspondence to Step 2 in any other field; their methods are still only based on various ways of approximating the space of acceptable parameters (Step 1).

17.6 Summary and Discussion

In this chapter we have considered the important topic of obtaining predictions with uncertainty. The focus has been on frequentist approaches, which considers the value of the likelihood function, or some other cost function, as the prime decider of whether a parameter is acceptable or not. This decision is determined by a cut-off value, δ , and with this cut-off decided, the space of acceptable parameters \mathcal{A} is defined. In this chapter, two recently presented approaches to analyse this space have been reviewed and contrasted: the core prediction approach, and PPL. In Step 1 of the core prediction approach, \mathcal{A} is approximated through a point cloud, and the extreme values of the prediction z are approximated by the extreme values in this point-cloud. In Step 2, these approximate boundaries are improved upon for a specific prediction, by formulating the constrained optimization problem (17.6), which can be solved directly using an optimization algorithm that can handle nonlinear constraints, or by using the simple algorithm mentioned herein. As is shown in [6], Step 2 clearly improves upon the boundaries for a complicated model for insulin signalling. In the core-prediction approach, there is a final and optional Step 3, which validates that the potential rejection suggested by a new experimental data point really holds. In PPL, the new data point is added as a parameter, whereafter its uncertainty is characterised using a normal PLH. If the added data point contains an uncertainty, the PPL is referred to as validation PPL. The stepping and cut-off in PPL are different, and these differences have been explained. For instance, it seems like the stepping strategy described for PPL is more stable. Also, because of the differences in where

and how the statistical considerations enters the picture, the extension to time-series is handled differently in the two approaches. In the core prediction approach, Step 2 is merely an optimization exercise, and time-series can just as easily be optimized for, as individual data points. However, even though potential rejections are checked in Step 3 one should be careful with the interpretation, as it is only the time-series as a whole that is extreme, not necessarily individual time-points. In contrast, for PPL, time-series cannot be considered, since the theory behind the cut-off is based on the independence of the parameters. Finally, even though this chapter provides a merging of these two approaches, it will be an important challenge for future research to also merge and relate these methods with the three other related alternatives: cluster Newton, Bayesian, and neutral parameters.

17.6.1 What Is the Price of Considering Predictions Without Uncertainty?

To end this chapter, we want to remind the reader of the conceptual and epistemological breakthrough that it means to now have predictions with uncertainty, and of the high price that is paid by not considering this uncertainty. For many years, systems biology modelling has been done with guessed parameter values, or using parameters from the literature that are not appropriate, or based on models based on earlier guesses. If the parameters are wrong, and a single parameter point is considered, predictions from such a model is of limited, if any, interest. However, even if the parameter point could be right, i.e. it is not unrealistic, it is still uncertain, and only a single point in the space of acceptable parameter points. Since the size and uncertainty of \mathcal{A} typically is big in biology, this means that the resulting prediction uncertainty (Δ_z^{\min} , Δ_z^{\max}) also could be big, and probably would be, at least for some predictions. If one does not know which predictions are well-determined, and which are not, one has to assume that all predictions are highly undetermined. Such undetermined predictions can only make statements of the character “it could be like this, or it could be in some other way”. Such statements are in [6] defined as suggestions, or beach statements, and they are weak statements. However, if one knows that a certain prediction must lie within certain narrow boundaries, that prediction is of the same epistemological level as a rejection, since the prediction is an implicit rejection: if the prediction is not sufficiently fulfilled when tested experimentally, the model will be rejected. Since a rejection is the strongest epistemological statement available in science, a core prediction is a statement that is final. This finality holds since the outer boundaries of a core prediction will not be refuted by the collection of more data as long as that data do not show that previous data was erroneous. Core predictions are therefore well-determined, and of the strongest possible epistemological character [6]. Finally, even if one knows that a prediction is highly uncertain, i.e. that each possible behaviour within that uncertainty just is a suggestion, that is still important knowledge. One reason for that is that such highly uncertain predictions are those

that most likely will be most beneficial to measure, if one wants to further determine the parameters in the model. In that way, predictions with uncertainty, but not predictions without uncertainty, feed back to the user in a powerful way, which allows you to know both how the model works, what you know and what you guess, and which allows you to close the experiment/modelling cycle in the most powerful way. When you do a new experiment, you can know exactly why what the new data will provide: a test of the entire model structure (Fig. 17.1b), a characterization of undetermined aspects of the model (Fig. 17.1c), or a discrimination between two competing model structures (Fig. 17.1d).

Conflict of Interest

The author declares that he has no conflict of interest.

References

1. Aoki, Y., Hayami, K., de Sterck, H., Konagaya, A.: Cluster Newton method for sampling multiple solutions of underdetermined inverse problems: applications to a parameter identification problem in pharmacokinetics. *SIAM J. Sci. Comput.* **36**(1), B14–B44 (2013)
2. Bjørnstad, J.F.: (1990) Predictive likelihood: a review. *Stat. Sci.* **5**, 242–265
3. Box, G., Tiao, G.: (1973) Bayesian inference in statistical analysis, Wiley Online Library
4. Brännmark, C., Palmér, R., Glad, T., Cedersund, G., Strålfors, P.: Mass and information feed-backs through receptor endocytosis govern insulin signaling as revealed using a parameter-free modeling framework. *J. Biol. Chem.* **94**, 121–163 (2010)
5. Brown, K.S., Sethna, J.P.: Statistical mechanical approaches to models with many poorly known parameters. *Phys. Rev. E Stat. Nonlinear Soft Matter Phys.* **68**, 021904 (2003)
6. Cedersund, G.: Conclusions via unique predictions obtained despite unidentifiability—new definitions and a general method. *FEBS J.* **279**, 3513–3527 (2012)
7. Cedersund, G., Roll, J.: Systems biology: model based evaluation and comparison of potential explanations for given biological data. *FEBS J.* **276**, 22–903 (2009)
8. Cedersund, G., Samuelsson, O., Ball, G., Tegnér, J., Gomez-Cabrero, D.: Optimization in biology parameter estimation and the associated optimization problem. In: Geris, L., Gomez-Cabrero, D. (eds.) *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
9. Eisenberg, M.C., Hayashi, M.A.L.: Determining identifiable parameter combinations using subset profiling. *Math. Biosci.* **256**, 116–126 (2014)
10. Fisher, R.A.: *Statistical Methods and Scientific Inference*. Oliver and Boyd, London (1956)
11. Geris, L., Gomez-Cabrero, D.: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
12. Geyer, C.: (1992) Practical Markov Chain Monte Carlo. *Stat. Sci.* 473–483
13. Jaulin, I., Kieffer, M., Didrit, O., Walter, E.: *Applied interval analysis: with examples in parameter and state estimation, robust control and robotics*. Springer, Heidelberg (2001)
14. Kirk, P., Silk, D., Stumpf, M.P.H.: Reverse Engineering under uncertainty. In: Geris, L., Gomez-Cabrero, D. (eds.) *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)

15. Kreutz, C., Raue, A., Timmer, J.: (2012) Likelihood based observability analysis and confidence intervals for predictions of dynamic models. <http://arxiv.org/abs/1107.0013>
16. Kreutz, C., Raue, A., Timmer, J.: Likelihood based observability analysis and confidence intervals for predictions of dynamic models. *BMC Syst. Biol.* **6**, 120 (2012)
17. Ljung, L.: *System Identification—Theory for the User*, 2nd edn. PTR Prentice Hall, Upper Saddle River (1999)
18. Ljung, L., Glad, T.: On global identifiability of arbitrary model parameterization. *Automatica* **30**, 265–237 (1994)
19. Mannakee, BK., Ragsdale, AP., Transtrum, M., Gutenkunst, RN.: Sloppiness and the geometry of parameter space. In: Geris, L., Gomez-Cabrero, D. (eds.) *Uncertainty in Biology*, Springer, Heidelberg (2015)
20. Mathiasen, P.E.: Prediction functions. *Scand. J. Stat.* **6**, 1–21 (1979)
21. Meeker, W., Escobar, L.: Teaching about approximate confidence regions based on maximum likelihood estimation. *Am. Stat.* **49**, 48–53 (1995)
22. Nyman, E., Brännmark, C., Palmér, R., Brugård, J., Nyström, F.H., Strålfors, P., Cedersund, G.: A hierarchical whole-body modeling approach elucidates the link between in Vitro insulin signaling and in Vivo glucose homeostasis. *J. Biol. Chem.* **286**, 26028–26041 (2011)
23. Raue, A., Kreutz, C., Maiwald, T., Bachmann, J., Schilling, M., Klingmüller, U., Timmer, J.: Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics* **25**, 1923–9 (2009)
24. Sahle, S., Mendes, P., Hoops, S., Kummer, U.: A new strategy for assessing sensitivities in biochemical models. *Philos. Trans. A Math. phys. Eng. Sci.* **366**, 3619–3631 (2008)
25. Sedoglavic, A.: A probabilistic algorithm to test local algebraic observability in polynomial time. *J. Symbolic Comput.* **33**, 735–755 (2002)
26. Sunnåker, M., Stelling, J.: Model extension and model selection. In: Geris, L., Gomez-Cabrero, D. (eds.) *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
27. Tafintseva, V., Tøndel, K., Ponosov, A., Martens, H.: *J. Chemometr.* **28**, 645–655 (2014)
28. Tucker, W.: Interval methods. In: Geris, L., Gomez-Cabrero, D. (eds.) *Uncertainty in Biology*, Springer, Heidelberg (2015)
29. Vanlier, J., Tiemann, C., Hilbers, P., van Riel, N.: An integrated strategy for prediction uncertainty analysis, *Bioinformatics*, In press (2012)

Chapter 18

Computational Modeling Under Uncertainty: Challenges and Opportunities

David Gomez-Cabrero, Jesper Tegnér and Liesbet Geris

Abstract Computational Biology has increasingly become an important tool for biomedical and translational research. In particular, when generating novel hypothesis despite fundamental uncertainties in data and mechanistic understanding of biological processes underpinning diseases. While in the present book, we have reviewed the necessary background and existing novel methodologies that set the basis for dealing with uncertainty, there are still many “grey”, or less well-defined, areas of investigations offering both challenges and opportunities. This final chapter in the book provides some reflections on those areas, namely: (1) the need for novel robust mathematical and statistical methodologies to generate hypothesis under uncertainty; (2) the challenge of aligning those methodologies in a context that requires larger computational resources; (3) the accessibility of modeling tools for less mathematical literate researchers; and (4) the integration of models with—omics data and its application in clinical environments.

Keywords Computational modeling · Uncertainty · Challenges · HPC · Hypothesis generation

D. Gomez-Cabrero (✉) · J. Tegnér
Unit of Computational Medicine, Department of Medicine, Karolinska Institutet,
Solna, Sweden
e-mail: david.gomezcabrero@ki.se

D. Gomez-Cabrero · J. Tegnér
Center for Molecular Medicine, Stockholm, Sweden

L. Geris
Biomechanics Research Unit, University of Liège, Chemin des Chevreuils 1 B52/3,
4000 Liège, Belgium
e-mail: liesbet.geris@ulg.ac.be

L. Geris
Prometheus, Division of Skeletal Tissue Engineering Leuven, KU Leuven,
Herestraat 49, Box 813, 3000 Leuven, Belgium

L. Geris
Biomechanics Section, KU Leuven, Celestijnenlaan 300 C, Box 2419,
3001 Leuven, Belgium

18.1 Introduction

There are two underlying rationales that motivate the chapters in this book. The first, is the *usefulness* and *necessity* of mechanistic mathematical and computational modeling in biomedical research. The usefulness has been widely shown in several chapters (see for instance [31] by Lejon and Samaey, [26] by Hug et al.) and (for instance) from classical groundbreaking works in neuron modeling (such as [23]). The necessity of mechanistic modeling originates essentially from the limitations w.r.t mechanistic understanding when solely using classical statistical analysis in the analysis of complex systems [51].

The second rationale is that mechanistic modeling in biology needs to address uncertainty in order to generate testable hypothesis. For instance in biology when a transcript is profiled—either by PCR, array or RNA-seq—there are several sources of variability to consider: *technical*, from the experimental procedure use, and *biological*, that is for instance when the same type of cell may react in different ways to the same perturbation. At the cell level, one explanation for observed transcriptomics biological variation is that the regulation is driven at several and different layers (e.g. genetic and epigenetic regulation), but large parts of these regulatory mechanisms are still only in part possible to decipher [3, 22, 27, 32]. Furthermore, the profiles of those “other regulatory layers” are in most cases not available during modeling. A second explanation for uncertainty is the stochastic nature of some biological processes as shown in intra-cellular chemical reactions, gene expression [33] and pharmacokinetics [12] among others. In both explanations, we need to clearly face *uncertainty during the modeling, in the parameters of the model and in the biological processes* when investigating model behaviors.

While the [16] by Geris and Gomez-Cabrero provides an overview, we find it useful to close the book with a chapter that summarizes major existing challenges and opportunities. We have identified four challenges that will be briefly discussed in the different sections of this Chapter. First, (1) there is a need for methodological development, (2) linking modeling and high-performance computing, (3) strengthen the accessibility of modeling tools targeting non-specialists and, (4) integrating omics and modeling tools for the benefit of personalized medicine. Additional challenges for the future of computational biomedicine, especially with respect to the clinical dimension, can be found in the Digital Patient Roadmap.¹

18.2 The Need for Methodological Development

In the last decade, we have observed a shift in biological modeling analysis. In initial attempts, mechanistic ordinary differential equation (ODE) models were generated by defining a set of equations, and investigators *manually* fine-tuned the parameters. The manual fine-tuning was conducted by exploring the parameter space “in the

¹http://www.digital-patient.net/files/DP-Roadmap_FINAL_N.pdf.

quest” of finding those parameters that agreed with experimental observed behavior (we will denote them by “good quality parameter sets”). Eventually, the manual search was made *automatic* by designing the fine-tuning problem as an optimization problem as shown in [5] by Salmuelsen et al. Furthermore, with the growth of computational resources the parameter space of larger models became intractable using theoretical analysis, it became clear that investigations of either exploring the surrounding areas of good quality parameter sets (see [34] by Mannakee et al. and [47] by Van Schepdael et al.) or by exploring the set of “good quality parameter sets” ([17] by Gomez-Cabrero et al. and [7] by Cedersund) became important. We consider that those types of methodologies are necessary and they are an active research field in computational biology, however it still requires a coordinated effort to generate a solid foundation for further development. We consider two major requirements:

(1) **Rigorous definitions** In order to develop useful methodologies and tools we need to provide a robust answer to the following question: *what is a useful output from the analysis under uncertainty of a biological theoretical model?* (QUES). In [17] by Gomez-Cabrero et al. the answer proposed is (briefly) *first the grouping and secondly group characterization of good quality parameter sets*. The idea is that by exploring the “good quality parameter set” space it is possible to find competing hypothesis (from the groups of “good quality parameter sets”) that could be tested at the laboratory. However, given the exploratory nature of the proposed methodology (that does not investigate all possible “good quality parameter sets” but a sample of them by an optimization methodology) the robustness of the competing hypothesis is not rigorously ensured. Reference [7] by Cedersund answers that the fundamental outputs are the set of predictions that can be then tested back in the laboratory. Furthermore [7] by Cedersund provides an initial classification of predictions: *core predictions* (well-determined predictions that allow to test the quality of the model) and *suggestions* (poorly determined predictions that may provide specific insights that can be tested in order to improve the overall quality model). Both results and proposals shown in [17, 34] by Gomez-Cabrero et al. and [7] by Cedersund represent part of the initial efforts generated to provide a formal answer to QUES; however we consider it necessary to develop further these efforts and work on generating a consensus and robust formulation for answering QUES. Relevant material on the topic can be found in [6, 9, 10, 18, 30, 46].

(2) **Development of software tools that implement such methodologies so they may become a standard** The shift from manual search to automatic search started during last decades of 20th century and actively continued during first decade of 21st century. Several teams worked on those ideas and several tools were developed at the same time; some of those tools aimed for specific areas such as Neuroscience (Neurofitter, [45]) while some other tools were more generic such as COPASI [24]. Many of those tools are still available (and there are active research groups continuously updating them) see [5] by Cedersund et al. for further detail. On one hand, the generation of that many tools raised the awareness and use of those new methodologies; on the other hand it was clear that the wheel was reinvented many times. When considering the generation of hypothesis under uncertainty we may argue to be at the beginning of user-friendly method development. Yet no tool is able to

perform automatically the analysis presented in [17] by Gomez-Cabrero et al. or [7] by Cedersund; in those cases customized coding solutions were generated. We consider it necessary to generate user-friendly solutions able to perform automatically (or under human supervision) those analyses. However we also believe it is necessary to generate coordinated working groups to avoid the generation of similar tools simultaneously.

A final complementary development to those methodologies is the generation of novel methodologies and (user-friendly) tools allowing automatic simplification and reduction of models as shown in [14] by Eriksson et al., [44] by Tucker or through Global Sensitivity Analysis [29, 39, 43].

18.3 Integration of Computational Modeling with High-Performance Computing Techniques

Computational resources have been both the key and bottleneck for computational modeling analysis. The automatic search for “good quality parameter sets” depended on the availability of machines able to run hundreds or thousands of simulations in brief periods of time. This was possible through medium sized (20+ cores) to large sized (named supercomputers such as Mare Nostrum in the Barcelona Supercomputing Center (www.bsc.es, Spain) or SNIC solutions (www.snic.vr.se, Sweden) machines; the former was mainly affordable by computational-oriented groups able to invest funding in the resource while the latter were available through national programs that provided (and still provide) a number of hours-per-month upon request. The first computational biology analysis competed for such computational resources with theoretical physics or computational chemistry (among many others) simulations, but at that time the required resources were minor compared to the rest of research areas. Over the years, and with both the development of automatic fine-tuning tools and larger models, the computational requirements grew and computational biology is starting to compete at a similar scale of requirements than the other research domains. The present and coming future shows that the demand of computational requirements are still to grow for several reasons, among them: (i) possible increased size of the models, (ii) increased amount of data to be considered (see later the omics’ section for further details) and (iii) an increased amount of users (see for instance the development of novel conferences such as HiCOMB, High Performance Computational Biology from 2002 until nowadays). For this reason the long-term resources are to be planned carefully in order to correctly assess the future needs of Biological and Medical Sciences.² We consider the following three aspects to be of major relevance:

(1) **High-Performance Computing (HPC) infrastructures** There is a general trend to avoid buying small-medium computational resources by every group and invest better into large-scale resources or cloud-based solutions; see for instance the

²<http://cordis.europa.eu/fp7/ict/e-infrastructure/docs/bms-agenda.pdf>.

action plan for the Digital Agenda for Europe.³ Small-medium sized solutions tend to be expensive and, in many cases, sub-optimally used. While cloud-based solutions, if prizes are competitive, may provide a cheaper solution that will optimally reflect the needs and uses of different research groups in real-time. Furthermore, as pointed out by Peter V Coveney,⁴ it is necessary to optimize the interoperability across large infrastructures and it is necessary to harmonize mechanisms such as access, advance reservation and urgent computing among others.

(2) **Parallelization** Both simulation and fine-tuning benefit from parallelization, that is the possibility to run a process as separate parallel batches therefore reducing the amount of time by using several CPUs simultaneously. Both optimization algorithms and methodologies to integrate Partial Differential Equations benefit from better and robust parallelizable algorithms. Interestingly, in the area of eScience (“the application of computer technology to the undertaking of modern scientific investigation, including the preparation, experimentation, data collection, results dissemination, and long-term storage and accessibility of all materials generated through the scientific process”, Shannon Bohle⁵) there is an effort to import to computational biology those methods already developed for other areas where large-scale modeling is actively used (such as Weather Forecast modeling). Among those efforts there is the Swedish e-Science Research Center (<http://www.e-science.se>).

(3) **Scalability** Both for computational resources and parallelization need to consider optimal scalability of the solutions developed, given that the number of users and computational requirements is expected to grow over time [15, 21].

18.4 To Widen the Use and Applicability of Modeling as a Tool for Non-specialists

Most of the chapters of this book have been written by statisticians, mathematicians, and engineers with a strong mathematical background. This may represent the background requirements for method development in computational biology, however it does not represent the requirements for *using* computational biology. Fortunately, in the last twenty numerous biologists have been exposed to the necessary background to develop and analyze their own models. We consider that to make the use of modeling in biomedicine it is important to make the necessary knowledge and tools as accessible as possible; on this direction we consider that following points are important.

(1) **The necessary theoretical background** When biologists decide to design a model of their system under study, it is necessary for them to learn the basics

³<https://ec.europa.eu/digital-agenda/en/pillar-v-research-and-innovation/action-53-financially-support-joint-ict-research-infrastructures>.

⁴<http://cordis.europa.eu/fp7/ict/e-infrastructure/docs/bms-pres-a-6.pdf>.

⁵http://www.scilogsg.com/scientific_and_medical_libraries/what-is-e-science-and-how-should-it-be-managed/.

of mathematical modeling. General and specific knowledge of modeling will be required depending on the system to investigate. The amounts of material (specially books) addressing this knowledge have been growing in both quantity and user-friendliness. Additionally, courses (such as Computational Biology in Cold Spring Harbor, directed by Professor Gregory Smith⁶) are becoming more common. We consider that it is necessary to continue this trend, but also that (i) courses where biological-strong and mathematical-strong participants are both enlisted are to be prioritized, because it allows exchanging of views and goals and creates a richer learning environment [4]; and (ii) the development of on-line courses addressing this topic needs to receive attention, so students may have introductory sessions without the need to wait for face-2-face courses.

(2) **Software environments** We consider it necessary to enhance the user-friendliness of existing (and novel tools) in order to enlist researchers in the use of modeling. Existing tools have certainly shown an increase in accessibility and friendliness, but any researcher with no experience will still need to invest large amounts of time to get confident with them. In Ph.D. programs where modeling may be a side project to investigate experimental results this situation may end in not considering modeling as a research tool. We consider that our aim must be to make “computational modeling” another accessible tool in the biologist tool-box, therefore improving user-friendliness is necessary. An example of generating a simulation environment for medical researchers is [25], which is part of the results from the European Project Synergy-COPD [19].

(3) **Syllabus implementation** When a clinician or a biologist may interact with modelers or discover a model of interest, existing syllabus usually do not provide the necessary background to understand them. We consider that initiatives such as Erasmus BioHealth Computing Program [4] and Medical Research Masters are initiatives of value where future biological and medical researchers are set to interact with modelers and computational biologists. This approach enhances the visibility of modeling in biology and biomedicine.

18.5 Forming Stronger Ties Between Omics Data and Computational Biology

Following the Human Genome Project, array-based and Next-Generation Sequencing-based technologies have pushed transcriptomics analysis to novel boundaries [1, 35]. SNP profiling of thousands of individuals have allowed the identification of genetic risk factors for many diseases such as Multiple Sclerosis [41] or Rheumatoid Arthritis [40], however the use of such information in Computational Models is limited to say the least. A very important open question is then: *how do we integrate and omics-based knowledge into modeling?*

⁶<http://meetings.cshl.edu/courses/2014/c-comp14.shtml>.

While omics-data is used in the generation of predictive models (such as patient classification or risk prediction) and integrative approaches are being continuously developed to improve such models [2, 11, 28, 38, 50] what we here refer to is the use of omics data in the analysis of biological systems through mechanistic models. Eventually those integrated mechanistic models may provide in the future relevant information to be included in better prediction models making use of simulation outputs. However, at the present time we focus on the challenge of *creating models that address the individual (personalized modeling)*. Lets consider for instance the development of a immune system model of Multiple Sclerosis Progression such as the one presented in [48]. If we gather information of DNA Methylation profiling and/or SNP genotype for a given individual, the challenge is now how we implement such information so the model is not anymore a generic model but individual specific. There exist several attempts on this direction as those shown in Synergy-COPD [19] and CombiMS, in the context of Systems Medicine and the Virtual Physiological Human. In order for omics data to be routinely used in computational biomedicine and, later on, in a clinical setting, a number of requirements need to be fulfilled, as recently identified by [49]. These include (1) the ability to work with sensitive data, (2) to work with complex and heterogeneous data (including non-textual information), (3) to work with a distributed data management under security and performance constraints, (4) to define methods allowing for the integration of bioinformatics and systems biology information with clinical observations on various length scales, and finally (5) to define tools able to define the ‘physiological envelope’ of a patient (ref white paper).

18.6 Conclusions

We find that Computational Biology is a crucial tool for biology and biomedicine, but to enhance its practical applicability there is an urgent need to address the uncertainty commonly observed in biological systems to ensure the uptake in the biological and clinical communities. The present chapter reviews the needs and challenges in computational biology, that are important to consider in the nearby development of the field. We summarize those needs in four major aspects:

1. Robust definitions for the generation of useful predictions,
2. Development of novel and optimization of existing HPC resources that address the state-of-the-art computational needs.
3. Development of user-friendly analysis tools and easily accessible computing resources,
4. Development of models and tools that incorporate information on the different omics widely profiled nowadays.

We hope that the reading of this book may motivate young and senior researchers to follow and work on those challenges.

References

1. Almomani, R., van der Heijden, J., Ariyurek, Y., Lai, Y., Bakker, E., van Galen, M., den Dunnen, J.T.: Experiences with array-based sequence capture; toward clinical applications. *Eur. J. Hum. Genet. EJHG* **19**(1), 50–5 (2011). doi:[10.1038/ejhg.2010.145](https://doi.org/10.1038/ejhg.2010.145)
2. Anderson, A.R., Quaranta, V.: Integrative mathematical oncology. *Nat. Rev. Cancer* **8**(3), 227–234 (2008). doi:[10.1038/nrc2329](https://doi.org/10.1038/nrc2329)
3. Bock, C.: Analysing and interpreting DNA methylation data. *Nat. Rev. Genet.* **13**(10), 705–719 (2012). doi:[10.1038/nrg3273](https://doi.org/10.1038/nrg3273)
4. Cascante, M., de Atauri, P., Gomez-Cabrero, D., Wagner, P., Centelles, J.J., Marin, S., Sabatier, P.: Workforce preparation: the Biohealth computing model for master and PhD students. *J. Trans. Med.* **12**(2), S11 (2014). doi:[10.1186/1479-5876-12-S2-S11](https://doi.org/10.1186/1479-5876-12-S2-S11)
5. Cedersund, G., Samuelsson, O., Ball, G., Tegnér, J., Gomez-Cabrero, D.: Optimization in biology parameter estimation and the associated optimization problem. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
6. Cedersund, G.: Conclusions via unique predictions obtained despite unidentifiability—new definitions and a general method **279**, 3513–3527 (2012). doi:[10.1111/j.1742-4658.2012.08725.x](https://doi.org/10.1111/j.1742-4658.2012.08725.x)
7. Cedersund, G.: Prediction uncertainty estimation despite unidentifiability: an overview of recent developments. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
8. Cedersund, G.: Conclusions via unique predictions obtained despite unidentifiability—new definitions and a general method. *FEBS J.* **279**(18), 3513–3527 (2012). doi:[10.1111/j.1742-4658.2012.08725.x](https://doi.org/10.1111/j.1742-4658.2012.08725.x)
9. Cedersund, G., Roll, J.: Systems biology: model based evaluation and comparison of potential explanations for given biological data. *FEBS J.* **276**(4), 903–922 (2009). doi:[10.1111/j.1742-4658.2008.06845.x](https://doi.org/10.1111/j.1742-4658.2008.06845.x)
10. Cedersund, G., Strålfors, P.: Putting the pieces together in diabetes research: towards a hierarchical model of whole-body glucose homeostasis. *Eur. J. Pharm. Sci. Off. J. Eur. Fed. Pharm. Sci.* **36**(1), 91–104 (2009). doi:[10.1016/j.ejps.2008.10.027](https://doi.org/10.1016/j.ejps.2008.10.027)
11. Chan, S.Y., Loscalzo, J.: The emerging paradigm of network medicine in the study of human disease. *Circ. Res.* **111**(3), 359–374 (2012). doi:[10.1161/CIRCRESAHA.111.258541](https://doi.org/10.1161/CIRCRESAHA.111.258541)
12. Donnet, S., Samson, A.: A review on estimation of stochastic differential equations for pharmacokinetic/pharmacodynamic models. *Adv. Drug Delivery Rev.* **65**(7), 929–939 (2013). doi:[10.1016/j.addr.2013.03.005](https://doi.org/10.1016/j.addr.2013.03.005)
13. Droste, P., Miebach, S., Niedenführ, S., Wiechert, W., Nöh, K.: Biosystems visualizing multi-omics data in metabolic networks with the software Omix—a case study. *BioSyst.* **105**(2), 154–161 (2011). doi:[10.1016/j.biosystems.2011.04.003](https://doi.org/10.1016/j.biosystems.2011.04.003)
14. Eriksson, O., Tegnér, J.: Modeling and model simplification to facilitate biological insights and predictions. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
15. Finak, G., Frelinger, J., Jiang, W., Newell, E.W., Ramey, J., Davis, M.M., Gottardo, R.: OpenCyto: an open source infrastructure for scalable, robust, reproducible, and automated, end-to-end flow cytometry data analysis. *PLoS Comput. Biol.* **10**(8), e1003806 (2014). doi:[10.1371/journal.pcbi.1003806](https://doi.org/10.1371/journal.pcbi.1003806)
16. Geris, L., Gomez-Cabrero, D.: An introduction to uncertainty in the development of computational models of biological processes. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
17. Gomez-Cabrero, D., Ardid, S., Cano-Colino, M., Tegnér, J., Compte, A.: Neuroswarm: a methodology to explore the constraints that function imposes on simulation parameters in large-scale networks of biological neurons. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)

18. Gomez-Cabrero, D., Compte, A., Tegner, J.: Workflow for generating competing hypothesis from models with parameter uncertainty. *Interface Focus* **1**(3), 438–449 (2011). doi:[10.1098/rsfs.2011.0015](https://doi.org/10.1098/rsfs.2011.0015)
19. Gomez-Cabrero, D., Menche, J., Cano, I., Abugessaisa, I., Huertas-Migueláñez, M., Tenyi, A., Tegnér, J.: Systems Medicine: from molecular features and models to the clinic in COPD. *J. Trans. Med.* **12**(2), S4 (2014). doi:[10.1186/1479-5876-12-S2-S4](https://doi.org/10.1186/1479-5876-12-S2-S4)
20. Gomez-Ramirez, J., Sanz, R.: On the limitations of standard statistical modeling in biological systems: a full Bayesian approach for biology. *Prog. Biophys. Mol. Biol.* **113**(1), 80–91 (2013). doi:[10.1016/j.pbiomolbio.2013.03.008](https://doi.org/10.1016/j.pbiomolbio.2013.03.008)
21. Gupta, A., Briat, C., Khammash, M.: A scalable computational framework for establishing long-term behavior of stochastic reaction networks. *PLoS Comput. Biol.* **10**(6), e1003669 (2014). doi:[10.1371/journal.pcbi.1003669](https://doi.org/10.1371/journal.pcbi.1003669)
22. Heyn, H., Esteller, M.: DNA methylation profiling in the clinic: applications and challenges. *Nat. Rev. Gen.* **13**(10), 679–692 (2012). doi:[10.1038/nrg3270](https://doi.org/10.1038/nrg3270)
23. Hodgkin, A.L., Huxley, A.F.: Currents carried by sodium and potassium ions through the membrane of the giant axon of *Loligo*. This information is current as of January 29, This is the final published version of this article?; it is available at?: This version of the article may not be. *J. physiol. Paris* **116**, 449–472 (1952)
24. Hoops, S., Sahle, S., Gauges, R., Lee, C., Pahle, J., Simus, N., Singhal, M., Xu, L., Mendes, P., Kummer, U.: COPASI—a COMplex PATHway SIMulator. *Bioinformatics* **22**(24), 3067–3074 (2006)
25. Huertas-Migueláñez, M., Mora, D., Cano, I., Maier, D., Gomez-Cabrero, D., Lluch-Ariet, M., Miralles, F.: Simulation environment and graphical visualization environment: a COPD use-case. *J. Trans. Med.* **12**(2), S7 (2014). doi:[10.1186/1479-5876-12-S2-S7](https://doi.org/10.1186/1479-5876-12-S2-S7)
26. Hug, S., Schmidl, D., Bo Li, W., Greiter, M.B., Theis, F.J.: Bayesian model selection methods and their application to biological ODE systems. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
27. Jones, P.A., Liang, G.: Rethinking how DNA methylation patterns are maintained. *Nat. Rev. Gen.* **10**(11), 805–811 (2009). doi:[10.1038/nrg2651](https://doi.org/10.1038/nrg2651)
28. Joyce, A.R., Palsson, B.Ø.: The model organism as a system: integrating “omics” data sets, 7(March), 198–210. doi:[10.1038/nrm1857](https://doi.org/10.1038/nrm1857)
29. Kent, E., Neumann, S., Kummer, U., Mendes, P.: What can we learn from global sensitivity analysis of biochemical systems? *PLoS One* **8**(11), e79244 (2013). doi:[10.1371/journal.pone.0079244](https://doi.org/10.1371/journal.pone.0079244)
30. Kuepfer, L., Peter, M., Sauer, U., Stelling, J.: Ensemble modeling for analysis of cell signaling dynamics. *Nat. Biotechnol.* **25**(9), 1001–1006 (2007). doi:[10.1038/nbt1330](https://doi.org/10.1038/nbt1330)
31. Lejon, A., Samaey, G.: Stochastic modeling and simulation methods for biological processes: overview. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
32. Luco, R.F., Allo, M., Schor, I.E., Kornblihtt, A.R., Misteli, T.: Review Epigenetics in alternative pre-mRNA splicing. *Cell* **144**(1), 16–26 (2010). doi:[10.1016/j.cell.2010.11.056](https://doi.org/10.1016/j.cell.2010.11.056)
33. Magklara, A., Lomvardas, S.: Stochastic gene expression in mammals: lessons from olfaction. *Trends Cell Biol.* **23**(9), 449–456 (2013). doi:[10.1016/j.tcb.2013.04.005](https://doi.org/10.1016/j.tcb.2013.04.005)
34. Mannakee, B.K., Ragsdale, A.P., Transtrum, M.K., Gutenkunst, R.N.: Sloppiness and the geometry of parameter space. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
35. Metzker, M.L.: Sequencing technologies—the next generation. *Nat. Rev. Gen.* **11**(1), 31–46 (2010). doi:[10.1038/nrg2626](https://doi.org/10.1038/nrg2626)
36. Miyoshi, N.S.B., Pinheiro, D.G., Silva, W.A., Felipe, J.C.: Computational framework to support integration of biomolecular and clinical data within a translational approach. *BMC Bioinform.* **14**, 180 (2013). doi:[10.1186/1471-2105-14-180](https://doi.org/10.1186/1471-2105-14-180)
37. Petersson, K.M., Nichols, T.E., Poline, J.B., Holmes, A.P.: Statistical limitations in functional neuroimaging. I. Non-inferential methods and statistical models. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **354**(1387), 1239–1260 (1999). doi:[10.1098/rstb.1999.0477](https://doi.org/10.1098/rstb.1999.0477)

38. Ramsey, S.A., Gold, E.S., Aderem, A.: A systems biology approach to understanding atherosclerosis **2**(3), 79–89 (2010). doi:[10.1002/emmm.201000063](https://doi.org/10.1002/emmm.201000063)**A**
39. Rand, D.A.: Mapping global sensitivity of cellular network dynamics: sensitivity heat maps and a global summation law. *J. Roy. Soc. Interface/Roy. Soc.* **5**(1)(00), S59–69 (2008). doi:[10.1098/rsif.2008.0084.focus](https://doi.org/10.1098/rsif.2008.0084.focus)
40. Raychaudhuri, S., Sandor, C., Stahl, E.A., Freudenberg, J., Lee, H.-S., Jia, X., de Bakker, P.I.W.: Five amino acids in three HLA proteins explain most of the association between MHC and seropositive rheumatoid arthritis. *Nat. Genet.* **44**(3), 291–296 (2012). doi:[10.1038/ng.1076](https://doi.org/10.1038/ng.1076)
41. Sawcer, S., Hellenthal, G., Pirinen, M., Spencer, C.C.A., Patsopoulos, N.A., Moutsianas, L., Compston, A.: Genetic risk and a primary role for cell-mediated immune mechanisms in multiple sclerosis. *Nature* **476**(7359), 214–219 (2011). doi:[10.1038/nature10251](https://doi.org/10.1038/nature10251)
42. Seoane, J.A., Aguiar-Pulido, V., Munteanu, C.R., Rivero, D., Rabunal, J.R., Dorado, J., Pazos, A.: Biomedical data integration in computational drug design and bioinformatics. *Curr. Comput. Aided Drug Des.* **9**(1), 108–117 (2013). <http://www.ncbi.nlm.nih.gov/pubmed/23294434>
43. Sumner, T., Shephard, E., Bogle, I.D.L.: A methodology for global-sensitivity analysis of time-dependent outputs in systems biology modelling. *J. R. Soc. Interface* 2156–2166 (2012)
44. Tucker, W.: Interval methods. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
45. Van Geit, W., Achard, P., De Schutter, E.: Neurofitter: a parameter tuning package for a wide range of electrophysiological neuron models. *Front. Neuroinform.* **1**(November), 1 (2007). doi:[10.3389/neuro.11.001.2007](https://doi.org/10.3389/neuro.11.001.2007)
46. Van Riel, N.A.W.: Dynamic modelling and analysis of biochemical networks: mechanism-based models and model-based experiments. *Brief. Bioinform.* **7**(4), 364–74 (2006). doi:[10.1093/bib/bbl040](https://doi.org/10.1093/bib/bbl040)
47. Van Schepdael, A., Carlier, A., Geris, L.: Sensitivity analysis by design of experiments. In: *Uncertainty in Biology, A Computational Modeling Approach*. Springer, Chem (2016, this volume)
48. Velez de Mendizabal, N., Carneiro, J., Sole, R.V., Goni, J., Bragard, J., Martinez-Forero, I., Villoslada, P.: Modeling the effector - regulatory T cell cross-regulation reveals the intrinsic character of relapses in Multiple Sclerosis. *BMC Systems Biology* **5**(1), 114 (2011). doi:[10.1186/1752-0509-5-114](https://doi.org/10.1186/1752-0509-5-114)
49. Viceconti, M., Hunter, P., McCormack, K., Henney, A., Omholt, S.W., Graf, N., Morley-Fletcher, E., Geris, L., Hose, R.: Big data, big knowledge: big data for personalised healthcare, White Paper from the VPH Institute (2014). http://www.vph-institute.org/upload/white-paper-big-data-and-vph-v7mc_54662d0e8ff52.pdf
50. Villoslada, P., Baranzini, S.: Data integration and systems biology approaches for biomarker discovery: Challenges and opportunities for multiple sclerosis. *J. Neuroimmunol.* **1–8**, (2012). doi:[10.1016/j.jneuroim.2012.01.001](https://doi.org/10.1016/j.jneuroim.2012.01.001)
51. Zhu, J., Zhang, B., Smith, E.N., Drees, B., Brem, R.B., Bumgarner, R.E., Schadt, E.E.: Complex. *Yeast Regul. Netw.* **40**(7), 854–861 (2009). doi:[10.1038/ng.167](https://doi.org/10.1038/ng.167)**Integrating**

Author Index

A

Ardid, Salva, [427](#)

B

Ball, Gordon, [33](#), [177](#)

Bastogne, Thierry, [127](#)

Bullinger, Eric, [127](#)

C

Cano-Colino, Maria, [427](#)

Carlier, Aurélie, [327](#)

Cedersund, Gunnar, [177](#), [449](#)

Chitforoushzadeh, Zeinab, [155](#)

Compte, Albert, [427](#)

D

D'Otreppe, Vinciane, [393](#)

E

Eriksson, Olivia, [301](#)

F

Fey, Dirk, [127](#)

Findeisen, Rolf, [127](#)

G

Geris, Liesbet, [3](#), [327](#), [467](#)

Gomez-Cabrero, David, [3](#), [177](#), [427](#), [467](#)

Greiter, Matthias B., [243](#)

Gutenkunst, Ryan N., [271](#)

H

Hug, Sabine, [243](#)

J

Janes, Kevin A., [155](#)

Jones, Alison Claire, [393](#)

K

Kirk, Paul, [15](#)

L

Lagani, Vincenzo, [33](#)

Laing, Carlo R., [367](#)

Lejon, Annelies, [75](#)

Li, Wei Bo, [243](#)

M

Mannakee, Brian K., [271](#)

Mengoni, Marlène, [393](#)

R

Ragsdale, Aaron P., [271](#)

S

Samaey, Giovanni, [75](#)

Samuelsson, Oscar, [33](#), [177](#)

Scheurich, Peter, [127](#)
Schliemann-Bullinger, Monica, [127](#)
Schmidl, Daniel, [243](#)
Shah, Millie, [155](#)
Sikora, Sebastien, [393](#)
Silk, Daniel, [15](#)
Stelling, Joerg, [213](#)
Stumpf, Michael P.H., [15](#)
Sunnåker, Mikael, [213](#)

T

Tegnér, Jesper, [33](#), [177](#), [301](#), [427](#), [467](#)

Theis, Fabian J., [243](#)
Transtrum, Mark K., [271](#)
Triantafillou, Sofia, [33](#)
Tsamardinos, Ioannis, [33](#)
Tucker, Warwick, [199](#)

V

Van Schepdael, An, [327](#)

W

Wilcox, Ruth Karen, [393](#)