

Nonlinear Dynamics and Statistics

Alistair I. Mees
Editor

Nonlinear Dynamics and Statistics

With 142 Figures



Springer Science+Business Media, LLC

Alistair I. Mees
Centre for Applied Dynamics and Optimization
University of Western Australia
Perth, Western Australia WA 6907
Australia

Library of Congress Cataloging-in-Publication Data

Nonlinear dynamics and statistics / [edited by] Alistair I. Mees.

p. cm.

Includes bibliographical references.

ISBN 978-1-4612-6648-8 ISBN 978-1-4612-0177-9 (eBook)

DOI 10.1007/978-1-4612-0177-9

1. Dynamics. 2. Time-series analysis. 3. Nonlinear theories. I. Mees, A.I.

QA845 .N648 2000

531'.11—dc21

00-039807

CIP

Printed on acid-free paper.

© 2001 Springer Science+Business Media New York

Originally published by Birkhäuser Boston in 2001

Softcover reprint of the hardcover 1st edition 2001



All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher Springer Science+Business Media, LLC, except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use of general descriptive names, trade names, trademarks, etc., in this publication, even if the former are not especially identified, is not to be taken as a sign that such names, as understood by the Trade Marks and Merchandise Marks Act, may accordingly be used freely by anyone.

ISBN 978-1-4612-6648-8

SPIN 10749185

Production managed by Louise Farkas; manufacturing supervised by Erica Bresler.

Typeset by the editor in TeX.

9 8 7 6 5 4 3 2 1

Contents

Preface	xiii
Contributors	xix
I Issues in Reconstructing Dynamics	1
1 Challenges in Modeling Nonlinear Systems: A Worked Example	
<i>Henry D.I. Abarbanel</i>	3
1.1 Summary and Preview	4
1.2 Small Neural Assemblies: What We Did	5
1.3 General Outlook	7
1.4 The Neuron Model	9
1.4.1 Ca^{2+} Dynamics	10
1.4.2 Voltage Dynamics	10
1.4.3 Parameters	12
1.4.4 Comment on This Model	12
1.5 Choosing T	12
1.6 Choosing d : Global Dimension d_E	14
1.7 Dynamical Dimension	15
1.8 Lyapunov Exponents; Predictability	21
1.9 “Black Box” Prediction in State Space	23
1.10 Summary and Challenges	26
2 Disentangling Uncertainty and Error: On the Predictability of Nonlinear Systems	
<i>Leonard A. Smith</i>	31
2.1 Introduction	31
2.2 Uncertainty	34
2.3 The Perfect Model Scenario	35
2.3.1 Forecasting with a Perfect Model	38
2.4 Ensemble Verification	44
2.4.1 Minimum Spanning Trees	45
2.4.2 Relevance to Operational Forecasting	48
2.5 Imperfect Model Scenarios	49
2.5.1 ϕ -Shadowing	52
2.5.2 Bounding Boxes	53

2.5.3	Applications to Climate	54
2.6	Multimodel CPT Ensembles	55
2.7	Discussion	57
2.8	Summary	59
3	Achieving Good Nonlinear Models: Keep It Simple, Vary the Embedding, and Get the Dynamics Right	
	<i>Kevin Judd, Michael Small and Alistair I. Mees</i>	65
3.1	Introduction	65
3.2	Minimum Description Length Models: Keep It Simple . . .	66
3.3	Variable Embedding: Cylinders Stack Up Better than Spheres	69
3.3.1	Uniform Embedding	70
3.3.2	Nonuniform Embedding	70
3.3.3	Variable Embedding	71
3.3.4	Cylindrical Basis Models	72
3.4	Systematic Errors: $\Psi\Phi$ -Models	73
3.5	Conclusions	79
4	Delay Reconstruction: Dynamics versus Statistics	
	<i>Jaroslav Stark</i>	81
4.1	Introduction	81
4.2	Conventional Delay Reconstruction	84
4.3	Delay Reconstruction for Stochastic Systems	87
4.3.1	Random Dynamical Systems	87
4.3.2	Conjugacy for Random Dynamical Systems	88
4.3.3	Stochastic Takens Embedding Theorem	89
4.3.4	The Time Series Model	91
4.3.5	Technical Conditions	92
4.4	Trading-Off Noise and Complexity	93
4.5	Reconstruction for Spatio-Temporal Systems	96
4.6	Conclusions	101
5	Some Remarks on the Statistical Modeling of Chaotic Sys- tems	
	<i>Dominique Guegan</i>	105
5.1	Introduction	105
5.2	How Do We Decide If a Process Is Generated by a Stochastic or a Deterministic System Based on Observed Data?	106
5.3	What Features Are Captured by Stochastic Processes? . . .	108
5.4	Recent Developments in the Statistical Study of Chaotic Systems	111
5.4.1	Invariance and Mixing Properties in Chaotic Systems	112
5.4.2	Nonparametric Methods for an Estimation Theory in Chaotic Systems	116
5.5	The Deconvolution Problem	118

5.6 Further Developments and Remarks 121

6 The Identification and Estimation of Nonlinear Stochastic Systems

Peter Young **127**

6.1 Introduction 127

6.2 Time Varying Parameter Regression (TVPR) Models 129

 6.2.1 Maximum Likelihood (ML) Optimization of Hyperparameters 134

6.3 State Dependent Parameter Regression Models 135

6.4 State-Dependent Parameter Dynamic Models 142

 6.4.1 The SDARX Model 143

 6.4.2 The SDAR Model 153

 6.4.3 SDTF Estimation 159

6.5 Conclusions 160

II Fundamentals **167**

7 An Introduction to Monte Carlo Methods for Bayesian Data Analysis

Christophe Andrieu, Arnaud Doucet and William J. Fitzgerald **169**

7.1 Introduction 169

7.2 The Bayesian Approach 171

 7.2.1 Bayesian Model and Posterior Distribution 171

 7.2.2 Bayesian Decision 171

 7.2.3 Model Choice 172

 7.2.4 Discussion 174

7.3 Basics of Monte Carlo Methods 174

 7.3.1 Monte Carlo Integration 175

 7.3.2 Monte Carlo Optimization for MAP Estimation . . . 176

 7.3.3 Simulation Problems 177

 7.3.4 Classical Methods 177

 7.3.5 Importance Sampling 178

 7.3.6 Markov Chain Monte Carlo Methods 181

7.4 An Example: The Autoregressive Model 192

 7.4.1 Bayesian Estimation of AR Process 193

 7.4.2 Bayesian Model Selection of AR Process 198

7.5 Applications 200

 7.5.1 Sequential Importance Sampling for Optimal Filtering 201

 7.5.2 Blind Deconvolution of Impulsive Process 204

 7.5.3 Robust Bayesian Spectral Analysis 208

7.6 Conclusion 211

8	Constrained Randomization of Time Series for Nonlinearity Tests	
	<i>Thomas Schreiber and Andreas Schmitz</i>	219
8.1	Introduction	219
8.2	Probing Nonlinearity	220
8.3	Generating Constrained Realizations	221
	8.3.1 Fourier-Based Methods	222
	8.3.2 A General Randomization Scheme	224
8.4	Remaining Caveats	229
8.5	Discussion	231
9	Removing the Noise from Chaos Plus Noise	
	<i>Steven P. Lalley</i>	233
9.1	Introduction	233
9.2	Homoclinic Pairs	234
9.3	Sensitive Dependence on Initial Conditions	237
9.4	Example: The Henon Mapping	239
9.5	The Weak Orbit Separation Property	241
9.6	Concluding Remarks	243
10	Embedding Theorems, Scaling Structures, and Determinism in Time Series	
	<i>Colleen D. Cutler</i>	245
10.1	Introduction	245
10.2	Dynamical Systems and Determinism	247
10.3	Time-Delay Embeddings and Determinism	254
10.4	Scaling Structures and Determinism	257
10.5	Examples	259
11	Consistent Estimation of a Dynamical Map	
	<i>Andrew Nobel</i>	267
11.1	Introduction	267
	11.1.1 Overview	268
	11.1.2 Outline	269
11.2	Two Models for Dynamical Data	269
11.3	Multivariate Histograms	271
11.4	Ergodic Systems without Noise	272
	11.4.1 Estimation of Irregular Maps	272
	11.4.2 Remarks and a Corollary	274
11.5	Ergodic Systems with Dynamical Noise	275
	11.5.1 Continuous Maps	275
	11.5.2 Absolutely Continuous Noise	276

12 Extracting Dynamical Behavior via Markov Models	
<i>Gary Froyland</i>	281
12.1 Introduction and Basic Constructions	281
12.1.1 What Do We Do?	282
12.1.2 How Do We Do This?	282
12.1.3 Why Do We Do This?	283
12.2 Objects and Behavior of Interest	284
12.2.1 Invariant Measures	284
12.2.2 Invariant Sets	285
12.2.3 Decay of Correlations	286
12.2.4 Lyapunov Exponents	286
12.2.5 Mean and Variance of Return Times	287
12.3 Deterministic Systems	287
12.3.1 Basic Constructions	288
12.3.2 Invariant Measures and Invariant Sets	288
12.3.3 Decay of Correlations and Spectral Approximation	292
12.3.4 Lyapunov Exponents and Entropy	295
12.3.5 Mean and Variance of Return Times	299
12.4 Random Systems	303
12.4.1 Basic Constructions	305
12.4.2 Invariant Measures	307
12.4.3 Lyapunov Exponents	310
12.4.4 Mean and Variance of Return Times	312
12.4.5 Advantages for Markov Modeling of Random Dynamical Systems	313
12.5 Miscellany	314
12.6 Numerical Tips and Tricks	315
12.6.1 Transition Matrix Construction	315
12.6.2 Partition Selection	316
13 Formulas for the Eckmann-Ruelle Matrix	
<i>Timothy D. Sauer</i>	323
13.1 Introduction	323
13.2 Attractor Reconstruction	325
13.3 The Noiseless Case	326
13.4 The Medium Noise Case	328
13.5 Spurious Lyapunov Exponents	332
13.6 Concluding Remarks	334

III	Methods and Applications	337
14	Noise and Nonlinearity in an Ecological System	
	<i>Paul A. Dixon, Maria J. Milicich and George Sugihara</i>	339
14.1	Introduction	339
14.2	Univariate Analysis: Nonlinearity in the Larval Phase . . .	342
14.2.1	S-Map Analysis	342
14.2.2	Residual Delay Maps	344
14.2.3	Prediction Decay	346
14.2.4	Surrogate Analysis	346
14.3	Multivariate Analysis: Forecasting Larval Supply	352
14.3.1	Lunar Phase	352
14.3.2	Cross-Shelf Wind	355
14.3.3	Average Daily Wind	356
14.4	Conclusions	361
15	Cluster-Weighted Modeling: Probabilistic Time Series Prediction, Characterization, and Synthesis	
	<i>Bernd Schoner and Neil Gershenfeld</i>	365
15.1	Introduction	365
15.2	Cluster-Weighted Modeling	368
15.2.1	Architecture	368
15.2.2	Model Estimation	370
15.2.3	Error Estimation and Characterization	374
15.3	Application: How to Build a Digital Strad	377
15.3.1	Cluster-Weighted Spectral Modeling	378
15.3.2	Cluster-Weighted Sampling	379
15.3.3	Higher Order Factorization: Hierarchical Mixture Models and Hidden-Markov Models	381
15.4	Summary	384
16	Data Compression, Dynamics, and Stationarity	
	<i>Matthew B. Kennel and Alistair I. Mees</i>	387
16.1	Introduction	387
16.2	Symbolic Prediction and Coding	390
16.3	Context Trees	392
16.3.1	Zero-Memory Estimation	393
16.3.2	The Structure of the Tree Machine	395
16.4	Using Tree Machine Models	399
16.4.1	Compression	399
16.4.2	Simulation	399
16.4.3	Converting Tree Machines to Markov Chains	400
16.4.4	Topological Entropy	402
16.5	Application: Simulating a Chaotic Laser	402
16.6	Application: Stationarity Testing	404
16.7	Conclusions	410

17 Analyzing Nonlinear Dynamical Systems with Nonparametric Regression	
<i>Henning U. Voss</i>	413
17.1 Introduction	413
17.2 Nonlinear Regression Analysis and Maximal Correlation . .	414
17.3 High-Dimensional Dynamics (I): Analysis of Delayed-Feedback Systems	416
17.4 Interpretation of Models (I): Analysis of a Fiber Ring Resonator System	419
17.5 High-Dimensional Dynamics (II): Analysis of Spatially Extended Systems	424
17.6 Interpretation of Models (II): Analysis of a Convection Experiment	426
17.7 Conclusions	431
18 Optimization of Embedding Parameters for Prediction of Seizure Onset with Mutual Information	
<i>A. M. Albano, C. J. Cellucci, R. N. Harner and P. E. Rapp</i>	435
18.1 Introduction	435
18.2 Mutual Information	436
18.2.1 Embedding	437
18.2.2 Computationally Constructed Multivariate Data Sets	439
18.3 Multichannel EEG	444
18.3.1 Data and Analysis	444
18.4 Concluding Remarks	450
19 Detection of a Nonlinear Oscillator Underlying Experimental Time Series: The Sunspot Cycle	
<i>Milan Paluš</i>	453
19.1 Introduction	453
19.2 A Test for Nonlinearity Based on Redundancies and Surrogate Data	455
19.3 The Null Hypothesis of Nonlinearity Tests and Its Negations	457
19.4 An Example of Surface Air Pressure Series	458
19.5 Nonlinearity in the Sunspot Cycle	462
19.6 Amplitude-Frequency Correlation in Nonlinear Oscillators .	465
19.7 Amplitude-Frequency Correlation in the Sunspot Cycle . .	468
19.8 Discussion	470

Preface

All models are lies. “The Earth orbits the sun in an ellipse with the sun at one focus” is false, but accurate enough for almost all purposes. This book describes the current state of the art of telling *useful* lies about time-varying systems in the real world. Specifically, it is about trying to “understand” (that is, tell useful lies about) dynamical systems directly from observations, either because they are too complex to model in the conventional way or because they are simply ill-understood.

Because it overlaps with conventional time-series analysis, building models of nonlinear dynamical systems directly from data has been seen by some observers as a somewhat ill-informed attempt to reinvent time-series analysis. The truth is distinctly less trivial. It is surely impossible, except in a few special cases, to re-create Newton’s astonishing feat of writing a short equation that is an excellent description of real-world phenomena. Real systems are connected to the rest of the world; they are noisy, non-stationary, and have high-dimensional dynamics; even when the dynamics contains lower-dimensional attractors there is almost never a coordinate system available in which these attractors have a conventionally simple description.

Instead of describing the dynamics with the concise, powerful equations beloved of classical physics and applied mathematics, we have to be content with a *reconstruction* of the dynamics from data. That is, we represent the dynamics with computational algorithms, which (contrary to popular belief) can be used in most of the same ways as conventional models, but can seldom be written in a simple closed form that is pleasing to the eye.

The name “reconstruction” is, of course, a pun on “deconstruction”: the non-scientist’s view of science is often that it is excessively prone to isolate what it studies and to reduce everything to meaninglessness at worst, clockwork at best. However unfair this view, it is true that we choose to study what we can hope to model. Reconstruction theory tries to expand what we can hope to model by starting out from the axiom that, because the world is interconnected, we can only make approximate models of isolated parts of it, and there is no reason to suppose they will be as simple as Newton’s equation. Instead of striving for simplistic descriptions, it concentrates on powerful algorithms; instead of writing equations, it tries to grasp the geometry of phase space. The idea is to have simple *qualitative descriptions* but to rely on computers for quantitative predictions and estimates. This point is made right from our first chapter, where Abarbanel compares a conventionally built model with a reconstructed geometric model.

We see that reconstructing dynamics is not, in fact, trying to emulate Newton; it is Kepler whose shoulders we are standing on. Kepler reduced the huge body of data collected by Tycho Brahe to stunningly simple geometry: the planets orbit the sun in ellipses, and follow certain other simple laws. Kepler's dynamical invariants, like the rate at which the radius vector sweeps out area, have their counterparts in today's dynamical invariants like Lyapunov exponents, entropies, and dimensions. Today we are drowning in data, and we need insights; we have immensely powerful but stupid assistants, in the form of computers, so we can perhaps try to automate something of what Kepler did. It may be that another Newton will discover how to neatly encapsulate our current geometric descriptions, but it seems unlikely, and in the meantime, there are huge numbers of questions to be answered¹.

Many of the questions are statistical, because we are after all analyzing data. Others involve learning to build different sorts of models: control theorists have long realized that one can often treat some of the non-modeled variability in terms of known inputs from the world to the system. It is time dynamicists started to listen to statisticians, control systems engineers, information theorists, and signal processors. This book represents an attempt to get these groups together.

The chapters are based on some of the papers presented at a workshop on Nonlinear Dynamics and Statistics, which took place in September 1998, at the Isaac Newton Institute, Cambridge University, as part of the programme "Nonlinear and Non-stationary Signal Processing".

The aim of the workshop was to bring together workers in theoretical and applied nonlinear dynamics, statistics, signal processing, and systems engineering. For about a decade there had been a very exciting and rapidly growing body of work in dynamical reconstruction: It was clear that the work had matured to the extent that the nonlinear dynamics community had a lot to offer practitioners, but it was also clear that it now needed to take on board the insights and experience of statisticians and engineers. It was also hoped that the statisticians and engineers would find the reconstruction work interesting and inspiring. The consensus was that the workshop was extremely successful, with the famous Newton Institute blackboards—they are everywhere, even in the lifts—being in constant use. This book should be thought of as no more than a sampling from the many outstanding and thought-provoking presentations at the workshop.

The chapters were all written with the benefit of hindsight, so they incorporate some of what was learned from the very lively discussions that took place. By request of the editor, some of the papers (notably those by Froyland, by Andrieu et al. and by Young) are partly tutorial in nature, because they describe work that needs to be better known among the applied dynamics community. All of the papers are targeted at a broad audience of the kind at the conference, and the book should be useful to anyone interested in nonlinear dynamics and time-series analysis. At the

same time, I believe the contributions in this book are representative of the best that can be done in nonlinear dynamical reconstruction at the turn of the millennium.

The book has three parts: Issues in Reconstructing Dynamics, Fundamentals, and Methods and Applications. The divisions and the classifications of the papers are, of course, arbitrary: the chapters in each part have a great deal to say about the topics of the other parts. The division may, however, help the reader know where to start.

The first part sets the scene, illustrating some of the problems we face when we try to reconstruct nonlinear dynamics from observed data.

- Abarbanel describes the situation as seen by the dynamics community: he shows what we want to know and what we do know, and he analyzes a data set in a way that will be familiar to most dynamicists, but less familiarly, he discusses what is good and bad about this way of doing things.
- Smith discusses a number of fundamental issues in the interaction between observational uncertainty and model error in reconstructed dynamics. Given that all models tell lies, the perfection we seek is beyond our grasp, and it is not easy to decide which model gives the most understanding. Perhaps we should keep many? The idea of keeping many models reappears later, when Kennel and Mees use weighted models based on work in data compression.
- Judd et al. point out that modeling dynamics usefully is more than just a question of getting good short-term predictions—an insight that is also a theme of Schoner and Gershenfeld’s contribution later in the book—indeed, it is being increasingly recognized as a key question. They also emphasize the importance of separating noise from signal and the consequent usefulness of information theoretic ideas.
- Stark makes it clear that the standard embedding process is fraught with danger and should be subject to careful scrutiny. In the process, he mentions many successful extensions: the take-home lesson is perhaps that embedding is even more powerful than is often supposed, but that it needs far more care than is usually exercised.
- Guegan takes the bull by the horns and discusses the statistical modeling of chaotic time series and its perils and rewards. This is, of course, the theme of the book: reconstruction requires both dynamics and statistics to be considered, and Guegan speaks from the informed statistician’s viewpoint.
- Young speaks from the point of view of a systems engineer and points out that most practical dynamical problems are best posed as input-output systems, yet they are modeled as autonomous systems in the

dynamics community. A merger of his methods with the advances in nonlinear dynamics is long overdue.

Chapters in the second part examine fundamental questions related to modeling nonlinear time series (or reconstructing dynamics). The point of view varies from chapter to chapter, with some taking a more statistical approach, others considering fundamental questions of ergodic theory, others looking at basic questions in dynamical reconstruction, and still others describing modeling methods. As set out herein, there is a certain drift from statistics to dynamics as we progress through the chapters, but of course this drift is highly stochastic.

- Andrieu et al. present a survey of the modern theory of Monte Carlo methods; the great advances made in the past few years mean that many dynamical modeling problems, which are often representable in Bayesian terms, are now relatively tractable: something that should improve nonlinear dynamical models immensely.
- Schreiber and Schmitz deal with nonlinearity tests, which have been significant in recent years as a precaution against claiming nonlinearity in the absence of strong evidence.
- Lalley shows that a chaotic system with unbounded noise cannot be de-noised (at least in the conventional sense), and then he shows how bounded noise can be tackled successfully.
- Cutler discusses definitions of deterministic and stochastic for time-series and relates these to dynamical systems. She shows that it is not difficult to produce simple examples where the system cannot be reconstructed owing to poor choice of observable. The Grassberger-Procaccia algorithm converges in these cases even though embedding is never achieved.
- Nobel investigates how to approximate dynamics and contrasts deterministic models with models that include dynamic noise, while taking into account the oft-neglected issue of statistical consistency. The investigation sheds light on the question of distinguishing deterministic systems from stochastic systems, a question also considered in Guegan's chapter.
- Froyland estimates a system's invariant density via Markov models, which is often a superior way to compute statistical properties. Knowing the invariant density is arguably both the best that can be done and the best thing to do. His chapter is a detailed review of the latest results in the area.
- Sauer shows us that a calculation that is implicit in a great deal of dynamical analysis—estimation of a local derivative—is more subtle

than was thought and has surprising (perhaps pleasantly surprising) properties in the presence of noise.

The final part of the book is devoted to modeling methods and nontrivial applications.

- Dixon et al. discuss a successful fisheries model and describe the modeling approach that they used, the S -map method. They add more weight to one of the familiar themes of this book, the need to understand the interplay between noise and nonlinearity.
- Schoner and Gershenfeld discuss another successful reconstruction method, cluster-weighted modeling, and apply it to an exceptionally challenging problem in synthesis of audio signals: the “digital Stradivarius”.
- Kennel and Mees borrow work from the data compression literature and demonstrate that it has a lot to offer the dynamics community; they show an application to stationarity testing of fluidized-bed reactors. The models that are produced are inherently probabilistic, and predictions are conditioned on a discrete analogue of variable embeddings.
- Voss uses nonlinear non-parametric regression for the analysis of experimental data, which can greatly simplify the modeling process if there is a certain amount of knowledge available about the system’s structure.
- Albano et al. show how estimation of average mutual information, applied in an unusual way, can be useful in analysis of EEG signals, with specific application to early warning of epileptic seizures.
- Paluš uses a deep understanding of both dynamical and statistical fundamentals to examine the well-known sunspot data series and concludes that it is indeed most likely generated by a nonlinear dynamical system in spite of the fact that previously the best linear stochastic models seemed to be as good as the best nonlinear ones.

All of the chapters have been refereed anonymously; I thank the referees for this essential and underappreciated service: you know who you are.

The statement about useful lies that opened this preface is paraphrased from Richard Bandler. I thank Lenny Smith and Paul Rapp for valuable feedback on the rest of the preface; any useless lies that remain are my own responsibility.

This volume and the workshop that inspired it would not have been possible without the generous support of the Isaac Newton Institute and its sponsors; I am particularly grateful to Bill Fitzgerald and Richard Smith, who invited me to participate in the Nonlinear and Non-stationary Signal

Processing Programme. The programme has its own official proceedings volume, edited by Fitzgerald, Smith, and Young, to be published by Cambridge University Press in 2000.

I thank the Australian Research Council for partial financial support, and the University of Western Australia for leave to attend the programme. I also thank the Department of Systems Engineering and Engineering Management at The Chinese University of Hong Kong for hospitality.

Perth, Western Australia
New Year's Day, 2000.

Alistair I. Mees

Contributors

- Henry D. I. Abarbanel** Department of Physics; and
Marine Physical Laboratory, Scripps Institution
of Oceanography,
University of California at San Diego, La Jolla,
CA 93093-0402, USA
`hdia@jacobi.ucsd.edu`
- Alfonso M. Albano** Department of Physics, Bryn Mawr College,
Bryn Mawr, PA 19010-2899, USA
`aalbano@brynmawr.edu`
- Christophe Andrieu** Signal Processing Laboratory, Department of
Engineering, University of Cambridge,
Cambridge CB2 1PZ, UK
`ca226@eng.cam.ac.uk`
- Christopher J. Cellucci** Department of Physics, Ursinus College,
Collegeville, PA 19426-1000, USA
`ccellucci@ursinus.edu`
- Colleen Cutler** University of Waterloo, Canada
`cdcutter@math.uwaterloo.ca`
- Paul A. Dixon** Scripps Institution of Oceanography,
University of California at San Diego, La Jolla,
CA 92093-0202
`paul@complex.ucsd.edu`
- Arnaud Doucet** Signal Processing Laboratory, Department of
Engineering, University of Cambridge,
Cambridge CB2 1PZ, UK
`ad2@eng.cam.ac.uk`
- William J. Fitzgerald** Signal Processing Laboratory, Department of
Engineering, University of Cambridge,
Cambridge CB2 1PZ, UK
`wjf@eng.cam.ac.uk`

- Gary Froyland** Department of Mathematics and Computer Science, University of Paderborn, Warburger Str. 100, Paderborn 33098, Germany
froyland@upb.de
- Neil Gerschenfeld** MIT Media Laboratory, 20 Ames St., Cambridge, MA 02139, USA
gersh@media.mit.edu
- Dominique Guegan** Université de Reims, CREST Laboratoire de Statistique, Timbre J340, 3 avenue Pierre Larousse, 92245 Malakoff Cedex France
guegan@ensae.fr
- Richard N. Harner** BrainVue Systems, 3901 Netherfield Road, Philadelphia, PA 19129, USA
brainvue@msn.com
- Kevin Judd** Centre for Applied Dynamics and Optimization, The University of Western Australia, Nedlands, Perth WA 6907, Australia
Kevin.Judd@uwa.edu.au
- Matt B. Kennel** Institute for Nonlinear Science, University of California at San Diego, La Jolla, CA 92093-0402, USA
mbk@lyapunov.ucsd.edu
- Steven P. Lalley** Department of Statistics, University of Chicago, 5734 University Avenue, Chicago IL 60637, USA
lalley@stat.purdue.edu
- Alistair I. Mees** Centre for Applied Dynamics and Optimization, The University of Western Australia, Nedlands, Perth WA 6907, Australia
Alistair.Mees@uwa.edu.au
- Maria J. Milicich** M2 Environmental Ltd, 15D, Trust Tower, 68-74 Johnston Rd, Wan Chai, Hong Kong
m2@netvigator.com

- Andrew Nobel** Department of Statistics,
University of North Carolina,
Chapel Hill, North Carolina 27599, USA
nobel@stat.unc.edu
- Milan Paluš** Institute of Computer Science,
Academy of Sciences of the Czech Republic
Pod vodárenskou věží 2, 182 07 Prague 8,
Czech Republic
mp@cs.cas.cz
- Paul E. Rapp** Alfred P. Noyes Clinical Research Center,
Norristown State Hospital,
Norristown, PA 19401-5397 USA
Paul.E.Rapp@drexel.edu
- Tim Sauer** Department of Mathematical Sciences, George
Mason University,
Fairfax, VA 22030, USA
tsauer@gmu.edu
- Bernd Schoner** MIT Media Laboratory, 20 Ames St., Cam-
bridge, MA 02139, USA
schoner@media.mit.edu
- Andreas Schmitz** Department of Physics, University of Wuppertal,
42097 Wuppertal, Germany
schmitz@theorie.physik.uni-wuppertal.de
- Thomas Schreiber** Department of Physics, University of Wuppertal,
42097 Wuppertal, Germany
schreibe@theorie.physik.uni-wuppertal.de
- Michael Small** Department of Physics, Heriot-Watt University,
Edinburgh, Scotland, UK
- Leonard A. Smith** Mathematical Institute, University of Oxford,
Oxford, UK
lenny@maths.oxford.ac.uk
- Jaroslav Stark** Centre for Nonlinear Dynamics and its Applica-
tions, University College London,
Gower Street, London WC1E 6BT, UK
j.stark@ucl.ac.uk

George Sugihara

Scripps Institution of Oceanography,
University of California at San Diego,
La Jolla, CA 92093-0202, USA
george@complex.ucsd.edu

Henning U. Voss

University of Freiburg, Physics Department,
Hermann - Herder Str. 3, D-79104 Freiburg,
Germany
hv@physik.uni-freiburg.de

Peter Young

Centre for Research on Environmental Systems
and Statistics, Lancaster University, UK
p.young@lancaster.ac.uk

Part I

**Issues in Reconstructing
Dynamics**

Chapter 1

Challenges in Modeling Nonlinear Systems: A Worked Example

Henry D. I. Abarbanel

ABSTRACT

The interaction between nonlinear dynamics and statistics has been rather limited over the two decades of concentrated work in nonlinear systems by physical and biological scientists. This chapter is meant to be a contribution to stimulating that interaction by presenting a discussion of a problem in biology which is addressed by tools of nonlinear dynamics and by posing, along the way, issues of statistical relevance not answered by the community of nonlinear dynamicists.

The overall issue is that of characterizing and modeling nonlinear systems using observed data. Typically this is in the initial absence of a model for the source of the data, but that often is the goal of the analysis. Models derived from these data can be black box or analytic. Black box models typically consist of a set of numerical rules for prediction or control of the system in the absence of any fundamental knowledge of the physics or biology of the system. Analytic models attempt to incorporate knowledge from the observations and their analysis into sets of differential equations or maps embodying the properties of the measured processes.

In this chapter we focus on the analysis of membrane voltage data from identified neurons of the stomatogastric ganglion of the California spiny lobster with the goal of modeling individual neurons and their oscillatory behavior in a variety of environmental circumstances. The membrane voltage dynamics of these neurons is typically low dimensional and chaotic. Hodgkin-Huxley models describing the ion currents which flow through the membrane are not sufficient to capture this behavior, but the addition of a slow background dynamics, which we attribute to the storage and release of calcium in the cell, permits an excellent description of the observations.

We will describe the experiments, the analysis of the data, and the model building connected with these statements, and hopefully we'll leave the reader with the sense that much has been done, but much more is required to transform what seems to be a working set of usable implements into a scientifically sharp collection of tools.

1.1 Summary and Preview

We begin with a summary of the issues, and then we turn to details of the experimental system at hand.

In the biological observations, as in a wide class of measurements, we observe a single scalar time series of signals $s(t) = s(t_0 + n\tau_s) = s(n)$ measured each τ_s seconds starting at some time t_0 . If the signal source is stationary in the statistical sense, t_0 is not important. The actual state of the dynamical system producing the signals is not scalar but is composed of many variables collected in a state vector $\mathbf{x}(n)$ which follows a dynamical rule

$$\mathbf{x}(n+1) = \mathbf{F}(\mathbf{x}(n)), \quad (1.1)$$

and the observations are given by another rule

$$s(n) = h(\mathbf{x}(n)). \quad (1.2)$$

The rules $\mathbf{F}(\bullet)$ and $h(\bullet)$ are generally unknown to us. Nonetheless, we wish to characterize the system by statistical quantities associated with the observed states $\mathbf{x}(n)$; $n = 1, 2, \dots, N$.

We are led to ask the following questions:

How, and with what accuracy, can we determine nonlinear statistical properties associated with the states $\mathbf{x}(n)$ from observations of the $\{s(n)\}$ alone?

The kind of statistical quantity we can determine which is a characteristic of the system, not of any individual orbit $\mathbf{x}(n)$ of the system includes:

- various dimensions D_q which classify the way in which local moments of the $\mathbf{x}(n)$ vary on the average over the set of observations. This set constitutes the attractor for the source of the observations,
- the spectrum of Lyapunov exponents λ_a $a = 1, 2, \dots, d_L$ associated with the dynamics,
- topological properties of the attractor,
- unstable periodic orbits which form the backbone of the attractor,
- and perhaps many more not currently being studied or pursued.

There are clearly some questions here for statisticians. The identification of which statistical quantities, such as D_q and λ_a are important have long been addressed in the nonlinear dynamics literature. We want to know the rate at which an algorithm converges to D_q or λ_a . The rate of convergence is not only in terms of number of data but also in terms of sampling time. One must not only have a sufficient number of data, but more critically the data must represent a full sampling of the properties of the system attractor,

and one should have a criterion for proper sampling times—oversampling allows the acquisition of lots of irrelevant data.

We want the expected error as well as the variance in the estimation of these quantities, but this is not enough. The harder question, without a definite answer to this date, concerns the main goal which is *model verification*. Suppose on the basis of various considerations we develop a model of the evolution of the state $\mathbf{x}(n) \rightarrow \mathbf{x}(n+1) = \mathbf{G}(\mathbf{x}(n))$. This could be a “black box” model built on the data and the properties of the attractor, or it could be a set of differential equations built on some reasoning about the physical or biological processes acting in the system.

How do we verify this model?

Clearly we need to make sure the values of the D_q or the λ_a are the same from the model as from the data. Is this sufficient? What constitutes an acceptable error ΔD_q or $\Delta \lambda_a$? Is there a *complete* set of invariants such as the λ_a which we should use to compare our model $\mathbf{x}(n+1) = \mathbf{G}(\mathbf{x}(n))$ to observations $\{s(n)\}$?

And once we have established good relations among the global quantities, what about *local* properties on the attractor? Are there local things such as a $D_q(\mathbf{x})$ which we should evaluate and compare?

Finally, while we have developed some sense of a reliable technology for many of these activities for dimensions up to seven or ten, what do we do about the many interesting systems with $d > 10$?

We are almost ready to begin the actual work. I do not know the answers to all the questions I just posed, but hopefully research will answer them over the coming years.

1.2 Small Neural Assemblies: What We Did

In many animals, including humans, there are a number of small neural assemblies responsible for generating rhythmic behavior for appropriate functions. These systems, called central pattern generators (CPGs) [1] by biologists, produce rhythmic electrical outputs to drive muscles for particular purposes. At the base of our spines is a CPG which aids in our ability to walk regularly.

In our laboratory at UCSD we have extensively studied a CPG in the stomatogastric ganglion of the California spiny lobster *Panrulis interruptus*. This CPG, the pyloric CPG, governs the passage of shredded food from the stomach region where it is shredded to the digestive tract where its nutrients are transferred into the lobster’s system. The pyloric CPG consists of fourteen neurons corresponding to six different individuals. One of these neurons, the lateral pyloric neuron or LP, has been a focus of our work. We are able to isolate the LP from its environment and measure long

time series of the electric potential across its cellular membrane. The measurements are performed by inserting a fine electrode of diameter about 1 micron at its tip through the cell wall. The resistance at the electrode is of order 10^7 ohm, and with voltage differences of about a millivolt, we measure currents of a nanoamp. A typical time trace from the LP neuron is shown in Figure 1.1 where $\tau_s = 0.5$ ms, and we are able to routinely take samples of many minutes resulting in a data set of $N \approx 10^7$ samples.

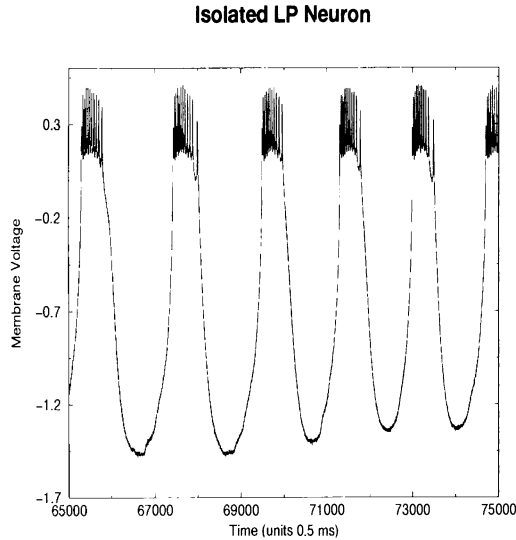


FIGURE 1.1. Time series of membrane voltage from an isolated LP neuron from the lobster stomatogastric ganglion; $\tau_s = 0.5$ ms.

The two essential features here are the slow oscillations, the bursts at about 1 Hz and the rapid spiking oscillations on the top of each burst. The time scale for a spike is about 10 ms. The spikes are the high-frequency oscillations which are transmitted down axons connecting neurons, and they are responsible for communication among connected neurons.

This chapter tells how we analyzed $\{s(t)\}$ from these data, and created and tested a model for the isolated LP neuron. The challenge is to understand:

- how accurate our analysis was,
- how meaningful our analysis was,
- how we should do this in the future; and
- how we get there.

1.3 General Outlook

The methods of analysis we use are based in the *time domain*. Fourier methods, so widely and productively used in analysis of linear systems, are inappropriate for nonlinear dynamical systems. We start from $s(n); n = 1, 2, \dots, N$; from this data set we

- create a d -dimensional state space of vectors which are a dynamical proxy for the motion of the observed system in its actual (unknown) space of states. These vectors

$$\mathbf{y}(n) = [s(n), s(n - T), s(n - 2T), \dots, s(n - T(d - 1))], \quad (1.3)$$

replace the scalar data $s(n)$ as the object of our investigations.

- determine the integer time lag T (in units of τ_s) which makes the components of $\mathbf{y}(n)$ independent in a nonlinear sense—this uses the statistical quantity *average mutual information*.
- determine the number (integer) $d = d_E$ of components of $\mathbf{y}(n)$ required by the data to unambiguously represent the trajectories of the system in state space—this uses the statistical quantity *global false nearest neighbors*.
- determine the number (integer) $d = d_L$ of dynamical degrees of freedom acting locally in the state space to evolve the system forward in time—this uses the statistical quantity *local false nearest neighbors*.
- determine the spectrum of Lyapunov exponents $\lambda_a; a = 1, 2, \dots, d_L$ which establish the predictability of the source of our observations and permit an estimation of the fractional dimension of the attractor in state space.
- use the geometrical structure of the attractor to create *black box* models of for prediction of future states of the system in the proxy state space.

We will take data from the pyloric CPG of lobster as well as data from a three-degree-of-freedom nonlinear electrical circuit built at UCSD¹ for testing many of the ideas we discuss.

¹N.F. Rulkov, private communication

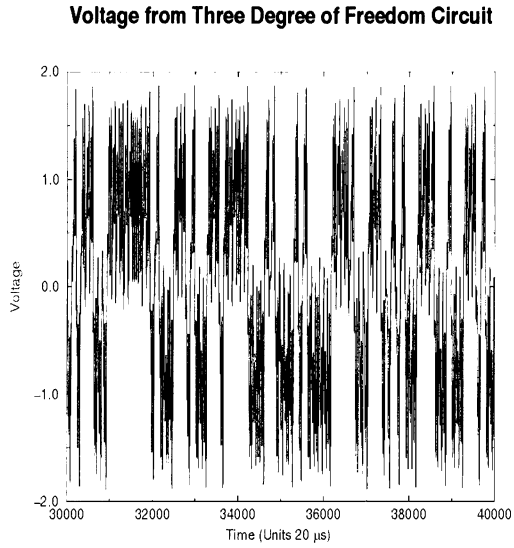


FIGURE 1.2. Time series of voltage from a three-degree-of-freedom chaotic electrical circuit; sampling time is $20 \mu\text{s}$.




FIGURE 1.3. Average mutual information for data in Figure 1.2.

Let's start with data from the nonlinear circuit. In Figure 1.2 we show a section of a time series of a voltage from this circuit. These data are collected with a sampling time of $\tau_s = 20\mu\text{s}$. In Figure 1.3 we show the average mutual information for this data set. It shows a clear first minimum at $T = 15$. We will have more to say about minima of average mutual information in a bit.

The key idea in the analysis is that signals which show complex temporal waveforms are likely to have originated in a multidimensional source when their Fourier spectra are not a set of sharp lines; in that case they are quasi-periodic and possibly linear sources. We need to reconstruct the state space for these systems from their scalar observations $\{s(n)\}$. We achieve this by recognizing that $s(n)$ and $S(n + T)$ for some integer T are independent samplings of the state of the nonlinear system. During time $T\tau_s$ the system evolves under the influence of all dynamical variables and the unobserved variables are now reflected in $s(n + T)$ in some unknown and nonlinear fashion. In constructing the data vectors

$$\mathbf{y}(n) = [s(n), s(n - T), s(n - 2T), \dots, s(n - T(d - 1))], \quad (1.4)$$

we are guaranteed by a geometrical theorem that (1) when we have an infinite amount of infinitely accurate data, the actual value of T is not important, and (2) when d is large enough, the attractor is unfolded from its projection on the observation $s(n)$ axis. This first fact is not useful, so we need some method for selecting a T .

1.4 The Neuron Model

For completeness, but not in its entirety, we present the model we developed for describing the oscillations of an individual neuron in the lobster pyloric CPG [2]. It is a model of Hodgkin-Huxley type. It consists of elementary circuit equations relating the circuit voltages in two sections of the neuron, the variables $V(t)$ and $V_1(t)$, to the currents flowing through the cell membrane. In addition to these currents, widely discussed over the past forty years, we add a slower dynamics for the intracellular concentration of calcium. The net result is the following detailed, but slightly ugly, model.

1.4.1 Ca^{2+} Dynamics

$$\dot{C} = j_{\text{rel}} - j_{\text{fil}} - j_{\text{out}} \quad (1.5)$$

$$\dot{C}_{\text{er}} = -(j_{\text{rel}} - j_{\text{fil}})/\sigma \quad (1.6)$$

$$\dot{h} = \frac{h_{\infty} - h}{\tau_h} \quad (1.7)$$

$$j_{\text{fil}} = V_{\text{erp}} \frac{C^2}{C^2 + K_{\text{erp}}^2} \quad (1.8)$$

$$j_{\text{rel}} = (P_{\text{leak}} + P_{IP_3} a_{\infty} b_{\infty} d_{\infty} h) (C_{\text{er}} - C) \quad (1.9)$$

$$j_{\text{out}} = \nu_{\text{pmp}} \frac{C^2}{C^2 + K_{\text{pmp}}^2} + \nu_{\text{pmex}} \frac{C^4}{C^4 + K_{\text{pmex}}^4} + \alpha(I_{Ca1} + I_{Ca2}) \quad (1.10)$$

$$\Gamma(x, y, z) = \frac{1}{1 + e^{\frac{x-y}{z}}} \quad (1.11)$$

$$a_{\infty} = \Gamma(\theta_a, C, k_a) \quad (1.12)$$

$$b_{\infty} = \Gamma(\theta_b, IP_3, k_b) \quad (1.13)$$

$$d_{\infty} = 0.2(1 + 4\Gamma(C_{\text{er}}, \theta_d, k_d)) \quad (1.14)$$

$$h_{\infty} = \Gamma(C, \theta_h, k_h) \quad (1.15)$$

$$k_a = \bar{k}_a \left(0.8 + \frac{IP_3}{IP_3 + 0.2} \frac{0.15^2}{0.15^2 + (IP_3 - 0.4)^2} \right) \times \frac{60}{60 + C_{\text{er}}} \quad (1.16)$$

$$k_h = \bar{k}_h \left(0.05 + \frac{IP_3^2}{IP_3^2 + 1 + \frac{180}{C_{\text{er}}}} \right) \quad (1.17)$$

$$\tau_h = \frac{\bar{\tau}_h}{b_{\infty} d_{\infty} \cosh \frac{C - \theta_i}{k_i}} \quad (1.18)$$

1.4.2 Voltage Dynamics

$$\dot{V} = (-I_{Ca1} - I_{Ca2} - I_l - I_{K(Ca)} - I_h - I_{V,V1})/c_m \quad (1.19)$$

$$\dot{V}_1 = (-I_{Na} - I_{l1} - I_{Kd} + I_{V,V1})/c_{m1} \quad (1.20)$$

$$I_i = g_i m_i^{q_i, m} h_i^{q_i, h} r_i(V) \quad (1.21)$$

$$\dot{n}_i = (e_{i,n} - n_i)/\tau_{i,n}, \quad (n = m, h) \quad (1.22)$$

I_i	n	$e_{i,n}$	$q_{i,n}$	$\tau_{i,n}$	g_i	$r_i(V)$
I_{Ca1}	m	$\Gamma(-V, 33.1, 13.18)$	3	60-40 $\Gamma(-V, 53.1, 20.5)$	0.172	$\frac{-V}{\exp \frac{2FV}{RT}} - 1.0$
	h	$\Gamma(V, -23.1, 5.5)$	1	150		
I_{Ca2}	m	$\Gamma(-V, -6.9, 17)$	3	37.14-25.86 $\Gamma(-V, 10.1, 26.4)$	3.75	$\frac{-V}{\exp \frac{2FV}{RT}} - 1.0$
$I_{K(Ca)}$	m	$\Gamma(V, 2.5-f(C-0.5), -13)x$	1	5/3	0.06	$\frac{-V}{(V+80)x}$
		$\Gamma(V, -30.5-f(C-0.5), -3.5)$				$\frac{C^4}{C^4 + K^4_{K(Ca)}}$
I_h	m	$\Gamma(-V, -43.3, 6.5)$	1	272+1499 $\Gamma(-V, 27.2, 8.73)$	0.024	V+20
I_t					0.0024	V+65
$I_{V,V1}$					0.072	V-V ₁
I_{I1}					0.024	V ₁ +65
I_{Na}	m	$\Gamma(-V_1, 4.5, 5.29)$	3	constant: $m_{Na} = m_{Na\infty}$	80	V ₁ -50
	h	$\Gamma(V_1, -28.9, 5.18)$	1	0.67(1.5+ $\Gamma(V_1, -14.9, 3.6)$)x $\Gamma(-V_1, 42.9, 10)$		
I_{Kd}	m	$\Gamma(-V_1, -7.7, 11.8)$	4	7.2-6.4 $\Gamma(-V_1, 8.3, 19.2)$	13	V ₁ +80

TABLE 1.1. The parameters of the model.

1.4.3 Parameters

Table 1.1 shows the parameters of the model. The voltage values are in mV and the g_i in μS . The parameter values are as follows: $\sigma=0.6$, $V_{\text{cell}}=2.671\text{mV}$, $f_{\text{cyt}}=0.01$, $\theta_a = 0.4 \mu\text{M}$, $\theta_b = 0.6 \mu\text{M}$, $\theta_d = 20 \mu\text{M}$, $\theta_h = 0.36 \mu\text{M}$, $\theta_t = 0.35 \mu\text{M}$, $k_b = 0.2 \mu\text{M}$, $k_d = 10 \mu\text{M}$, $k_t = 0.18 \mu\text{M}$, $\bar{k}_a = 0.14 \mu\text{M}$, $\bar{k}_h = 0.46 \mu\text{M}$, $K_{K(Ca)} = 0.5 \mu\text{M}$, $K_{\text{erp}} = 0.2 \mu\text{M}$, $K_{\text{pmp}} = 0.1 \mu\text{M}$, $K_{\text{pmex}} = 0.9 \mu\text{M}$, $\nu_{\text{pmp}} = 0.0145 \mu\text{Ms}^{-1}$, $\nu_{\text{pmex}} = 0.145 \mu\text{Ms}^{-1}$, $P_{\text{leak}} = 0.0286 \text{ s}^{-1}$, $P_{IP_3} = 3.571 \text{ s}^{-1}$, $V_{\text{erp}} = 3.762 \mu\text{Ms}^{-1}$, $\bar{\tau}_h = 1.25 \text{ s}$, $\alpha = 0.0194 \mu\text{M} (\text{nAs})^{-1}$, $c_m = 0.5 \text{ nF}$, $c_{m1} = 0.33 \text{ nF}$, $f = 2 \text{ V}\mu\text{M}^{-1}$, $F/RT = 0.04095 \text{ mV}^{-1}$, $T = 283\text{K}$.

1.4.4 Comment on This Model

From a physicist's point of view, this model is strikingly unpleasant. It is complex and has many remarkably accurately stated constants which are probably unknown to the stated accuracy. It assumes detailed forms of the vector field for the differential equation. Its virtue is that it tries to identify the particular ion currents which are responsible for the various aspects of the bursting (slow oscillations of about 1 Hz) and rapid spikes (fast oscillations of about 100 Hz) seen in the wave forms of Figure 1.1. Many of us, including this author, would prefer to see significantly simplified models, perhaps starting from this level of description. At present, we do not have a workable projection algorithm which would take this thirteen dimensional model and reduce it to a few dimensional model with much the same membrane voltage characteristics. We do, fortunately, have a systematic way to determine the dimension of that simplified model, if it exists, and to establish some of the properties it must have to correspond to the data [3]. We now turn to this.

1.5 Choosing T

Our task is to establish a criterion which selects the time delay T so that the signal at time $t_0 + n\tau_s$, $s(n)$ is independent enough of the signal at time $t_0 + (n + T)\tau_s$, $s(n + T)$. We want this independence to be such that the role of the other dynamical variables in the signal source, the ones we do not observe, is significantly reflected in the value of $s(n + T)$.

A standard linear criterion for this is to evaluate the autocorrelation function of the data and select that T where it first vanishes. This yields a criterion for the linear independence of the data at n and $n + T$. It is questionable what value this linear independence may be.

A standard estimator of nonlinear dependence comes from the answer to the question how much in *bits* do we learn about $s(n + T)$ from measurements of $s(n)$, on average over the data? This is given by the average

mutual information $I(T)$ introduced by Shannon fifty years ago:

$$I(T) = \sum_{\{s(n)\}, \{s(n+T)\}} P(s(n), s(n+T)) \log_2 \left[\frac{P(s(n), s(n+T))}{P(s(n))P(s(n+T))} \right], \quad (1.23)$$

where $P(s(n))$ is the normalized histogram of values taken by the signal $s(n)$, and $P(s(n), s(n+T))$ is the normalized joint histogram of the signals $s(n)$ and $s(n+T)$. $I(T)$ serves as a nonlinear correlation function for the data which can replace the standard linear autocorrelation function.

As a useful criterion for selecting T from this nonlinear correlation function, we follow Fraser and Swinney [4] who suggested that the first minimum of $I(T)$ would serve well. This is not an optimum in any sense, just a useful heuristic criterion. It serves our purpose as it provides a time lag not too long when numerical or measurement inaccuracies dominate the observations or too short when the other dynamical variables have not had a chance to act significantly.

For the nonlinear circuit $I(T)$ was shown in Figure 1.3, and our criterion selects $T = 15$ as a useful time lag for creating data vectors. Because the theorem we are using suggests that *any* T will do, one should use several T 's in the vicinity of the one selected to verify that subsequent quantities are independent of this choice.




FIGURE 1.4. Average mutual information for the data from an isolated LP neuron; $\tau_s = 0.5$ ms.

For the data from our isolated LP neuron and from our model of this neuron, we have $I(T)$ in Figures 1.4 and 1.5.




FIGURE 1.5. Average mutual information from our detailed Hodgkin-Huxley conductance model of the isolated LP neuron; $\tau_s = 0.5$ ms.

1.6 Choosing d : Global Dimension d_E

Our measured data $s(n)$ are points projected down from a larger dimensional space onto the axis of observations. Can we hope to “unproject” these data? Without knowledge of the full state of the system, we cannot really find the original point in state space from which our observation came. However, we can make a proxy space from the data vectors $\mathbf{y}(n)$ and ask when points in that space are determined in an unambiguous fashion. We assume that our data come from a statistically stationary source; namely, we say there are no external forces acting on this source and that during the time over which we make the observations all system parameters, conductances, resistances, inductances, capacitances, temperatures, and so on are constant. Then we have a set of differential equations which determine the signal, and the solutions to those differential equations are unique. This means the trajectories in state space do not overlap, and by choosing enough coordinates we can “unoverlap” the observations.

When two points on the attractor, the time asymptotic representation of the signal in state space, are in fact far apart but they appear close because of projection, these are *false neighbors* in the projected space. When two points on the attractor are close because of dynamics, they, too, will

appear close in the observation space, but they will be true neighbors. To distinguish between these, we construct our data vector $\mathbf{y}(n)$ in some dimension d and identify the nearest neighbor of each point in this space. Call the nearest neighbor of $\mathbf{y}(n)$ $\mathbf{y}^{NN}(n)$. The time index on this nearest neighbor could be far away from n , but in state space in dimension d they are closest. Now add a dimension and ask the same question of $\mathbf{y}(n)$ and $\mathbf{y}^{NN}(n)$ as seen in $d + 1$. In going from $d \rightarrow d + 1$

$$\begin{aligned} \mathbf{y}(n) &\rightarrow [\mathbf{y}(n) + s(n + dT)] \\ \mathbf{y}^{NN}(n) &\rightarrow [\mathbf{y}^{NN}(n), s^{NN}(n + dT)]. \end{aligned} \quad (1.24)$$

So if $|s(n + dT) - s^{NN}(n + dT)|$ is small, the two nearest neighbors in dimension d are close in dimension $d + 1$. They are true neighbors; if not, they are false neighbors. When the number of false neighbors drops to zero, in principle, we have identified the dimension in which one has unfolded the attractor. In practice, if the percentage of false nearest neighbors drops below some threshold, say 1%, then we accept that dimension. We call this dimension d_E ; it represents the global integer dimension required to unfold the attractor of the signal source.

What is wrong with working in too low a dimension? For some things, such as the calculation of the D_q it doesn't matter. However, if one wants to make predictive models, the two false neighbors will move into different domains of state space in time, so we would make serious mistakes in creating a model that moved them ahead more or less in the same way. Similarly if one wants to evaluate the spectrum of Lyapunov exponents, choosing d too small is sure to lead to errors.

In Figure 1.6 we have the plot of false nearest neighbors versus dimension for the data from our circuit. It is clear that $d_E = 3$ is selected. In Figures 1.7 and 1.8, we have the same plot for observed data from the isolated LP neuron and for model output for the membrane voltage. Each of these is compatible with $d_E = 5$. This is actually quite a striking result. The model has thirteen degrees of freedom and the actual neuron at least that many. Yet, in measurements or calculations of the membrane voltage fewer dimensions play a dynamical role. Models that wish to reproduce all thirteen (or more) degrees of freedom are, in some sense, too complicated or, perhaps better said, have more information in them than required for the role membrane voltage plays in the neural dynamics. The impetus for simplified modeling is quite strong.

Because d_E is small enough, we may learn something (visual at least) by plotting the three-dimensional vectors $[s(n), s(n + T), s(n + 2T)]$ for each data set. We do this in Figures 1.9, 1.10 and 1.11.

1.7 Dynamical Dimension

The global dimension d_E tells us how many coordinates we need to unfold the attractor from its projection onto the observations. This may not be the number of integer dimensions d_L required in the dynamical equations




FIGURE 1.6. Global false nearest neighbors for data from the nonlinear electrical circuit.

to produce the observed time evolution of the system: $d_L \leq d_E$. As an example, suppose we have a map of the plane to itself, so $(x_n, y_n) \rightarrow (x_{n+1}, y_{n+1}) = F(x_n, y_n)$, but the global motion lies on a torus. To unfold the torus requires $d_E = 3$ while $d_L = 2$ for the dynamics.

To find d_L , we search through **local** rules (maps) in state space which take $\mathbf{y}(n) \rightarrow \mathbf{F}(\mathbf{y}(n)) = \mathbf{y}(n+1)$. We do not know these maps, so we use the neighbors of $\mathbf{y}(n) : \mathbf{y}^{(1)}(n), \mathbf{y}^{(2)}, \dots, \mathbf{y}^{(NB)}(n)$ to define a spatial region. We then represent the maps as

$$\mathbf{y}^{(r)}(n+1) = \mathbf{F}(\mathbf{y}^{(r)}(n)) = \sum_{l=1}^M c(l)\phi_l(\mathbf{y}^{(r)}(n)), \quad (1.25)$$

where $\phi_i(\mathbf{x})$ is any convenient set of basis functions in the state space. We typically use polynomials, but radial basis functions are also a useful choice. This representation is essentially a local Taylor series expansion of the unknown $\mathbf{F}(\mathbf{x})$ using our knowledge of how neighbors move into one another under action of the dynamics. The coefficients $c(l)$ are found by a least squares minimization of

$$\sum_{r=1}^{NB} |\mathbf{y}^{(r)}(n+1) - \sum_{l=1}^M c(l)\phi_l(\mathbf{y}^{(r)}(n))|^2. \quad (1.26)$$

We now use these **local** maps to test the quality of local predictions as a function of NB and $d \leq d_E$ until the predictions are independent of these two quantities.




FIGURE 1.7. Global false nearest neighbors for data from an isolated LP neuron.




FIGURE 1.8. Global false nearest neighbors for data from our model of an isolated LP neuron.




FIGURE 1.9. Three-dimensional representation $[s(n), s(n+T), s(n+2T)]$ for data from our nonlinear electrical circuit.

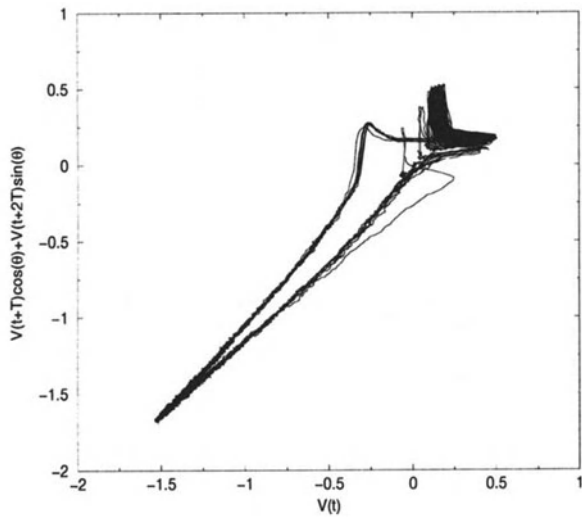


FIGURE 1.10. Three dimensional representation $[s(n), s(n + T), s(n + 2T)]$ for data from our observations on an isolated LP neuron.

In Figure 1.12 we show the percentage of bad predictions as a function of d and for $NB = 40, 60, 80,$ and 100 for data from our three degree of




FIGURE 1.11. Three dimensional representation $[s(n), s(n + T), s(n + 2T)]$ for output of our model for an isolated LP neuron.




FIGURE 1.12. Local false nearest neighbors for data from the three degree of freedom nonlinear electrical circuit. $d_L = 3$.

freedom nonlinear circuit. $d_L = 3$ is clearly selected. In Figure 1.13 we do the same for data from the isolated LP neuron, and in Figure 1.14 we do this for data from our Calcium dynamics model for this neuron. In the latter two cases $d_L = 3$ is also selected, although one could live with $d_L = 4$




FIGURE 1.13. Local false nearest neighbors for observed data from an isolated LP neuron. $d_L = 3$.




FIGURE 1.14. Local false nearest neighbors for membrane voltage output for our calcium dynamics model of an isolated LP neuron. $d_L = 3$.

from the experimental data.

1.8 Lyapunov Exponents; Predictability

The key estimate in determining whether a measured signal is from a chaotic oscillator is to establish whether there is at least one positive global Lyapunov exponent for motions on its attractor. These exponents, and there will be d_L of them, are determined by the evolution under the dynamics of a small perturbation to an orbit $\mathbf{y}(1), \mathbf{y}(2), \dots, \mathbf{y}(N)$. When we perturb the orbit at time n $\mathbf{y}(n) \rightarrow \mathbf{y}(n) + \Delta(n)$ we wish to know if $\Delta(n)$ shrinks or grows in time. If it grows as

$$\Delta(n + L) \approx \Delta(n)e^{L\lambda}, \tag{1.27}$$

then the signal source is chaotic, and the allowed λ are the Lyapunov exponents.

$\Delta(n)$ satisfies the linearized evolution

$$\Delta(n + 1) = \frac{\partial \mathbf{F}(\mathbf{x})}{\partial \mathbf{x}} \Big|_{\mathbf{x}=\mathbf{y}(n)} \cdot \Delta(n) = \mathbf{DF}(\mathbf{y}(n)) \cdot \Delta(n), \tag{1.28}$$

when $\mathbf{y}(n + 1) = \mathbf{f}(\mathbf{y}(n))$. $\mathbf{DF}(\mathbf{x})$ is a $d_L \times d_L$ matrix. After L steps, $\Delta(n + L)$ is determined by the composition of matrices $DF^L(\mathbf{y}(n)) = \mathbf{DF}(\mathbf{y}(n + L - 1)) \cdot \mathbf{DF}(\mathbf{y}(n + L - 2)) \cdots \mathbf{DF}(\mathbf{y}(n))$. The eigenvalues of this composite matrix give us the d_L Lyapunov exponents $\lambda_1, \lambda_2, \dots, \lambda_{d_L}$. To evaluate these quantities it is convenient to form the orthogonal matrix due to Oseledec [5]

$$\mathbf{OSL}(\mathbf{y}(n)) = \lim_{L \rightarrow \infty} \left[\mathbf{DF}^L(\mathbf{y}(n))^T \cdot \mathbf{DF}^L(\mathbf{y}(n)) \right]^{\frac{1}{2L}}. \tag{1.29}$$

The eigenvalues we want are $e^{\lambda_1}, e^{\lambda_2}, \dots, e^{\lambda_{d_L}}$.

Oseledec proves that this limit exists, is independent of $\mathbf{y}(n)$ within the basin of attraction for the orbit, and is independent of the coordinate system in which it is evaluated when those coordinate systems are connected by smooth transformations. The last property means that we can evaluate these exponents in our time delay proxy space as it represents a smooth transformation of the original (unknown) state space. One only needs to be careful that the correct number d_L of exponents are evaluated, not $d_E \geq d_L$. So a subspace of the d_E -dimensional space where the attractor is unfolded needs to be identified. Using the local false neighbors statistic does this in an efficient practical sense.

If any of the λ_a ; $a = 1, 2, \dots, d_L$ are positive, we have a chaotic system. If any of the λ_a is zero, there is a high likelihood we have differential equations describing the dynamics. The latter is because in a flow, dynamics

described by a differential equation, one perturbation can always be displacement along the orbit itself. This is just a phase change and does not grow in time; instead it follows the original orbit but is displaced in time. Because of dissipation,

$$\sum_{a=1}^{d_L} \lambda_a < 0. \quad (1.30)$$

If we know the λ_a , we have several pieces of useful information:

- because the perturbations on an orbit behave as

$$\Delta(n + L) \approx e^{L\tau_s \lambda_1} \Delta(n), \quad (1.31)$$

where λ_1 is the largest exponent, this places a limit on the time over which any effective predictions can be made. When Δ reaches the size of the attractor R_A , prediction is definitely out of the question, so at times of order $T_{\text{PRED}} \approx \log(\frac{R_A}{|\Delta(0)|})$ we lose predictability. $|\Delta(0)|$ is an estimate of the size of an initial perturbation or error due to noise or whatever disturbs the system.

- we can define a dimension, the so-called Lyapunov dimension associated with this spectrum of exponents. If

$$\sum_{k=1}^K \lambda_k > 0; \quad \sum_{k=1}^{K+1} \lambda_k < 0, \quad (1.32)$$

then the Lyapunov dimension is defined as

$$D_L = K + \frac{\sum_{k=1}^K \lambda_k}{|\lambda_{K+1}|}, \quad (1.33)$$

gives an estimate of the dimension of the attractor.

- the λ_a are invariants associated with the dynamical system producing the measured signal. They can be used to classify the system.

To numerically determine the λ_a we need accurate local estimates of the Jacobian matrices $\mathbf{DF}(\mathbf{x})$ at various locations on the attractor, then we require a systematic and accurate method for determining the eigenvalues of the product of Jacobians entering Oseledec's formula. The matrix involved is very ill conditioned, and we use a recursive QR decomposition technique pioneered by Eckmann and Ruelle [6].

To estimate the Jacobians, we once again make local maps on the attractor, namely $d_L \times d_L$ maps

$$\mathbf{x} \rightarrow \mathbf{F}(\mathbf{x}) = \sum_{m=1}^M c(m) \phi_m(\mathbf{x}), \quad (1.34)$$




FIGURE 1.15. Local Lyapunov exponents for voltage data from our three-degree-of-freedom electrical circuit. The local exponents are evaluated for L steps along the orbit following a perturbation. For $L \rightarrow \infty$ we arrive at the global exponents; here our maximum L is 2^{10} which is adequate.

evaluate the coefficients by a least squares minimization over NB neighbors, and then our estimate for $\mathbf{DF}(\mathbf{x})$ is

$$\mathbf{DF}(\mathbf{x}) = \sum_{m=1}^M c(m) \frac{\partial \phi_m(\mathbf{x})}{\partial \mathbf{x}}. \quad (1.35)$$

In our evaluations we use local polynomials of second or third order, and then the linear term gives us the Jacobian.

In Figure 1.15 we show the Lyapunov exponents for the three-degree-of-freedom electrical circuit. Figures 1.16 and 1.17 show the same quantities for the observed isolated LP neuron and for our membrane voltage output from our calcium dynamics model for the LP neuron.

1.9 “Black Box” Prediction in State Space

In the case of the LP neuron, we have a model which does remarkably well in capturing both the time series and the nonlinear statistical quantities we have evaluated. This means we could use this model for prediction of the membrane voltage time course of a neuron, assuming we could know all the state variables at some time. This is likely to be unavailable for this system, as we cannot accurately measure those variables.

In the case of the low-dimensional electrical circuit, we have the dynamical equations, or an excellent approximation to them, but even then to predict, for example, the circuit voltages, we require knowledge of the state variables at some time. It is likely that we will not have that information.

Fortunately, from the data or from data generated by a model of our choosing which has been tested by comparison with the nonlinear statistics from observations, we can construct “black box” models which allow prediction within the horizon set by the largest Lyapunov exponent λ_1 .

To make this kind of model, we once again rely on the information we have in phase space to help us see forward in time. The idea is that if we know how a whole neighborhood of observed points near $\mathbf{y}(n)$ move forward to a whole neighborhood near $\mathbf{y}(n+1)$, we can model that region in the usual way

$$\mathbf{y}(n) \rightarrow \mathbf{y}(n+1) = \mathbf{F}(\mathbf{y}(n)) = \sum_{m=1}^M c(m)\phi(\mathbf{y}(n)). \quad (1.36)$$

Using NB neighbors of the points $\mathbf{y}(n)$, we can determine the coefficients $c(m)$ **locally** near $\mathbf{y}(n)$. This we can do with each neighborhood on the attractor, and we can then create what amounts to a huge lookup table.

Prediction proceeds in the following fashion: Suppose from another sample of the data, a new point is given and it results in the data vector \mathbf{z}_0 in d_E -dimensional space. We take the original data set—the one on which we learned the local maps; we call this the training set—and we look for the




FIGURE 1.16. Local Lyapunov exponents for membrane voltage data from an isolated LP neuron. The local exponents are evaluated for L steps along the orbit following a perturbation. For $L \rightarrow \infty$ we arrive at the global exponents; here our maximum L is 2^{10} which is adequate.




FIGURE 1.17. Local Lyapunov exponents for membrane voltage output from our calcium dynamics model of an isolated LP neuron. The local exponents are evaluated for L steps along the orbit following a perturbation. For $L \rightarrow \infty$ we arrive at the global exponents; here our maximum L is 2^{10} which is adequate.

nearest neighbor of \mathbf{z}_0 in that set; call it $\mathbf{y}(Q)$. Now take the **local** map associated with the neighborhood of $\mathbf{y}(Q)$ and use that map to predict where \mathbf{z}_0 will go; call that \mathbf{z}_1 :

$$\mathbf{z}_1 = \sum_{m=1}^M c_Q(m) \phi_m(\mathbf{z}_0). \tag{1.37}$$

This moves us one step in time (one sampling time τ_s) along the orbit starting from \mathbf{z}_0 . To proceed to the next step, determine the nearest neighbor of \mathbf{z}_1 in the learning set, call it $\mathbf{y}(P)$, and we determine \mathbf{z}_2 to which \mathbf{z}_1 goes as

$$\mathbf{z}_2 = \sum_{m=1}^M c_P(m) \phi_m(\mathbf{z}_1). \tag{1.38}$$

We can now construct an orbit $\mathbf{z}_0 \rightarrow \mathbf{z}_1 \rightarrow \mathbf{z}_2 \rightarrow \dots$ which should give us good predictions up to the prediction horizon—on average over the attractor.

Figure 1.18 shows this prediction process in action. We took 20,000 data points from the nonlinear circuit, and with those made local maps of dimension $d_L = 3$ in $d_E = 3$. Then we used these models to predict L steps ahead from point 31000 through point 36000. The predictions were then compared to the actual values L steps ahead of each starting location for these




FIGURE 1.18. Prediction for the nonlinear electrical circuit; 20,000 points were used to ‘train’ the maps to capture the dynamics of the circuit in reconstructed state space. Predictions were made five steps in τ_s ahead using these maps. We show the predictions, the actual values of the circuit voltage and the errors.

5000 locations on the attractor. Shown in Figure 1.18 are the actual, the predicted, and the error of predictions for $L = 5\tau_s$. Figures 1.19 and 1.20 show this for $L = 12\tau_s$ and $L = 25\tau_s$ respectively. Clearly the accuracy of prediction degrades, as it must, for larger L , yet one can clearly see that the method works quite well for most regions of phase space. Where the *local* Lyapunov exponents are large, the prediction horizon is, naturally, smaller, and in those regions the method will not work as well.

1.10 Summary and Challenges

The “worked example” of analysis of membrane potential activity of an isolated neuron from a small assembly is instructive both for what it does and what it does not do. What it has done is provide a framework within which one can decide on the number of degrees of freedom required to describe a signal source and some quantitative statistical quantities with which that comparison can be made. What is not done here is to provide a “complete” set of comparison statistics nor do we suggest how to make models of the appropriate dimension to capture the dynamics seen. These latter two items are quite important and have no clear answer known to me. The former, a larger set of statistical quantities for exploring nonlinear




FIGURE 1.19. Prediction for the nonlinear electrical circuit; 20,000 points were used to ‘train’ the maps to capture the dynamics of the circuit in reconstructed state space. Predictions were made twelve steps in τ_s ahead using these maps. We show the predictions, the actual values of the circuit voltage and the errors.

signal sources, is amenable to existing technology and approaches; the latter is much more elusive. As an example of the latter, let me note that the complicated equations for the Hodgkin-Huxley models presented here can surely be replaced by three or four simplified equations of motion. We have done this in our work at UCSD where we used equations from Hindmarsh and Rose’s work in 1983 arising from a fit of the current-voltage characteristics of a large class of neurons without regard for the details of their ion channels. This leads to three-dimensional equations for most processes and four dimensions, if one wishes to include slow calcium dynamics as well.

This underlines the critical role of experiment and observation in working with models of nonlinear systems. Mathematics and theory are useful, but in my opinion somewhat sterile, when pursued in the absence of experiment—at least in this arena. The surprises we see which make this such a delightful area to work in come from how physical and biological systems have solved problems using capabilities of nonlinear systems we have just begun to uncover.




FIGURE 1.20. Prediction for the nonlinear electrical circuit. 20,000 points were used to ‘train’ the maps to capture the dynamics of the circuit in reconstructed state space. Predictions were made 25 steps in τ_s ahead using these maps. We show the predictions, the actual values of the circuit voltage and the errors.

Acknowledgments

The results and ideas discussed here were presented as part of a workshop held at the Isaac Newton Institute of Cambridge University, Cambridge, England in September, 1998. I would like to thank the organizers, Alistair Mees, Richard Smith, and Bill Fitzgerald for inviting me and putting together such an excellent extended discussion session. The hospitality of the Institute Director, Keith Moffatt was impeccable; thank you. Over the past several years these ideas have been developed in collaboration with U. Lall, M. Kennel, L. Tsimring, M. Rabinovich, T. Frison, N. Rulkov, R. Elson (a former Cambridge undergraduate), A. Selverston, M. Sushchik, R. Huerta, and R. Brown. This work was supported in part by U.S. Department of Energy grant DE-FG03-90ER14138 and in part by NSF grant NCR-9612250. The author received partial support from the CIA/Office of Research and Development through project No. 98-F135000-000, and from the Australian Research Council.

References

- [1] Shepherd, Gordon M., *Neurobiology*, 3rd Edition, Oxford University Press, 1994.
- [2] M. Falcke, R. Huerta, M.I. Rabinovich, H.D.I. Abarbanel, R. Elson, A. Selverston. 1999. Modeling Observed Chaotic Oscillations in Bursting Neurons: The Role of Calcium Dynamics and IP_3 . Submitted to *Biological Cybernetics*, February, 1999.
- [3] H.D.I. Abarbanel, *Analysis of Observed Chaotic Data* (Springer-Verlag), 1996.
- [4] Fraser, A. M. and Swinney, H. L., *Phys. Rev.*, **33A**, 1134 (1986); Fraser, A. M., *IEEE Trans. on Info. Theory*, **35**, 245 (1989); Fraser, A. M., *Physica*, **34D**, 391 (1989).
- [5] Oseledec, V. I., "A Multiplicative Ergodic Theorem. Lyapunov Characteristic Numbers for Dynamical Systems" *Trudy Mosk. Mat. Obsc* **19**, 197 (1968); *Moscow Math. Soc.* **19**, 197 (1968).
- [6] Eckmann, J.-P. and D. Ruelle "Ergodic Theory of Chaos and Strange Attractors" *Rev. Mod. Phys.* **57**, 617 (1985).

Chapter 2

Disentangling Uncertainty and Error: On the Predictability of Nonlinear Systems

Leonard A. Smith

ABSTRACT

Chaos places no a priori restrictions on predictability: Any uncertainty in the initial condition can be evolved and then quantified as a function of forecast time. If a specified accuracy at a given future time is desired, a perfect model can specify the initial accuracy required to obtain it, and accountable ensemble forecasts can be obtained for each unknown initial condition. Statistics which reflect the global properties of infinitesimals, such as Lyapunov exponents which define "chaos", limit predictability only in the simplest mathematical examples. Model error, on the other hand, makes forecasting a dubious endeavor. Forecasting with uncertain initial conditions in the perfect model scenario is contrasted with the case where a perfect model is unavailable, perhaps nonexistent. Applications to both low (2 to 400) dimensional models and high (10^7) dimensional models are discussed. For real physical systems no perfect model exists; the limitations of near-perfect models are considered, as is the relevance of the recurrence time of the system in terms of the likely duration of observations. It is argued that in the absence of a perfect model, a perfect ensemble does not exist and hence no accountable forecast scheme exists: Accurate probabilistic forecasts cannot be made even when the statistics of the observational uncertainty are known exactly. Nevertheless, ensemble forecasts are required when initial conditions are uncertain; returning to single best guess forecasts is not an option. Both the relevance of these observations to operational forecasts and alternatives to aiming for exact probabilistic forecasts are discussed.

2.1 Introduction

All my means are sane, my motive and my object mad.
Captain Ahab [42]

This chapter discusses the limits that uncertainty in the initial condition and error in the model place on both individual forecasts and predictability

in general. The systems of interest will be nonlinear, potentially chaotic. The methods of analysis and means of computation are sane and may be assumed exact without altering the limits discussed herein. The issue is whether our questions are well posed: Is the object of our search unobtainable even in the best of circumstances?

It has long been known (see, for example, Brillouin [12]) that even in a well-understood and accurately examined physical system, the combination of *observational uncertainty* and *model error* places severe limits on what we can say about the future of the system. While the following remarks hold for systems as simple as an analog circuit, they will be interpreted in the jargon of weather forecasting, even though the Earth's atmosphere/ocean system is not particularly well observed, nor are current models near-perfect. Nevertheless, numerical weather prediction (NWP) is an appropriate choice because, due to its economic importance, operational forecasts must be made every day and a great deal of thought has gone into attempting to improve the forecasts using any means available. Unlike the armchair forecasts of nonlinear dynamics or theoretical economics, operational weather forecasters must face their failures daily. This led Thompson [63] to contrast the relative contributions of uncertainty in the initial condition and model error in the 1950s. In 1965, variations in the reliability of individual forecasts led Lorenz [38] to suggest one (now operational) approach to quantifying the likely impact of uncertainty in initial condition on each particular forecast. Shortly thereafter, Epstein [16] and Leith [32] investigated both computational and analytic limits to maintaining initial uncertainty throughout a forecast. Many issues of current interest to nonlinear dynamicists are old chestnuts of the weather forecasting community.

For many years now, operational centers have made ensemble forecasts: a collection of initial conditions, each consistent with the observational uncertainty, are integrated forward in time. The role of uncertainty is introduced in Section 2.2. In Section 2.3, ensemble forecasting is explored within the perfect model scenario, and some jargon normalization is provided. The ensemble approach to forecasting deterministic systems replaces the single "best guess" initial condition of the traditional approach with a relatively small ensemble of different initial conditions, each member of the ensemble being consistent with the observational uncertainty in the initial state of the system. The idea here is that any initial uncertainty in the initial condition is reflected in the evolution of the ensemble, which in turn reflects the importance of that uncertainty in today's forecast. By observing how quickly the ensemble spreads out (or shrinks), one obtains a local estimate of the stability of forecasts made in this region of the system's state space; global measures like Lyapunov exponents are useless here [59, 57] except in the most simple, uniform systems. Even localized Lyapunov exponents [38, 3, 67] are misleading [70, 60], because they are based on the linearized dynamics over a pre-defined period of time, while the ensemble members may well sample the relevant nonlinearities and indicate when it

is that they appear. Indeed, chaos places no a priori limits on predictability: Given a perfect model, ensembles can accurately reflect the likelihood of observing various future conditions (i.e., provide a series of accountable probability forecasts). Such ensembles will slowly evolve toward the invariant measure of the system; but the time scale on which this happens is independent of the measures used to define chaos which are, in turn, based on the statistics of infinitesimals. Because there is always uncertainty in the initial condition, all nonlinear forecasts should be ensemble forecasts, and the issues discussed herein should find application to low-dimensional chaotic systems and high-dimensional weather forecasts.

The stated aim of ensemble forecasts ranges from estimating the ideal forecast probability density function (PDF) to simply obtaining a rough guide to the reliability of today's "best guess" or the control forecast. While the second aim remains in sight, the first cannot be fully realized. A major conclusion of this chapter is that just as uncertainty in the initial condition severely limits the utility of a single forecast even in a perfect model, so model error severely limits attempts to obtain "the" forecast PDF. This clarifies the limited applicability of results drawn from within the perfect model scenario. All models are wrong but some are more useful than others. If imperfect models are judged by a standard they cannot achieve, then the more useful models may be discarded. A similar situation holds when judging between single forecast models by using forecast error: Even a perfect model of a chaotic system will have a larger forecast error than a model which predicts the observed mean, at least in the far future. Predicting the mean may be desirable, if one really wants to minimize single forecast error, but this approach is obviously a poor guide to improving the physics of the model.

A basic difficulty in evaluating ensemble forecasts comes from the fact that the ensemble forecast estimates a probability density function in state space, while the verification (the true state of the system at the forecast time) is a point in state space¹. It is not possible to verify a single probability forecast, and each forecast involves a different initial condition. Further, no two initial conditions will ever be close in a dynamical system where the time required for the system to return to a point near the current state (i.e., the recurrence time) is longer than the likely duration of observations; thus the details of each PDF will differ for each forecast. The evaluation of a series of probability forecasts, given that each forecast PDF is different and that only a single realization of each forecast exists, is discussed in Section 2.4, where the one-dimensional method due to Talagrand is generalized to higher-dimensional spaces. But once it is accepted that

¹Worse still, there are at least three relevant spaces here: Forecasts lie in the model-state space, the system lies in the "true" state space, and observations explore yet another.

an accurate forecast PDF cannot be obtained even in near-perfect models, then new methods both of inter-model comparison and multi-model forecasts are called for; this may prove especially important in guiding model development. After a realistic look at the ambiguities introduced by model error in Section 2.5, two alternatives to computing a forecast PDF are introduced: (i) aiming for a bounding box and (ii) aiming for a ϕ -shadowing orbit. Each of these can be used to determine admissible predictability times. Fully embracing the limitations discussed herein suggests a new method for combining (rather than selecting the best of) imperfect models: The cross pollination in time (CPT) ensemble strategies introduced in Section 2.6 can outperform all the models available in terms of the two aims stated earlier. Standard multiple-model inter comparisons search for the best model in the same way that standard data assimilation routines search for the true state of the system; if no unique state can be identified empirically even under ideal conditions, then there is no “true” state, and each of these standard approaches may hinder the resulting forecast. This holds regardless of how sane and sophisticated the techniques used in the endeavor may be.

2.2 Uncertainty

Consider an intelligence which knew all the laws of nature precisely and had accurately (but not exactly) observed an isolated chaotic system for an arbitrarily long time. Such an agent — even if sufficiently vast to subject all this data to computationally exact analysis — could not determine the current state of the system, and thus the present, as well as the future, would remain uncertain. While our agent could not predict the future precisely, the future would hold no surprises: The predictability of the current “state” that could be seen [28, 56]. By forming an ensemble forecast from the plausible initial conditions consistent with both the system and the observations, could be estimated the probability density function (PDF) of future states to any desired accuracy. And these ensemble forecasts would be accountable: As the number of members in the ensemble grew, the accuracy of the PDF would improve proportionately. Further, for each particular initial state, the accuracy of observation required to allow a desired level of accuracy in the final state could be seen [53, 56, 57]. It is not only a perfect model, but also a perfect ensemble: a set of initial conditions both consistent with all observations and “on the attractor.” The true trajectory can be viewed as just another member of the distribution that is sampled to form the ensemble.

Operational forecasters at major weather centers in both Europe and North America, attempt an impersonation of this intelligence daily when they perform ensemble forecasts (see Palmer et al. [50], Toth and Kalnay [65],

and references therein). The predictability of the atmosphere varies from day to day, and a single “best guess” forecast is incomplete without a daily estimate of its likely accuracy. Ensemble forecasts aim to foresee variations in predictability by quantifying the time required for a given day’s ensemble members to splay out along significantly different trajectories, thereby quantifying the point at which that day’s “best guess” forecast is unlikely to be accurate. Ideally, one could also use the ensemble to quantify the probability of various events. But no physical model is perfect, and as we shall see, model error may make accountable probability forecasts unreachable, just as observational uncertainty makes a single forecast of little value. Our agent achieves an accountable forecast by evolving a perfect ensemble under a perfect model; once imperfect models are in use, no perfect ensemble exists. Accepting this forces us to change the interpretation and goals of forecasts. In fact, it calls into question what is meant by the state of a physical system.

Traditionally, the current state of a deterministic system is regarded as a point in state space, the exact location of which is obscured by observational uncertainty. This scenario only arises in computer experiments where we determine a trajectory and then pretend to forget where it was after adding some simulated observational noise. Even in that case, given only the noisy observations it would not be possible to identify a true state if we did not already know the answer: There would be a range of initial conditions, parameter values, and even distinct model structures which provided equally valid descriptions of the data. Clearly the traditional notion of “the state” of the system must be empirically suspect if even our idealized agent could not identify this “state” given a perfect model. In reality, of course, all models are wrong. It is our models which have states; there is no need for the hypothesis that physical systems have them.

2.3 The Perfect Model Scenario

What is the perfect model scenario? Let the role of the physical system be played by a set of equations proposed by Lorenz [37] as a parody of some atmospheric variable. As shown schematically in Figure 2.1, the system consists of m slow large-scale variables (the \tilde{x}_i) and $m \times n$ fast small-scale variables (the $y_{i,j}$) and thus has a state space dimension of $m(n + 1)$. The notation \tilde{x} is used to distinguish variables in the system state space from those in the model-state space, which will be denoted as x . Details can be found in Lorenz[37], Hansen [24], Orrell [48], Hansen and Smith [25] and the references therein. The equations are:

$$\frac{d\tilde{x}_i}{dt} = -\tilde{x}_{i-2}\tilde{x}_{i-1} + \tilde{x}_{i-1}\tilde{x}_{i+1} - \tilde{x}_i + F - \frac{h_{\tilde{x}}c}{b} \sum_{j=1}^n \tilde{y}_{j,i} \quad (2.1)$$

$$\frac{d\tilde{y}_{j,i}}{dt} = cb\tilde{y}_{j+1,i} (\tilde{y}_{j-1,i} - \tilde{y}_{j+2,i}) - c\tilde{y}_{j,i} + \frac{h_{\tilde{y}}c}{b} \tilde{x}_i. \quad (2.2)$$

where $i = 1, \dots, m$ and $j = 1, \dots, n$ and with cyclic boundary conditions on both the \tilde{x}_i and the $\tilde{y}_{j,i}$ (that is $\tilde{x}_{m+1} = \tilde{x}_1$, $\tilde{y}_{(n+1,i)} = \tilde{y}_{(1,i)}$ and so on). In the computations that follows $F = 10$, $m = 8$ and $n = 4$. The constants b and c are both equal to 10, so the small-scale dynamics are 10 times faster (and a factor of 10 smaller) than the large-scale dynamics, while the coupling coefficients $h_{\tilde{x}}$ and \tilde{y} are both set to unity.




FIGURE 2.1. Schematic of the Lorenz two-scale system.

Now I'll introduce some jargon.

When the forecast model is used to generate the observations which are to be forecast, one is in the *perfect model scenario*. The actual state of the system will be called *truth*, while our best estimate of that state, given only limited, noisy observations, is commonly referred to as the *analysis*. To test our model, the forecast is contrasted with the *verification*, which is in practice a future analysis; in (and only in) a perfect model experiment can the verification be truth itself. For a single set of simultaneous observations, the uncertainty in the analysis is related to the observational uncertainty. Given a time series of observations, the analysis corresponds to our best guess at the state, because this uses all the available observations (and a model), the analysis uncertainty in this case may be much lower than the observational uncertainty in the individual measurements. In the perfect model scenario, the analysis uncertainty is less than the observational uncertainty.

Operationally, an analysis may be generated via a four-dimensional variational assimilation (4DVAR) technique [62, 52]. 4DVAR attempts to locate the free running model trajectory which minimizes the difference between the model trajectory and the observations over a given duration (called the assimilation window), while allowing the observations to be spread out in both space (3-D fields) and time (+1-D). Achieving this in real-time with disparate data sources, each of which has different observational uncertainties and intermittently vanishes, is nontrivial. The search for a solution is also hampered by local minima in a 10^7 -dimensional space, but the key point here is that the resulting analysis can be much more accurate than the measurement uncertainty in a single set of simultaneous observations *as long as* the model is sufficiently accurate. We shall quantify “sufficiently accurate” later; here we note that this approach searches for “the” true state; this is somewhat troubling if we have accepted that there is no unique solution even within the perfect model scenario. An alternative approach to generating a best guess analysis and then creating ensemble members by adding perturbations is to generate an ensemble directly. This approach has been illustrated in simple low-dimensional models [28] and an operational method based on multiple analyses has been investigated by Houtekamer et al. [26]. Issues surrounding what makes the best analysis or the best operational ensemble are widely debated within the atmospheric community; many other options exist [7, 9, 22, 23, 43].

Traditionally, a weather forecast consisted of a single trajectory, started at the analysis and run at the highest available resolution. Such a traditional “best guess” is often run alongside an ensemble forecast, but because it is run at higher resolution, it lies in a different model-state space from that of the ensemble members. The *control* forecast is the ensemble member starting at the current analysis. Typically, roughly equal computational resources are invested in constructing the analysis and running the ensemble, with the high resolution run taking up most of the remainder ($\approx 10\%$). Open questions include the issue of whether additional computational resources should go to increasing the model resolution at which the ensemble members are run (i.e., obtaining a better PDF), running more ensemble members at the current resolution (i.e., a better approximation of an inferior PDF), or running the current system further into the future. Alternatively, resources could be directed to obtaining a better analysis. This could be approached through either a more computationally intensive assimilation technique or obtaining additional observations, the locations of which may change² daily [24, 25, 34].

²The idea being to take data at locations where the current level of uncertainty most hinders the forecast at some future time.

2.3.1 *Forecasting with a Perfect Model*

Figure 2.2 shows three ensemble forecasts in the perfect model scenario: A new ensemble is initiated every four time units (as denoted by the circle superimposed upon truth). Although the initial condition is not known exactly, we will assume a perfect model in this section; that is, the equations and parameter values are known exactly and the same integration scheme is used by the model and the system. We also assume that the system is chaotic, although this assumption is not necessary if we have only a finite duration of observations.

Brillouin [12] clearly showed how observational uncertainty limits our knowledge of both the current state and of the future; general arguments [28] establish that the current state is often not uniquely defined given uncertain observations over any duration. A simple way to see that this is true is to consider a special case of a chaotic dynamical system for which stable and unstable manifolds of the current state exist where the observational uncertainty is due only to quantization (i.e., truncation error). Clearly, there are portions of the unstable manifold within the current quantization box, which are also in the same series of previously observed boxes; that is, a set of trajectories which agree with all previous observations exactly, say, equal in the first three digits. This implies an infinity of states consistent with the observations. Thus no unique current state is defined by the observations, and therefore there can be no unique future state. Accountable forecasts must consider this infinity of states and attempt to maintain the initial uncertainty, quantifying its evolution during the forecast.

The forecast approach shown in Figure 2.2 will fail in this aim. The perfect model is used, and the initial conditions used are consistent with the uncertainty in the current observation. Because the model is perfect, the ensemble may contain trajectories which remain indistinguishable from the observations arbitrarily far into the future; such a model is said to ι -shadow the system [20, 56, 58]. Further the forecast PDF is a valid Monte Carlo approximation of the Fokker-Plank equations, given the observational uncertainty. In what way then is the forecast PDF incorrect?

When making ensemble forecasts we can estimate the probability of future events simply by counting the number of ensemble members in which the event occurs; for example, counting the number of ensemble members in which there are clear skies over Oxford for a 24-hour period of interest. By grouping together various forecasts (made on different days) which happen to have the same predicted probability, we can determine the relative frequency with which the event occurred on the days where the predicted probability was, say, about 10%. Ideally, this relative frequency should be near 0.10. To achieve this ideal requires a model capable of producing a realistic trajectory and an initial ensemble which gives the correct relative weight to physically relevant points consistent with the observational uncertainty. Of course, evaluating the accuracy of extremely low probability events, like the preceding example, may require extremely long data sets




FIGURE 2.2. Perfect model ensemble forecasts for the Lorenz system of Equations 2.1 and 2.2 showing an \tilde{x}_i component the true trajectory (solid) and of the forecast trajectories (dashed) from three perfect model ensembles. The members of each ensemble are consistent with the initial observational uncertainty. In this case the model and the system are identical but the values of the \tilde{x}_i are imperfectly known; for convenience, the true \tilde{y} values are used in each case. Every four steps an ensemble of initial conditions is forecast (each initiation is denoted by a circle). Visually, one can identify the time at which any one best guess forecast is likely to become unreliable. Yet one cannot obtain an accurate probability forecast from these ensembles, because the probability that an initial state is mistaken of the true state differs from the probability that that state is the true state.

to collect enough statistics to obtain a reliable estimate of the relative frequencies (see Murphy [44, 45] and references therein; Smith [57] provides a low dimensional dynamicist's point of view and examples).

Obviously a perfect model contains initial conditions consistent with the observational uncertainty which t -shadows for an arbitrarily long time. This is not the question, however. The difficulty lies in determining the subset of initial conditions which are physically relevant. Suppose that the system evolves on a manifold with dimension less than that of the system state space. Physically relevant points are restricted to the manifold, while the observational uncertainty will, in general, extend into the full state space: We are required to select only initial conditions from a set of zero measure on a manifold we do not know a priori. In this case the probability that a state \mathbf{y} cannot be distinguished from the true state is not equal to the probability that it is the true state; the true state will lie on the manifold, but \mathbf{y} need not, as illustrated in Figure 2.3. If we do not restrict our ensemble members to this manifold, then the predicted probabilities will not match the relative frequencies; this is nicely demonstrated by an example due to Gilmour [20], shown in Figure 2.4. The evolving probability distri-

butions in the left column (Figure 2.4 a) reflect ensembles consistent with the analysis uncertainty but *not* constrained to lie on the attractor, as in Figure 2.2. Contrast the unconstrained ensembles in the left column for $1 < t < 2$ and $2 < t < 3$, with the corresponding perfect ensembles in the right column; in each case the unconstrained forecast grows much too wide, due to including initial conditions which cannot be distinguished from the true state given the observations, but also cannot be the true state since they are inconsistent with the (unknown) long-term dynamics (i.e., they are not on the attractor). Interpreted in terms of the schematic in the right panel of Figure 2.3, the unconstrained ensembles choose members from the two-dimensional plane weighted by the isopleths, while the perfect ensemble only admits points on the attractor (the dots), again weighted by their relative likelihood given the isopleths of uncertainty. The unconstrained ensembles succeed in giving a general estimate of when the (unconstrained) analysis will become unreliable, but unlike the perfect ensemble these unconstrained ensembles cannot yield accountable probability forecasts.

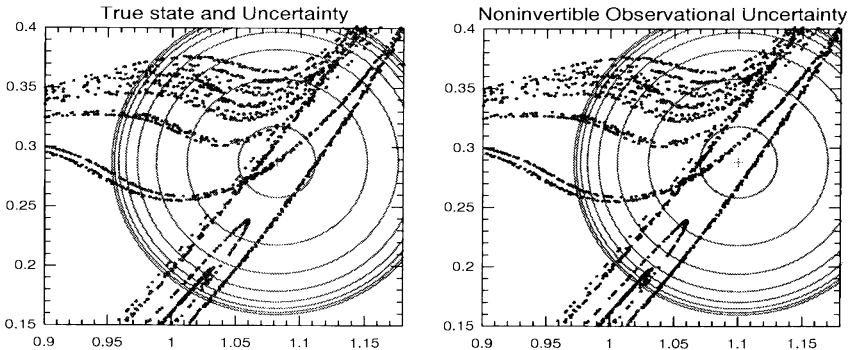


FIGURE 2.3. Left: Isopleths of the probability of an observation, given that the true state (+) is (1.0818408, 0.28764392), note that (+) lies on the attractor. Right: Isopleths of the probability that a state would give rise to the observation (+) given the observational uncertainty; but without knowing whether a point in state space is on the attractor, one cannot compute the probability of its being the true state. In this panel, the observation (+) is *not* on the attractor.

As indicated on the left panel of Figure 2.3, we can compute the probability of an observation \mathbf{x}_{obs} given the true state and the statistics of observational uncertainty (or the analysis uncertainty if any noise reduction is attempted). As indicated on the right panel, we can compute the probability that a given point \mathbf{x} could give rise to the observation \mathbf{x}_{obs} , but we **cannot** compute the probability that a given \mathbf{x} is the true state given *only* the observation and the noise process if the system evolves on a lower-dimensional manifold. To obtain the probability that \mathbf{x} is the true state

requires³ additional information (the manifold or the invariant measure in the case of a strange attractor). Without this additional information, the initial PDF will assign positive probability to points in state space which cannot correspond to the true state, and thus the initial PDF will be incorrect. If the initial PDF is wrong, then the final PDF is almost certainly wrong. We may be able to state approximately the probability of falling outside a region of state space, but we cannot obtain an accurate probability forecast. This again emphasizes that it is misleading to think of “uncertainty in the initial condition” in terms of a single well-defined state to which a random variable is added to yield the analysis. It is often better to think of “truth” as a random choice from the physical states consistent with the observations. To obtain a perfect ensemble (one with accurate predicted probabilities) one must choose ensemble members from the same distribution with the same relative weighting. In general, if the system is evolving on a lower-dimensional manifold (or attractor) this *cannot* be done, at least not without a perfect model and a huge computational effort.

We will return to that point in a moment; but first we stress that there is nothing “low dimensional” about this manifold: In the 10^7 -dimensional systems common in NWP, a $10^7 - 1$ -dimensional manifold counts as lower dimensional. Further, many practical forecasting systems (including NWP) are likely to fall into a Catch 22: If the system evolves on a lower-dimensional manifold, then obtaining perfect ensembles may prove intractable; but if the number of active degrees of freedom is equal to the dimension of the state space, then there is an insufficient number of observations to initialize the model in the first place. In practice, models can be initialized given the observations, so physical constraints implicit in the equations of the model must lower the effective number of degrees of freedom; but if the system evolves on a lower-dimensional manifold then

Of course, high-dimensional modeling (e.g., those with high spatial resolution) assumes that “the physics” restricts the effective number of active degrees of freedom. In practice, weather models tend to evolve the equations of motion of a fluid in a three-dimensional space (either in a grid point form, a spectral form, or both); given this restriction on model structure a high-resolution model may be required to obtain a good representation of a low-order manifold. Ideally we might use a (lower-order) model structure whose model-state space consisted only of the manifold, but currently our understanding of the physics is based on a spatial representation of atmospheric fields, and there is not yet sufficient data to construct the desired manifold empirically. We cannot formulate (much less integrate) the physical equations restricted to this manifold. Good low-order behavior may

³The extent to which this is relevant to NWP is discussed in Section 2.4.2, it is clearly relevant to forecasting systems whenever the projection of the attractor (or manifold) into the model-state space is lacunar on the length-scales defined by the observational uncertainty. Stephenson [61] notes implications this holds for quantifying analysis error.

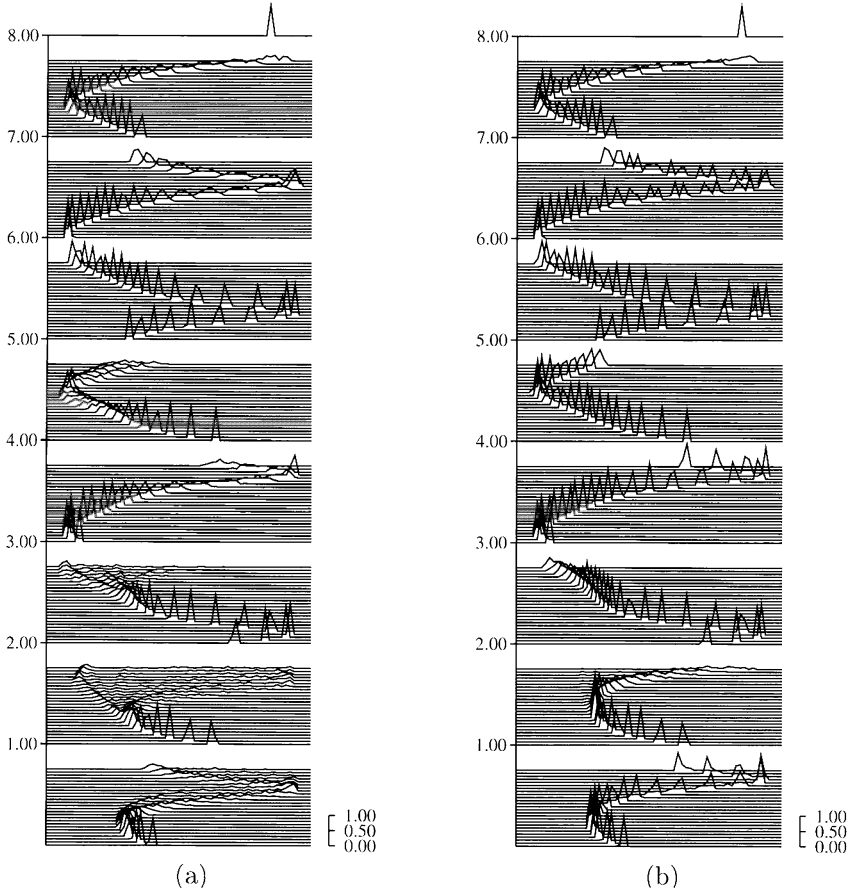


FIGURE 2.4. Comparison of (a) unconstrained ensembles and (b) perfect ensembles based on the same observations. Each 64 member ensemble is evolved under a perfect model of the Marzec Spiegel system [40] and projected onto $x \in [-1, 1.5]$. Time increases upwards. The gaps in the vertical indicate when new ensembles have been formed about the corresponding observation. Note that the distributions of the perfect ensembles just prior to the gaps tend to be tighter and more closely aligned with the distribution just after the gap (i.e., in closer agreement with the verifying observations). The distribution of the observational uncertainty was $\mathcal{U}_3(0.01)$. Figure from Gilmour [20] by permission.

require high-resolution models⁴.

There is something of a symmetry here between NWP forecasts and forecasting on strange attractors via delay embedding, as is common in nonlinear dynamics [54]. A major aim of the dynamic reconstruction in delay space is to model only the manifold or only the lowest-dimensional space within which the manifold can be embedded. But once in this low-dimensional space, there is no simple way to return to the physical state space of the system, that is, there is no method for interpreting model states in terms of physical variables other than those observed. Even though the dynamics of the reconstruction are diffeomorphic to dynamics in the full state space (on the attractor), an interpretation in terms of physical variables is much simpler in the full state space. In NWP the difficulty is in restricting the “physical variable” model to the manifold, while in delay reconstructions it lies in interpreting points on the manifold in terms of physical variables.

The existence of the right-hand column of Figure 2.4 indicates that perfect ensembles are not always unobtainable. Given a perfect model, the issue is one of computational expense which is, in turn, determined by the resolution of the observations and the recurrence time of the system. To build a perfect ensemble, we simply wait for an analog. The relevant question is: How long must we evolve the model before we obtain two states which are indistinguishable given our observational uncertainties? The sixty-four-member perfect ensembles of Figure 2.4 were obtained by collecting analogs in this way [56]. For third-order chaotic systems, this is often computationally feasible; for the Earth’s atmosphere, however, a single return to within the current observational accuracy over a large area like the northern hemisphere has been estimated to require 10^{30} years [66]; this is significantly longer than the lifetime of the atmosphere (and likely to exceed that of the Universe, for that matter). The model *must* be perfect: An arbitrarily good weather model can have a horrid climate, and it is the climate (the attractor) we must sample to obtain good probability weather forecasts. Our agent can do this because it is a perfect model and has unlimited computational power. In Section 2.5 we note that if the model is imperfect, no perfect ensembles exist (almost certainly).

We close this section with an epistemological question. In a recurrent system, perfect ensembles can be constructed with an analog approach (assuming that the successive returns are completely decorrelated!); in a non-recurrent system, or a system whose recurrence time is long compared to its likely lifetime, what meaning can be given to an ensemble forecast? Taking uncertainty in the initial condition seriously also raises a practical question: if we hold “truth” to be a point in state space, then we are

⁴I am grateful to P. Young and A. Lorenc for persuasively arguing the merits of the low-order approach and of the high-resolution approach, respectively.

forecasting a probability distribution in state space which we must verify with a single point. How might we do this?

2.4 Ensemble Verification

For each initial condition, an ensemble of initial states is forecast but only a single state exists with which to verify the forecast⁵. How might we evaluate that ensemble forecast? An individual ensemble forecast cannot be verified, but the consistency of a series of ensemble forecasts can be verified. For forecasts of scalar quantities the standard approach is to use rank histograms [5, 6, 21] commonly referred to as *Talagrand diagrams*. Assume for the moment that we have a perfect ensemble: Our ensemble was chosen from the same distribution as “truth”; in this case nothing can distinguish “truth”, it is just another ensemble member. This fact may be exploited, for example, by counting the number of forecasts which are greater than “truth”. This is illustrated in Figure 2.5 which shows the evolution of some scalar quantity; time runs from left to right and we have adopted the meteorological technique of denoting the “true” trajectory as a straight horizontal line. Eight-member ensembles of model trajectories appear at regular intervals and diverge from “truth” at a rate that depends on the local nonlinear structure of the model. Given a perfect ensemble, the number of ensemble members above “truth”, N_{over} , should be uniformly distributed between 0 and N ; better still the variance of any one bin in such a histogram is easily estimated. In operational NWP, the first and last bins tend to be overpopulated: Truth falls outside of the ensemble too often.

For the imperfect ensembles in the left-hand column of Figure 2.4, which are consistent with the observational uncertainty but not constrained to lie on the attractor, the Talagrand diagrams are under-populated at the extremes; this is to be expected when the ensemble regularly contains initial conditions not on the attractor and which diverge rapidly. For the perfect ensembles in the right column of Figure 2.4, the Talagrand diagrams are consistent.

Note that the Talagrand diagram can only be used for scalar forecasts since it relies on the rank ordering of the forecast values. Attempting to combine diagrams of different forecast values (say the temperature in London, Berlin, and Paris; or the geopotential height at each grid point in some region of interest) is ill-advised unless the predictands are truly independent, an unlikely case. Given a perfect ensemble, these combined diagrams would still be flat asymptotically, but we could no longer compute the expected rate of convergence (i.e., the variance), and hence we could not

⁵One might treat the verification as a PDF consistent with the observational uncertainty and centered upon the analysis, but the results below are easily generalized to that case.




FIGURE 2.5. A schematic of ensemble evaluation in one dimension: count N_{over} , the number of forecasts greater than truth for each lead time. If perfect ensembles are used, then N_{over} should be uniformly distributed; in N_{exp} experiments, we expect the relative frequency of a particular value of N_{over} to have mean N_{exp}/N_{bins} and variance $N_{exp}(N_{bins} - 1)/N_{bins}^2$, where N_{bins} is just the number of members in the ensemble plus one.

determine whether diagrams based on a finite amount of data were consistent with those expected from perfect ensembles.

2.4.1 Minimum Spanning Trees

The essence of the one-dimensional approach can be generalized to high-dimensional spaces by using minimum spanning trees (MST) [4] to detect whether the ensemble members are simply additional draws from the distribution that generated “truth”. The idea is shown in Figure 2.6. Consider a finite set of points in any metric space. A *spanning tree* is a collection of line segments which connects all the points in a set with no closed loops. The minimal spanning tree is that spanning tree in which the sum of the lengths of the segments is smallest. The MST test then, is to take all N member subsets of the $N + 1$ points (the N ensemble members and the control). If “truth” and the ensemble members are drawn from the same distribution, then no computation can distinguish the spanning tree from which “truth” was omitted [68]: we simply count N_{over} , now the number of the N spanning trees where an ensemble member was omitted whose length is longer than that of the MST where “truth” was omitted. It is not possible to evaluate a single ensemble in this way, but given a collection of n ensemble forecasts, a wide variety of systematic errors in ensemble formation could be identified. Histograms of N_{over} should follow the same statistics as the histograms of the Talagrand diagram, with a relative frequency ap-

proaching $\frac{1}{N+1}$ for each of the $(N + 1)$ possible results $(0, 1, 2, \dots, N)$ and variance $\sigma^2 = \frac{1}{n} \frac{N}{(N+1)^2}$, as earlier. ⁶.



FIGURE 2.6. A minimal spanning tree from the combined set of eight ensemble members (dark dots) and the verification (light dot) which is also on the attractor (and in this experiment “truth”).

Four examples are shown in Figure 2.7. The upper-left panel shows an acceptably flat distribution when both the verification and ensemble members are chosen from the distribution in Figure 2.6. The upper-right distribution reflects that when the verification is randomly distributed within the frame of the figure, it is often too far from its nearest neighbor, leading to a small MST when it is omitted and thus an increasing histogram as shown. The lower-left panel shows the histogram which results when each verification

⁶Note that this is the variance in a given bin over many realizations, because the relative frequency in each bin is not independent (they must sum one), the variance of the different bins in a single realization will differ from this, particularly when only a small number of bins are used.




FIGURE 2.7. Each panel shows a histogram of N_{over} , the number of MSTs omitting one ensemble member which were longer than the tree omitting the verification. In every case, the ensemble members were taken from the distribution shown in Figure 2.6. Histograms reflect when the verification was taken from: the same distribution (top left panel), a uniform distribution in two-dimensions (top right), a uniform distribution in one-dimension (lower left), and with independently chosen x and y components, where the distribution of each component matched that shown in Figure 2.6 (lower right).

is taken from a line lying near the attractor; this graph is easy to reject but its shape is less easy to interpret: Again the verification is too often too far from its nearest neighbor, but on those occasions when an ensemble member is chosen from that part of the attractor near the line on which the verification must lie, then the MST length of the tree omitting the verification tends to fall in the middle range. Finally the lower-right panel shows the result when the variables are chosen independently, but each from the corresponding correct distribution: The x -component of the verification is

taken from a correct distribution of x values and the y -component of the verification is taken from the correct distribution of y values. In this case, both the Talagrand diagram for x and the Talagrand diagram for y would have been acceptable, but the MST test rejects it because the conditional distribution of x given y is incorrect.

2.4.2 Relevance to Operational Forecasting

Hopefully, the last few sections have made clear the difficulty of obtaining perfect ensembles, even given a perfect model. This is without a doubt a concern when forecasting low-dimensional systems described by strange attractors; if the perfect ensemble is lacunar and the operational ensemble is not, then accurate probability forecasts will not be obtained. But is this really an issue in operational weather forecasting? In operational forecasts where the system is evolving on a lower-dimensional manifold or attractor (not low, just lower) and the structure of the manifold is not isotropic on the length scales resolved by the uncertainty in the analysis, these issues are important. For example, let the true state lie on a line and the analysis uncertainty correspond to a uniform distribution on a disk the line intersects. In that case sampling the disk to form an ensemble consistent with the analysis uncertainty yields an ensemble very different from the perfect ensemble, which will only contain points from on both the disk and the line. Alternatively, if the manifold consisted of many parallel lines, effectively filling the plane on length-scale defined by the radius of the disk, then the unconstrained ensembles might prove similar to perfect ensembles, as long as they did not contain too many members. In general, these difficulties may prove less important in systems where the invariant measure is smooth and slowly varying in state space (or its projection into the model-state space is uniform), or where the manifold is so contorted on the scale of the observational uncertainty that it can be treated as uniform. There may also be cases where the resolution of the model is so coarse it makes the variations un-resolvable.

Of course, it is also possible that the model error is so large that the forecasts are very wrong before the effects come into play. But in the limits of accurate short-term prediction models and small uncertainty in the initial condition, these issues will prove relevant for both low-dimensional dynamical systems and high-dimensional weather models of NWP.

Meteorologists tend to distinguish forecasts made with large models (NWP) from those made using less complicated empirical models and personal insight. While the NWP models get the most press, the simpler methods are sometimes quite good. This is most often true on small spatial scales and short forecast times (hours) at locations for which there are long historical records [71], and on very long time scales (seasonal or greater) where the biases of NWP models may become evident [11, 51]. It would be interesting to contrast the performance of ensembles in these empirical models with those under NWP for, say, seasonal time scales.

2.5 Imperfect Model Scenarios

Only hypothetical agents are allowed perfect models; we must deal with realistic models. This fact alters the philosophy of nonlinear forecasting as fundamentally as the acceptance of uncertainty in the initial condition. To see this, we introduce a model for the two-level system of the previous section which will play a role analogous to that played by weather models in relation to the Earth's atmosphere/ocean system. Keeping equations 2.1 and 2.2 as the system, we will consider models of the form:

$$\frac{dx_i}{dt} = -x_{i-2}x_{i-1} + x_{i-1}x_{i+1} - x_i + P_i(\mathbf{x}, t), \quad i=1, m \quad (2.3)$$

These equations for the model variables \mathbf{x} are structurally similar to 2.1 and 2.2 which determined the large-scale $\bar{\mathbf{x}}$ dynamics of the system, they differ in that the dynamics of the small-scale fast variables, the $\bar{\mathbf{y}}$, have been parameterized by the function P . A wide range of parameterizations may be entertained; options we have explored for $P_i(\mathbf{x}, t)$ include:

$$P_i(\mathbf{x}, t) = \left\{ \begin{array}{ll} \alpha_0 & \text{constant} \\ \alpha_0 + \alpha_1 x_i & \text{linear} \\ \alpha_0 + \boldsymbol{\alpha} \cdot \mathbf{x} & \text{m-linear} \\ H_1(\mathbf{x}) & \text{nonlocal1} \\ H_2(\mathbf{x}, \frac{\Delta \mathbf{x}}{\Delta t}) & \text{nonlocal2} \\ \text{I.I.D}_{obs} & \text{IID} \\ \gamma_1 P_i(\mathbf{x}, t-1) + N(0, \gamma_0) & \text{AR(1)} \end{array} \right.$$

These parameterizations range from simple variations on linear models (a constant, a linear parameterization based on only the local variable x_i , a linearization based on all m components⁷ of \mathbf{x}) through nonlinear variations suggested by prediction studies in low-dimensional nonlinear dynamical systems [55, 17, 13, 27, 2, 1] (here H_1 is nonlinear and nonlocal in physical space, while H_2 is also nonlocal in time) and finally to simple stochastic parameterizations (either choosing a value for P_i at random from the observed historical forcing or fitting an autoregressive model to those observations and using that AR model).

One property each of these various parameterizations share is that they are wrong: Given that $\mathbf{x} \in R^m$ while $(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in R^{m(n+1)}$, there is, in general⁸,

⁷If the x_i s are interpreted as being distributed in physical space, then this last model is nonlocal in physical space because it requires input from other grid points; it is a serious complication given the computational structure of current weather models, but it may prove worth the difficulty of implementation as the spatial resolution of those models improves.

⁸Of course the inclusion of parameterizations H_1 and H_2 was motivated by our knowledge [54] that *if* the attractor is restricted to a manifold of dimension Q *and* the parameterization is evaluated only for states on the attractor, *then* perfect parameterizations of the form H_1 and H_2 (almost certainly) exist if $2Q < m$ or $Q < m$, respectively.

no perfect model with the form of Equations 2.3 and thus no perfect ensemble. Each model will have one distribution from which ensembles may be drawn which will verify at one day; and a different distribution yielding ensembles which will verify at two days, and so on. Even these distributions will vary from model to model. The forecast quality of each of these models will be discussed elsewhere; the point of introducing them here is to consider the question of what to do with them: Should one search for the best model? Consider an ensemble over models? Or something even more radical? And what is the aim of ensemble forecasting in this context?

We start with an easier question: What is the correct value for α_0 in the constant parameterization? An obvious choice is $\alpha_0 = \tilde{\alpha}$ where

$$\tilde{\alpha} \equiv \left\langle F - \frac{h_{\tilde{x}} c}{b} \sum_{j=1}^J \tilde{y}_{j,i} \right\rangle_{\tilde{x}}, \quad (2.4)$$

that is, the average value of the forcing term where the average is taken over the invariant measure of the true system. But in a nonlinear system, this value has no special claim to optimality; why not take the value which minimizes the one-step forecast error? Or that minimizes the two-step forecast error? Or that yields the longest mean ι -shadowing time? Or that best reproduces the invariant measure of the true attractor projected into the model-state space [41]? The model-state space differs from the state space, $\mathbf{x} \neq \tilde{\mathbf{x}}$ even if both variables are called x . Thus the correct method for determining the free parameters in the imperfect models depends on the goal of the forecaster. There need be no unique set of “true” parameter values; standing water need not at freeze at exactly zero degrees C in a good weather model.

In a perfect model, there is a unique perfect ensemble corresponding to all potential states of the system, each weighted by its probability given the observations. In an imperfect model no perfect ensembles exist, and it is doubtful whether a unique optimal ensemble is well-defined for the same reasons that optimal parameters are not. None of these models will ι -shadow indefinitely; as we get more data we will find (almost certainly) that the probability of the data given the model goes to zero; not just for these particular models, but for every model in the model class(es) under consideration. Although it is not clear if there is a natural definition of the best model, at least on time scales much less than the recurrence time, it does seem likely that an ensemble over models will outperform the best model for most reasonable definitions of “best.”

There is no simple stochastic fix. While adding a random component to a deterministic model may imply that a trajectory which stays near the verification exists, such a trajectory cannot be said to ι -shadow unless the random innovations required are consistent with the source of stochasticity specified by the model. For the AR parameterization, the innovations must be consistent (at some confidence level) with I.I.D. drawn from an a priori

specified Gaussian distribution. When model trajectories are restricted to remain on an unspecified manifold, the construction of stochastic terms which respect this constraint appears nontrivial. In practice, the stochastic models we have explored in this context are consistently over-dispersive in model-state space.




FIGURE 2.8. The trajectory of Figure 2.2, this time showing truth (solid) and three forecasts of four imperfect models: Linear (dark dotted), constant (light dotted), IID (dark dashed), and AR1 (light dashed). At each of the circles an ensemble of model trajectories is initiated using the exact values of \tilde{x} , one trajectory for each model.

Forecast ensembles over models are shown in Figure 2.8; here there is no observational uncertainty when the model is initialized: each initial condition is exact (i.e., the analysis corresponds to the true state projected into the model's state space). Because each of the individual models is wrong, the PDF will be incorrect; furthermore there need be no initial condition for any imperfect model which will ι -shadow for the duration of the forecast. Finally, no analog of the perfect ensemble exists because an arbitrarily good weather model can have a horrid climatology. No model ensemble scheme will verify accountably, nor can a level of accuracy can be determined a priori which will guarantee a bound on the uncertainty of the forecast at a fixed future time. The only option for an accountable ensemble is to wait for physical analogs, as suggested by Lorenz [35] in terms of a single forecast; for daily weather, that may take some time.

Of course, a forecast trajectory in the future need not be held to the same standards as analysis trajectory of the past. While an ι -shadowing trajectory must exist for coherent variational assimilation, forecast models unable to produce an ι -shadow trajectory may still be of great value. A less constrained notion of shadowing is required: ϕ -shadowing.

2.5.1 ϕ -Shadowing

A model ϕ -shadows for a time τ_ϕ if the model contains an initial condition, consistent with the initial observational uncertainty, which resembles the future closely enough for a forecaster: A ϕ -shadow need only be useful. This differs from ι -shadowing, for example, in that relatively large errors in the time of onset of a rain storm may be accepted (as long as a storm is forecast), or relatively small variations in the strength and location of spatially coherent structures may be acceptable (although these errors often result in huge contributions to the RMS error of a predicted field). The idea is to adopt an operationally relevant definition of what constitutes a useful forecast trajectory (see Murphy [44] for a discussion of what constitutes a good forecast); determining the median of the distribution of times $\tau_\phi(\bar{\mathbf{x}})$ over which a given model can ϕ -shadow provides a bound on predictability beyond which the use of Monte Carlo ensembles is, at best, questionable. It is difficult to conceive of a useful purpose for ensemble forecasts beyond the time horizon over which the model can ϕ -shadow: if there is no initial condition which will reflect the future even roughly, what can be gained from a distribution of such model-error-dominated trajectories? This only argues for the existence of a ϕ -shadow; ideally the higher the probability of finding them, the more useful the model.

Requiring a ϕ -shadowing trajectory to exist is a much looser constraint than requiring a model to ι -shadow. A model ι -shadows for a time τ_ι if it contains a trajectory which is consistent with the observational uncertainty at *all* times t , $0 \leq t \leq \tau_\iota$. The tolerance is set by the uncertainty in the observations⁹, and will be much more restrictive than just requiring a useful forecast [20, 56, 58]. If the observations are limited only by quantization uncertainty due to truncation, then to ι -shadow a trajectory must fall into the quantization box corresponding to the observation at each observation time. For Gaussian uncertainties, some confidence level α must be chosen. Experience indicates that shadowing time is not very sensitive to this choice; at least in low-dimensional systems, once things go wrong, they go badly wrong rather quickly¹⁰. Given two diffeomorphisms, Anosov [8] and Bowen [10] determined sufficient conditions to guarantee the existence of another type of shadowing, but this ϵ -shadowing contrasts the trajectories of two well-defined mathematical systems; it is based on assumptions which make it irrelevant (although nevertheless comforting) when contrasting im-

⁹In real applications, of course, the real measurements are rarely equivalent to the variables in the model-state space; there is an entire field of endeavor dedicated to relating point wise physical measurements to grid point model variables and the relation of both to the three-dimensional fields they are taken to represent. On which natural length scale can one coherently define wind or identify it with the velocity variables evolved within a model?

¹⁰Note, however, that the sci-fi models of Judd Small and Mees [27] in this volume address this problem explicitly.

perfect models and real data (see Gilmour [20] for additional discussion). In practice, what we require are more vague but still quantifiable shadows, more along the lines of Eddington’s use of the term [15]. Finding shadows is not so enlightening as realizing when they do not exist.

2.5.2 Bounding Boxes

Examining worst-case scenarios is another common goal in weather forecasting. One approach would be to use the ensemble to define a region of model-state space within which the future is likely to fall¹¹. The obvious method (see Figure 2.9) is to construct a convex hull from the ensemble members, and this is very useful in low-dimensional models. But inasmuch as it requires $m + 1$ points to define a convex hull in m dimensions, this approach is untenable with ensembles of about 10^2 members in a typical model where $m \approx 10^7$.

An alternative to the convex hull is the *bounding box*, which has the advantage that it requires only two points and a coordinate system, regardless of the dimension of the model-state space. Consider a model on a spatial grid; at each grid point (for each variable) take the maximum value over all ensemble members; this co-dimension one plane is one “side” of the box, while the plane corresponding to minimum value forms the side opposite. Repeating this for all variables defines a volume of state space. Figure 2.9 uses the ensemble of Figure 2.6 to show both the convex hull (solid boundary) and bounding box (dotted boundary) defined by this ensemble. Note that the verification (which is truth, in this case) falls just outside the bounding box, but because the ensemble members were drawn from the same distribution as truth, truth is no more likely to fall outside than any member of the ensemble. In cases where both are defined, the bounding box always has a greater volume than the convex hull, of course, and hence it provides a larger (i.e., easier) net with which to bag the verification. We have not found this advantage to result in over-confidence in the model when the bounding box test is applied to forecasts of real data.

In this scenario, it is straightforward to estimate the number of ensemble members required to have, say, a 95% chance that truth falls within the bounding box defined by the ensemble. This can be done analytically when the distributions are Gaussian as a function of standard deviation and bias [29]. Indeed, we plan to use this result to estimate bias of operational NWP ensembles. Of course, adopting the spatial grid carries the added bonus that so long as the verification is within the bounding box of the ensemble forecast, nothing unexpected can happen. The target here contains much less information than an accurate PDF forecast, but in the imperfect model scenario there is no accurate PDF to be had.

¹¹ More precisely, to define a region which will contain the verification at the α level of confidence; ideally α is 100% but finite ensembles make this unlikely even in the perfect model scenario.



FIGURE 2.9. The convex hull (solid boundary) and bounding box (dotted boundary) for the ensemble of Figure 2.6.

2.5.3 Applications to Climate

Bounding boxes may also be of use in climate modeling. Typically one does not expect to find a ϕ -shadow in the climate model context where the goal is to reproduce the general statistics of likely states rather than a particular forecast trajectory. Lorenz [36] refers to this goal of climate modeling as predictions of the second kind. Yet the atmospheres of climate models may still be quite sensitive to initial condition, even when forced by observed sea surface temperatures (SST). Further, climate models are often run in an ensemble mode over historical periods, say, an ensemble of fifty-year runs where each member is started from the analysis corresponding to a different day in 1950. While one should not expect to run large enough ensembles to produce even a ϕ -shadow over a fifty-year period, it is reasonable to ask how large an ensemble would be required so that the analysis (or reanaly-

sis [30]) for January 15, 1970 falls within the bounding box of all climate states of all ensemble members taken on all days between, say, January 1 and January 31, 1970. The point is that if reality (or even the analysis) consistently falls outside this bounding box, then the (dynamical) statistics of the climate model would be placed in doubt. Identifying specific historical periods where the model consistently fell outside the bounding box might aid in the identification of physical processes (active during those periods) which were insufficiently reproduced in the model. Over the historical record, one might hope for return of skill in the climate ensemble, inasmuch as each member is guided by the observed SSTs; in a free-running fully coupled model the minimum size of the ensemble required to obtain a bounding box would again be of interest in estimating the additional length of time (or number of ensemble members) that would have to be run to explore the additional degrees of freedom released.

Both the MST and the bounding box can be used to investigate natural variability of the climate system, either over time or in establishing whether the January 1 anomalies, variance adjusted, are the same in distribution as those of August 1. Given the short duration of many climate records, it is not uncommon to combine data from different seasons, once each data point has been adjusted to “remove seasonality.” Examining the relative frequency with which the data from one calendar day fall into the bounding box defined by data from a day some months later, or the MST equivalent, would provide a useful check on whether simply adjusting the mean and variance is sufficient.

2.6 Multimodel CPT Ensembles

In this section a new method of truly multi-model ensemble forecasting is presented which attempts to take the limitations discussed earlier seriously. If we accept that each of our models is incorrect, that the “correct initial PDF” is as ill-defined as the “true initial state”, then we can construct a multi-model forecasting scheme which will outperform any individual model in terms of both ϕ -shadowing and the duration for which the verification remains within the bounding box defined by the ensemble members, at least in the limit of huge ensembles.

The simplest reaction to having M models is to identify the best one, discard the others, and compute $M \times N$ member ensembles under this single “best” model. If the models are of comparable quality, then it is likely that different models will tend to do better in different regions of state space (i.e., on different days), due to variations in the particular processes that are important locally. In practice, there is rarely enough data to identify which will be the best on a given day, and a reasonable alternative is to compute M , N -member ensembles, one ensemble under

each model. Note that neither approach can produce a ϕ -shadow longer than the longest ϕ -shadow found within the individual models. If the M models really do have independent shortcomings (ideally, if they fail to ι -shadow in different regions of state space), then it is possible to cross pollinate trajectories between models in order to obtain truly multi-model trajectories that explore important regions of state space the individual models just can't reach. This cross pollination in time (CPT) approach can outperform both of the preceding methods.

The basic CPT approach first takes the M N -member ensemble forecasts made under each model and combines them to form one large set of $N \times M$ points in the model-state space. This large ensemble is then pruned back to N member states, attempting to maintain a large bounding box while deleting one member in each pair of relatively close ensemble members (the details of the PDF are wrong anyway). These N conditions are then propagated forward under each of the M models and so on.

Inasmuch as the CPT ensemble model implicitly contains all trajectories of each of its constituent models, CPT can ϕ -shadow as long as or longer than any of the individual models. Similarly, in the limit of large ensembles, or in the absence of pruning, a CPT bounding box will contain the bounding box of the best model; it is expected that given a good pruning scheme, the bounding box of the CPT ensemble will be more likely to contain the verification over a longer duration than those of an $M \times N$ member ensemble under the "best" model. While the optimal pruning scheme is still an object of research, the simple approach of taking the nearest pair of points and deleting the point with the smallest second nearest neighbor distance, has been found to work fairly well in some simple examples. Note that the aim of pruning is quite different from that of resampling from an estimated PDF [9].

This approach assumes that either all the models share the same model-state space, or the one-to-one maps that link their individual state spaces exist; neither needs to be the case in weather forecasting. And, of course, when parameterizations of physical processes are involved one must consider the time scale of pruning; for example, we would wish not to switch between parameterizations in the midst of growing a cloud. But at least the preliminary step of being able to run a collection of operational models on the same computer system has been achieved at the European Centre for Medium-Range Weather Forecasting (ECMWF). Switching between models is a nontrivial process, and some may find it objectionable in principle as there is no longer a "set of equations" being solved, although one might argue that the process is, in fact, solving a rather large iterated function system. Ideally, future research will resolve these issues while retaining a closed form for the model if not for the solutions. But it is interesting to note that physics of late has gotten along rather well without always building the mechanical model Lord Kelvin [33] held to be the prerequisite for understanding a hundred years ago. Perchance the twentieth century will

be remembered as the “Century of the Equation.”

2.7 Discussion

Chaotic systems are often thought to be unpredictable because they have the property of on-average exponential growth of infinitesimal uncertainties, at least when the (geometric) average is taken over a trajectory which explores the entire attractor. Yet this “exponential on-average” growth places no bounds on (i) the growth of a finite uncertainty, or (ii) predictability over any finite time horizon, or (iii) the average uncertainty doubling time. The dynamics of uncertainty are much richer than simple uniform exponential growth [47, 60]. Lyapunov exponents are only *effective* rates, nothing need actually grow like $\epsilon_0 e^{\Lambda t}$.

Of course simple mathematical models of chaos, designed with tractability in mind, tend to have fairly uniform growth rates (by construction); this can yield a very biased picture of predictability. In the Baker’s Map [49] for example, the fastest uncertainty doubling time of each initial condition is one iteration and the Lyapunov exponent is equal to one bit per iteration. Yet within the family of Baker’s Apprentice Maps [59], all of which have a Lyapunov exponent greater than one, there are maps with arbitrarily large average doubling times [57]. Even in the Lorenz 1963 model, there are regions within which all perturbations must shrink for a (finite) time [60, 46]. Vannitsem and Nicolis [67] investigate these inhomogeneities in an atmospheric model.

Only in systems where the dynamics linearized about a trajectory accurately reflect the true nonlinear dynamics at macroscopic scales of interest do Lyapunov exponents have any impact on predictability. As long as uncertainty stays infinitesimal, it cannot limit predictability, and once it is finite the Lyapunov exponents need not provide a reliable guide for uncertainty growth. Whether locally defined [3, 70] or global [14], Lyapunov exponents are only *effective* rates, and even when infinitesimal perturbations really do grow exponentially in time, the uncertainty growth may saturate at an amplitude much less than the diameter of the attractor. This is clear from the macroscopic structure visible in Figure 2.4. In general, there need not be a “Lyapunov horizon”.

Orrell contrasts forecast uncertainty growth due to chaos with that due to model error [48]. Assume for the moment that the model is perfect and that forecast error does grow as

$$\langle \epsilon(t) \rangle = \epsilon_0 e^{\Lambda t}. \quad (2.5)$$

In this case, the value of $\langle \epsilon(t) \rangle$ at any fixed t can be made small by taking a sufficiently small ϵ_0 . If the model is not perfect, however, there will be a difference between the velocity of the model trajectory in the model-state

space and that of the system trajectory (projected into the model-state space). This difference remains even when $\epsilon_0 = 0$ lead to an initial error growth which is linear in time and thus (initially) greater than $\epsilon_0 e^{\Lambda t}$. Thus for an imperfect model

$$\langle \epsilon(t) \rangle = \bar{v}t \quad (2.6)$$

where the value \bar{v} is the magnitude of this velocity difference averaged over the projection of the system's invariant measure into the model-state space¹². For an imperfect model, $\bar{v} > 0$ and therefore the forecast uncertainty due to model error will always dominate the forecast uncertainty due to chaos for sufficiently small ϵ_0 . While “chaos” can make the error growth greater still, as $\epsilon_0 \rightarrow 0$, model error will dominate. Worse, it is not clear how to correct this with ensemble forecasts.

Given a perfect model, one might construct a perfect ensemble; but even if the model structure is correct and only the values of model parameters are uncertain, accurate PDF forecasts seem beyond reach. One may sample the parameter space in a sensible way and construct a perfect ensemble for each realization, but the resulting ensemble PDF will not accurately reflect the likelihood of finding the properties of the future trajectory which will be observed. It is not obvious how to construct ensembles over a model class.

Is the model class of deterministic systems too small? Perhaps [64], but it is not clear how to best introduce stochastic dynamics in structures where a strong deterministic nonlinear component is easily extracted; this is particularly the case when the deterministic dynamics are known to lie on a lower-dimensional manifold, the details of which are not known. Other contributions in this volume [18, 19, 69] suggest a number of avenues. The operative question is how to best model the phenomenon: the issue of whether a real system “really is” deterministic or stochastic cannot be resolved from real data [39, 57]. Determining how to best model a phenomenon turns on the issue of how we decide to evaluate our models. This paper is intended to stimulate debate on sane methods of model evaluation; there may be no best.

Is an accountable probability forecast a viable goal? Perhaps not. The rank histogram evaluation techniques of Section 2.4 assume that “truth” is indistinguishable from the members of the ensemble and that all are drawn from the same distribution. This is never the case in practice, where we begin with uncertainty over the initial condition, the boundary condition, the parameter, and even the model structure. The verification is never “just another member” drawn from this distribution. Perhaps the one thing we are certain of is that our model class is incorrect: The very structure of our models will change with additional observations. This eventuality need

¹²Orrell [48] illustrates this relationship in the Lorenz 1996 system, deriving the variation in \bar{v} as a function of the parameter F in equation 2.1.

not stop us from decreasing our uncertainty and refining our probability forecasts, but it will prevent our forecast PDF from producing flat rank histograms.

2.8 Summary

Chaos poses no difficulties for LaPlace's demon [31], whose abilities were such that given one exact snapshot of a dynamical system, a perfect forecast of the future can be calculated. Such a forecast is beyond the powers of a modern incarnation with the same abilities but without access to exact observations; even given imperfect measurements which stretch back into the distant past, our agent cannot determine the current state of the system from among a set of indistinguishable states. It can, however, foresee the probability of any eventuality. For mortals with imperfect models, even the foresight of exact relative probabilities is lost; we must expect to be surprised, occasionally, as there will be events which cannot even be foreshadowed.

As has long been recognized, uncertainty in the initial condition limits the utility of single deterministic forecasts of nonlinear systems like the Earth's atmosphere. If this uncertainty is accepted, then internal consistency requires that an ensemble of initial conditions, each consistent with the observations, be evolved forward under the model. Methods for selecting these initial conditions [38] were advanced by Lorenz in 1965 and competing operational approaches dating back to the early 90's are used in European and American weather forecasting centers. Assuming that the model physics is perfect, these methods aim at a weighted selection of the perfect ensemble [56, 57], where the weighting scheme depends on the aim of the forecaster.

Even under ideal conditions, uncertainty in the initial condition also limits the utility of single deterministic predictions of deterministic nonlinear systems; in practice ensembles of initial conditions are forecast with the dual aims of (1) estimating the reliability of that forecast and (2) estimating some aspects of the probability density function (PDF). Current rank histogram verification techniques are limited to scalar forecasts; Section 2.4.1 introduced a method using minimum spanning trees to allow computationally efficient verification in higher-dimensional spaces, including the 10^7 -dimensional weather model forecasts. Given a perfect model, one may construct an accountable ensemble forecast system by sampling from a perfect ensemble; this scheme can yield accurate probability density estimates. In general, no perfect ensembles exist for imperfect models. If accountable estimates of the forecast PDF are unobtainable, we should question whether current skill scores provide a reliable guide for model improvement.

Ensemble prediction systems can consider Monte Carlo ensembles over initial conditions, parameters and model structure. If accurate probability

forecasts prove untenable, what viable aims exist? Two were discussed in Section 2.5: obtaining at least one good forecast trajectory (a ϕ -shadow) and constructing an ensemble whose bounding box is likely to contain the verification. In the long term, the bounding box of a large ensemble will evolve toward that of the climatology, containing all the observations and hence almost certainly containing whatever it is we are attempting to forecast; ideally we wish the box to grow as slowly as possible, but no slower.

Several shades of shadowing trajectory have been distinguished, and each has applications in operational forecasting. The distribution of ι -shadowing times reflects the longest time scales over which there exists a model trajectory consistent with the observational uncertainties. How long can operational weather models ι -shadow? Inasmuch as variational data assimilation assumes ι -shadows exist and may degrade the analysis if there is no ι -shadow over the entire assimilation window, knowledge of these time scales is of operational value, because ι -shadowing times would reveal limits to variational assimilation. ϕ -shadows need not stay so near the verification; indeed some practitioners at ECMWF already look for something similar to a ϕ -shadow when evaluating operational forecasts (Tim Palmer, personal communication). Their real value may come from examining historical data: If due to model error no useful forecast exists beyond some time scale, then what can model forecasts (ensemble or otherwise) possibly tell us regarding times beyond that horizon? This predictability horizon, the time scales at which the contribution of model error to the forecast is large compared to the natural variability of the system, is independent of time scales derived from Lyapunov exponents. Sometimes greater, sometimes not. But the question is no longer the classic issue of not being able to find the correct initial condition; it is now an issue of there being no correct initial condition to find.

Accepting the fact that an accurate PDF cannot be obtained allows consideration of other methods of evaluation. Two options are to examine the distribution of ϕ -shadowing times of each model and to estimate the ensemble sizes required to obtain an ensemble bounding box which contains the verification at various lead times. A somewhat more drastic result follows from accepting the ensemble paradigm completely and considering not only ensembles over trajectories from different models, but even individual trajectories which are evaluated using multiple models, the CPT approach introduced in Section 2.6 being a naive first step in this direction. Nevertheless, CPT multi-model ensembles can outperform any individual model in terms of both ϕ -shadowing and producing a good bounding box, while unashamedly producing an ensemble mean that does not resemble the verification, and an MST rank histogram that is inconsistent with an accurate probability forecast. Accepting the limits which exist even in ideal scenarios will force us to reevaluate the aims and evaluation of operational forecasting. Failure to do so is madness: There is no sane approach to an ill-posed goal other than to alter the object of the exercise.

Acknowledgments

It is a pleasure to acknowledge many enlightening discussions with participants at the Newton Institute over the duration of the program. I am particularly indebted to I. Gilmour, J. Hansen, A. Hero, K. Judd, B. Malamud, A. Mees, R. Smith, and P. Young. My understanding of the means and ends of operational NWP ensembles was much improved through enjoyable discussions with T. Palmer and Z. Toth and by the epistemological worries of M. Allen. D. Orrell provided computations on the two-level Lorenz system and its models and new ideas for estimating model error. I am also grateful for the shared insights of M. Berliner, D. Broomhead, R. Buizza, T. Hamill, P. McSharry, A. Provenzale, C. Ziehmann, and many participants who contributed to the high-resolution/low-order (the tastes great/less filling) debate. I remain grateful to C. Sparrow for introducing me to Cambridge in the first instance. This work was supported by Pembroke College, Oxford and the ONR Predictability DRI under grant number N00014-99-1-0056.

References

- [1] H. Abarbanel. Challenges in modelling nonlinear time series. In this volume.
- [2] H.D.I. Abarbanel, R. Brown, and J.B. Kadtko. Prediction in chaotic nonlinear-systems: Methods for time-series with broad-band Fourier spectra. *Physical Review A*, 41(4):1782–1807, February 1990.
- [3] H.D.I. Abarbanel, R. Brown, and M. B. Kennel. Local Lyapunov exponents computed from observed data. *Journal of Nonlinear Science*, 2(3):343–365, 1992.
- [4] R. Ahuja, T. Magnanti, and J. Orlin. *Network Flows*. Prentice Hall, Upper Saddle River, NJ, 1993.
- [5] J. L. Anderson. A method for producing and evaluating probabilistic forecasts from ensemble model integrations. *Journal of Climate*, 9:1518–1530, 1996.
- [6] J. L. Anderson. The impact of dynamical constraints on the selection of initial conditions for ensemble predictions: Low-order perfect model results. *Monthly Weather Review*, 125(11):2969–2983, 1997.
- [7] J. L. Anderson and W. F. Stern. Evaluating the potential predictive utility of ensemble forecasts. *Journal of Climate*, 9(2):260–269, 1996.
- [8] D.V. Anosov. Geodesic flows and closed Riemannian manifolds with negative curvature. *Proc. Steklov Inst. Math.*, 90, 1967.
- [9] C. H. Bishop, B.J. Etherton, and S.J. Majumdar. Adaptive sampling with the ensemble transform Kalman Filter. Part I: Theoretical Aspects. *Monthly Weather Review*, 1999. in review.
- [10] R. Bowen. ω -limit sets for axiom A diffeomorphisms. *J. Diff. Eqns.*, 18:333–339, 1975.
- [11] C. Brankovic and T.N. Palmer. Estimates of seasonal predictability and predictive skill from ecmwf provost ensemble integrations. *Q. J. Royal Meteorol. Soc.*, 1999. in press.

- [12] L. Brillouin. *Scientific Uncertainty and Information*. Academic Press, New York, 1964.
- [13] M. Casdagli, S. Eubank, J.D. Farmer, and J. Gibson. State space reconstruction in the presence of noise. *Physica D*, 51:52–98, 1991.
- [14] J.-P. Eckmann and D. Ruelle. Ergodic theory of chaos and strange attractors. *Rev. Mod. Phys.*, 57:617–656, 1985.
- [15] A. Eddington. *The Nature of the Physical World*. J. M. Dent and Sons, London, 1935. Everyman’s Library, Vol 922.
- [16] E. S. Epstein. Stochastic dynamic prediction. *Tellus*, XXI(6):739–759, 1969.
- [17] S. Eubank and J.D. Farmer. An introduction to chaos and randomness. In E. Jen, editor, *Lectures in Complex Systems*, volume Lecture II of *SFI Studies in the Sciences of Complexity*. Addison-Wesley, 1990.
- [18] W. Fitzgerald. An introduction to Monte Carlo methods for Bayesian data analysis. In this volume.
- [19] G. Froyland. Extracting dynamical behaviour via Markov models. In this volume.
- [20] I. Gilmour. *Nonlinear model evaluation: ι -shadowing, probabilistic prediction and weather forecasting*. D. Phil. Thesis, Oxford University, 1998.
- [21] T. Hamill and S. J. Colucci. Verification of Eta-RSM short-range ensemble forecasts. *Mon. Wea. Rev.*, 125:1312–1327, 1997.
- [22] T. Hamill and C. Snyder. A hybrid ensemble Kalman filter/3D-Variational analysis scheme. *Mon. Wea. Rev.*, 1999. In review (5 Oct 1999), 43 pages.
- [23] T. Hamill, C. Snyder, and R. Morris. A comparison of probabilistic forecasts. *Mon. Wea. Rev.*, 1999. In review (24 Feb 1999).
- [24] J. A. Hansen. *Adaptive Observations in Spatially-extended, Nonlinear Dynamical Systems*. D. Phil. Thesis, Oxford University, 1998.
- [25] J. A. Hansen and L. A. Smith. The role of operational constraints in selecting supplementary observations. *J. Atmos. Sci.*, 1999. in press.
- [26] P.L. Houtekamer, L. Lefaire, J. Derome, H. Ritchie, and H. Mitchell. A system simulation approach to ensemble prediction. *Monthly Weather Review*, 124(6):1225–1242, 1996.
- [27] K. Judd, M. Small, and A. I. Mees. Achieving good nonlinear models. In this volume.
- [28] K. Judd and L. A. Smith. Indistinguishable states I: Perfect model scenario. *Physica D*, 2000. in Review (3 Jan 2000), 21 pages.
- [29] K. Judd and L. A. Smith. Towards forecasting bounding boxes: Applications to both weather and climate. 2000. In preparation for *J. Atmos. Sci.*, 8 pages.
- [30] R. Kistler, E. Kalnay, W. Collins, S. Saha, G. White, J. Woollen, M. Chelliah, W. Ebisuzaki, M. Kanamitsu, V. Kousky, H. van den Dool, R. Jenneand, and M. Fiorino. 2000: The NCEP/NCAR 50-year reanalysis. *Bull. Amer. Meteor. Soc.*, 1999. In press.
- [31] Marquis de Laplace, Pierre-Simon. *Théorie Analytique des Probabilités*. Paris, 1820. Reproduced in the *Oeuvres complètes de Laplace*, Paris, Volume 11, 1886.
- [32] C. E. Leith. Theoretical skill of Monte Carlo forecasts. *Monthly Weather Review*, 102(6):409–418, 1974.
- [33] William Thomson Lord Kelvin. *Baltimore Lectures on Molecular Dynamics and the Wave Theory of Light*. Cambridge University Press, Cambridge, England, 1904.

- [34] E. Lorenz and K. Emanuel. Optimal sites for supplementary weather observations: Simulation with a small model. *J. Atmos. Sci.*, 55:399–414, 1998.
- [35] E. N. Lorenz. Deterministic nonperiodic flow. *J. Atmos. Sci.*, 20:130–141, 1963.
- [36] E. N. Lorenz. Climate predictability. In *GARP Publication No. 16*, pages 132–136. WMO, 1975. Appendix 2.1.
- [37] E. N. Lorenz. Predictability - A problem partly solved. In *Predictability*. ECMWF, Seminar Proceedings, Shinfield Park, Reading, RG2 9AX, UK, 1995.
- [38] E.N. Lorenz. A study of the predictability of a 28-variable atmospheric model. *Tellus*, XVII, 3:321–333, 1965.
- [39] E. Mach. *Knowledge and Error*. Reidel, Boston, 1976. See page 208.
- [40] C.J. Marzec and E.A. Spiegel. Ordinary differential equations with strange attractors. *SIAM J. Appl. Math.*, 38(3):403–421, 1980.
- [41] P. McSharry and L. A. Smith. Better nonlinear models from noisy data: Attractors with maximum likelihood. *Phys. Rev. Lett.*, 1999. in press.
- [42] H. Melville. *Moby Dick*. Oxford University Press, Oxford, 1998. Oxford World’s Classics, (Opening quote from pg 189; see also pg 508).
- [43] F. Molteni, R. Buizza, T.N. Palmer, and T. Petroliaigis. The ECMWF ensemble prediction system: methodology and validation. *Q. J. R. Meteorol. Soc.*, 122:73–120, 1996.
- [44] Allan H. Murphy. What is a “good” forecast? *Weather and Forecasting*, 8:281–293, 1993.
- [45] Allan H. Murphy and R. L. Winkler. A general framework for forecast verification. *Monthly Weather Review*, 115:1330–1338, 1987.
- [46] J.M. Nese. Quantifying local predictability in phase space. *Physica D*, 35:237–250, 1989.
- [47] J.S. Nicolis, G. Meyer-Kress, and G. Haubs. Non-uniform chaotic dynamics with implications to information processing. *Zeitschrift für Naturforschung*, 38 a:1157–1169, 1983.
- [48] D. Orrell. *A shadow of a Doubt: Model Error, Uncertainty, and Shadowing in Nonlinear Dynamical Systems*. 1999. Transfer of Status Thesis, *University of Oxford*.
- [49] E. Ott. *Chaos in dynamical systems*. CUP, Cambridge, 1993.
- [50] T. N. Palmer, R. Buizza, F. Molteni, Y.-C Chen, and S. Corti. Singular vectors and the predictability of weather and climate. *Phil. Trans. R. Soc. Lond.*, A 348(1688):459–475, 1994.
- [51] T.N. Palmer, C. Brankovicand, and D. Richardson. A probability and decision-model analysis of provost seasonal multi-model ensemble integrations. *Q. J. Royal Meteorol. Soc.*, 1999. in press.
- [52] C. Pires, R. Vautard, and O. Talagrand. On extending the limits of variational assimilation in nonlinear chaotic systems. *Tellus*, 48A:96–121, 1996.
- [53] K. R. Popper. *The Open Universe*. Routledge, London, 1982. (Originally published in 1956).
- [54] T. Sauer, J. A. Yorke, and M. Casdagli. Embedology. *J. Stat. Phys.*, 65:579–616, 1991.
- [55] L. A. Smith. Identification and prediction of low-dimensional dynamics. *Physica D*, 58:50–76, 1992.

- [56] L. A. Smith. Accountability in ensemble prediction. In *Predictability*, volume 1 of *ECMWF Workshop Proceedings*, pages 351–368, Shinfield Park, Reading, UK, 1996. ECMWF.
- [57] L. A. Smith. The maintenance of uncertainty. In G. Cini Castagnoli and A. Provenzale, editors, *Past and Present Variability in the Solar-Terrestrial System: Measurement, Data Analysis and Theoretical Models*, volume CXXXIII of *International School of Physics “Enrico Fermi”*, pages 177–246, Bologna, 1997. Il Nuovo Cimento.
- [58] L. A. Smith and I. Gilmour. Accountability and internal consistency in ensemble formation. In *Predictability*. ECMWF, Seminar Proceedings, Shinfield Park, Reading, RG2 9AX, UK, 1998.
- [59] L.A. Smith. Local optimal prediction: Exploiting strangeness and the variation of sensitivity to initial condition. *Phil. Trans. Royal Soc. Lond. A*, 348(1688):371–381, 1994.
- [60] L.A. Smith, C. Ziehmann, and K. Fraedrich. Uncertainty dynamics and predictability in chaotic systems. *Q. J. Royal Meteorol. Soc.*, 125:2855–2886, 1999.
- [61] D.B. Stephenson. Correlation of spatial climate/weather maps and the advantages of using the mahalanobis metric in predictions. *Tellus*, 49 A(5):513–527, 1997.
- [62] O. Talagrand and P. Courtier. Variational assimilation of meteorological observations with the a djoint vorticity equation - part I. *Q.J.R. Meteorol. Soc.*, 113:1311–1328, 1987.
- [63] P. D. Thompson. Uncertainty of initial state as a factor in the predictability of large-scale atmospheric flow patterns. *Tellus*, 9:275–295, 1957.
- [64] H. Tong. *Non-Linear Time Series Analysis*. Oxford Univ. Press, Oxford, 1990.
- [65] Z. Toth and E. Kalnay. Ensemble forecasting at NMC: the generation of perturbations. *Bull. Am. Meteorol. Soc.*, 74(12):2317–2330, 1993.
- [66] H. M. van den Dool. Searching for analogues, how long must we wait? *Tellus*, 46 A(3):314–324, 1994.
- [67] S. Vannitsem and C. Nicolis. Lyapunov vectors and error growth patterns in a T21L3 quasigeostrophic model. *J. Atmos. Sci.*, 54:347–361, 1997.
- [68] R. von Mises. *Probability Statistics and Truth*. George Allen and Unwin, London, 1957.
- [69] P. Young. The identification and estimation of nonlinear stockastic systems. In this volume.
- [70] C. Ziehmann, L.A. Smith, and J. Kurths. The bootstrap and Lyapunov exponents in deterministic chaos. *Physica*, D 126:49–59, 1999.
- [71] C. Ziehmann-Schlumbohm. *Vorhersagestudien in chaotischen Systemen und in der Praxis - Anwendung von Methoden der nichtlinearen Systemanalyse*. PhD. Thesis, Freie Universität Berlin, 1994. Meteorologische Abhandlungen N.F. Serie A Monographien.

Chapter 3

Achieving Good Nonlinear Models: Keep It Simple, Vary the Embedding, and Get the Dynamics Right

Kevin Judd¹
Michael Small
Alistair I. Mees

ABSTRACT This chapter presents an overview of three fundamental notions in modeling nonlinear dynamical systems from time series. They are the use of the minimum description length (MDL) principle in model selection; the use of variable embedding and cylindrical basis models to build models that better capture the dynamics; and the use of $\Psi\Phi$ -models to eliminate systematic error when making long-term prediction. Their purposes are to separate what can be modeled (“determinism”) from what cannot (“noise”); to capture varying time-scales and different geometric features in embedding space; and to make models that have good long-term dynamical behavior as well as short-term predictive ability.

3.1 Introduction

If William of Ockham were alive today and were asked what to keep in mind when constructing a mathematical model of a dynamical system, he might answer, “Keep it simple. . .”. The danger of having an overly complex model is that it might display dynamical behavior totally unlike the behavior displayed by the system. One needs to adjust the complexity of a model so it can display most of the observed behavior of the system, while minimizing the potential of the model misbehaving. Our suggested method for doing this is to apply the principle of minimum description length (MDL), which is derived from information theory. We have found that MDL models capture dynamics better than models built without taking account of

¹ Author for correspondence.

information theoretic aspects.

Constructing or fitting a model also implies working within some restricted *model class*. So another related issue is what class of models has sufficient complexity to capture the dynamical behavior of a large variety of systems. What we suggest is unusual. For some time it has been standard practice to first embed a time series (a Takens embedding), then to build a model (say, a radial basis model) in the embedded space. We introduce the notion of a *variable embedding*, which can be thought of as an embedding that changes with the state of the system. The equivalent of a radial basis model in a variable embedding scheme is a *cylindrical basis model*. The authors have found that variable embedding and cylindrical basis models capture dynamics better than standard uniform embedding and radial basis models. In short, cylinders stack up better than spheres.

A main focus in this chapter is the capturing of the dynamics of the system in a model, since this is a better measure of success than just prediction error. However, despite one's best efforts, a model will not be perfect and may still make systematic errors. The most likely cause of systematic errors (assuming everything else is done well) is that the "true" system does not lie in the model class one has chosen; that is, there is no model in the model class that has exactly the dynamics of the system. This most clearly shows itself when making long-term predictions. Even here there is something to be done; we suggest a simple technique of stacking models (called $\Psi\Phi$ -models [6]) that correct systematic errors and allow better long-term prediction with little additional effort.

3.2 Minimum Description Length Models: Keep It Simple

The minimum description length principle is an application of Ockham's Razor in a modeling context. It defines the best model for a time series to be the one that achieves the most concise description of the data. To understand how the principle works, suppose you (the "sender") have collected an experimental time series $x(t)$, $t = 1, \dots, n$ measured to an accuracy of (say) 12 bits and you wish to communicate this data to a colleague (the "recipient"). You could send the raw data. Alternatively, you could construct a dynamical model from the data that enables the recipient to predict a value of $x(t)$ from earlier values. If you and your colleague have previously agreed on a class of models, then you could communicate the data by sending the parameters of a model, enough initial data to start predicting future values of the time series, and the errors between the true time series and the values predicted by the model. Given this information, the recipient can reconstruct the experimental data to its full measured accuracy. An important point is that the parameters and errors need only

be specified to finite accuracy. Furthermore, if the model is good, then the total number of bits required to transmit parameters, initial values and errors will be less than the number of bits of raw data.

In practice the minimum description length principle requires calculating an approximation to the *description length* of the time series and model, which is effectively the number of bits required to transmit the model plus the number of bits required to transmit the errors. (The initial conditions are included in the parameter count, although their effect only matters when we are comparing different embedding dimensions.) Under fairly general assumptions one can write:

$$\begin{aligned} (\text{Description length}) \approx & \\ & (\text{number of data}) \times \log (\text{mean square prediction error}) \\ & + (\text{penalty for number and accuracy of parameters}) . \end{aligned}$$

As the number of parameters in a model increases the (in-sample) prediction errors decrease, but eventually, the penalty for introducing another parameter outweighs the benefit it has in reducing (the in-sample) prediction errors. The model that attains the minimum description length is the optimal model within the class of models considered. We do not have space here to discuss in detail why this is successful; extensive discussions are to be found elsewhere [10, 4, 14].

In special model classes, explicit approximations to the description length can be calculated. A particularly useful class of parameterized nonlinear autoregressive model consists of those we call *pseudo-linear* models, also called general linear models, which have the form

$$x(t+1) = \sum_{i=1}^m \lambda_i f_i(v(t)) + \epsilon_t, \quad (3.1)$$

$$v(t) = (x(t), x(t-1), \dots, x(t-d)) \quad (3.2)$$

for some selection of nonlinear functions f_i , unknown parameters λ_i and unknown independent and identically distributed random variates ϵ_t . (Observe in passing that choosing $v(t)$ amounts to using a particular embedding.) Define

$$V_i = (f_i(v(1)), \dots, f_i(v(n)))^T, i = 1, \dots, m, \quad (3.3)$$

$$y = (x(1), \dots, x(n))^T, \quad (3.4)$$

$$\lambda = (\lambda_1, \dots, \lambda_m)^T, \quad (3.5)$$

and let V be the matrix whose columns are V_i , $i = 1, \dots, m$. If the ϵ_t are assumed to be Gaussian and λ has been chosen to minimize the sum of squares of the prediction errors $e = y - V\lambda$, then [4] the description length

is bounded by

$$\left(\frac{n}{2} - 1\right) \ln \frac{e^T e}{n} + (k + 1) \left(\frac{1}{2} + \ln \gamma\right) - \sum_{j=1}^k \ln \delta_j, \quad (3.6)$$

where k is the number of non-zero components of λ , γ is related to the scale of the data (for example, a constant used to scale the observations to lie in the unit interval) and δ solves $[Q\delta]_j = 1/\delta_j$ where

$$Q = \hat{V}^T \hat{V} / e^T e$$

and \hat{V} is composed of just those columns of V that correspond to nonzero elements of λ . The variables δ can be interpreted as the relative precision to which the parameters λ are specified.

The attraction of pseudo-linear models is that the parameters λ are easily calculated, because the sum of squares of the prediction errors $e = y - V\lambda$ can be minimized efficiently using singular value decomposition or any of its many equivalents. What makes general pseudo-linear models different from, and more powerful than, special cases such as linear or global polynomial models, is that the basis functions f_i can be chosen in many ways.

The critical problem is how to select the basis functions f_i . In general these will be nonlinear functions depending on various additional parameters that should be optimized over. Unfortunately, this optimization is nonlinear and so is in general difficult, slow and prone to capture by local minima. (This problem is well known in modeling via single-layer neural nets, a particular pseudo-linear approach.) Instead of optimizing the parameters of a few basis functions, we can generate many fixed basis functions, not only at the start but also adaptively as the model building progresses, and select a subset of them that optimizes the description length.

This alternative scheme requires an efficient combinatorial optimization method to select an optimal subset of the basis functions. It would appear that we have made the problem worse, because combinatorial optimization is notoriously hard, but in fact the following subset selection algorithm, described in detail elsewhere [4] is very successful in most of the applications we have considered. The algorithm selects subsets that are near-optimal according to the description length criterion and hence produces good pseudo-linear models. It operates by adding and removing candidate functions from a given basis set according to a local optimality criterion, and accepting a set of given size as optimal if the same candidate is removed as was just added. The size of the basis set is increased until the description length criterion says it has become too large, and then the best set found so far is selected as the overall optimum.

In the algorithm, B represents any set of $k < m$ indices in $\{1, \dots, m\}$. We write V_B for the $n \times k$ matrix formed from the columns of V with indices in B , λ_B for the least squares solution to $y = V_B \lambda$, and $e_B = y - V_B \lambda_B$.

Algorithm 1

1. Normalize the columns of V to have unit length.
2. Let $S_0 = (\frac{n}{2} - 1) \ln(y^T y/n) + \frac{1}{2} + \ln \gamma$ (the description length of the raw data).
3. Let $B = \{j\}$ where V_j is the column of V such that $|V_j^T y|$ is maximum (this selects as the first basis function the one that most closely matches the data y ; note that $\lambda_B = V_j^T y/V_j^T V_j$ in this case).
4. Let $\mu = V^T e_B$ and i be the index of the component of μ with maximum absolute value. Let $B' = B \cup \{i\}$ (the components of the vector μ measure how closely each of the basis functions not currently in use will match the error of the current model; extend the current model with the basis function that best matches the current error).
5. Calculate $\lambda_{B'}$. Let o be the index in B' corresponding to the component of $\lambda_{B'}$ with smallest absolute value (here o is the index of the basis function that makes the smallest contribution to the current extended model).
6. If $i \neq o$, then put $B = B' \setminus \{o\}$ and go to step 4. Otherwise, set $B = B'$. (Throw out the “worst” basis function o if it is not i , the last one we brought in; then go back and try again. Otherwise, the extended basis B' is taken to be the “locally” optimal basis.)
7. Define $B_k = B$, where $k = |B|$. Find δ such that $(V_B^T V_B \delta)_j = 1/\delta_j$ for each $j = \{1, \dots, k\}$ and calculate $S_k = (\frac{n}{2} - 1) \ln \frac{\hat{e}^T \hat{e}}{n} + (k + 1)(\frac{1}{2} + \ln \gamma) - \sum_{j=1}^k \ln \hat{\delta}_j$. (At this stage we have found the best model of size k that can be built from the best model of size $k - 1$ by “bringing in the best and throwing out the worst.”)
8. If $S_k < S_{k-1}$, then go to step 4. (Continue until the description length stops decreasing.)
9. Take the basis B_k such that S_k is minimum as the optimal model.

3.3 Variable Embedding: Cylinders Stack Up Better than Spheres

For some time it has been a standard practice to embed a time series (a Takens embedding by delay reconstruction); then build a model (say, a radial basis model) in the embedded space. Unfortunately, there are many types of dynamical behavior that are not modeled well by this technique.

3.3.1 Uniform Embedding

Given a time series $x(t) \in \mathbb{R}$ one might form an embedded time series

$$z(t) = (x(t), x(t - \ell), \dots, x(t - (d - 1)\ell)) \in \mathbb{R}^d$$

where ℓ is called the *lag*. In anticipation of our discussion we will refer to this embedding as a *uniform* embedding. The lag is introduced to improve the observability, for example, of noisy time series. The lag is chosen to optimize the spread of the embedded time series without confusing the dynamics and to obtain an embedding that is independent of the sampling rate (of an oversampled continuous time system, for example). There are two principal methods of choosing the lag: the first zero of the autocorrelation function [1] and the minimum of the mutual information [3].

Uniform embeddings for modeling purposes are at their most effective when embedding a time series with a single dominant periodicity or recurrence time. Both of the above mentioned methods for calculating lags give similar lags for such time series and the lag is approximately one-quarter of the dominant period. For this lag the embedded time series is ring-shaped; shorter and longer lags result in elliptical rings, or—if the lags are too far from a good value—a scrambled mess. A good lag in this case keeps states that correspond to similar phases close together and anti-phase states as far apart as possible. Uniform embeddings are quite suitable for classic chaotic systems such the Rössler and Lorenz systems, which have a single dominant periodicity or recurrence time.

3.3.2 Nonuniform Embedding

Uniform embeddings can fail when there are multiple strong periodicities with differing time scales. For example, consider a quasi-periodic time series with differing frequency components or with very close frequencies that lead to a “carrier” frequency and a “modulation” of differing period. Figure 3.1 shows three time series, from different systems, that all possess short and long period recurrences. Uniform embedding fails for such time series because a short lag would be optimal for the high-frequency component and a long lag would be optimal for the low-frequency components and modulation, while a compromise lag is inadequate for both time scales.

One way to avoid this problem is to use nonuniform embedding strategies. For example, with the sunspot time series one would most likely choose to use the uniform embedding $(x(t), x(t - 3), x(t - 6))$, but the authors find that the embedding $(x(t), x(t - 2), x(t - 8))$ gives better results when constructing radial basis models [4, 5]. Note how compared to the uniform embedding the nonuniform embedding has two components that are more closely spaced and yet a broader window overall; alternatively, one could say the uniform embedding is a compromise between a shorter lag and a wider window.




FIGURE 3.1. Three time series that have multiple, strong periodicities: (a) the average annual sunspot number; (b) a sampling at 50Hz of a measurement that is proportional to the cross-sectional area of the abdomen of a child during quiet sleep: the phenomenon observed here is called “periodic breathing”; (c) a 12kHz recording of the Japanese vowel [a].

One can use a nonuniform embedding in any situation where one uses a uniform embedding; there is just a little additional bookkeeping required to ensure enough of the past time series is retained and that correct observations are used to make each forward prediction. Of course, to make long-term predictions or simulations of the dynamics, the prediction step must be a divisor of *all* the lags.

3.3.3 Variable Embedding

Going a step further, we suggest it is often advantageous to vary the embedding strategy with the state of the system. To visualize why it is useful to do this, consider modeling the Lorenz system with its butterfly-shaped attractor. When the system state is out on the “wings” of the attractor, a two-dimensional embedding is sufficient to model and predict motions—one would require a lot of very high-quality data to discern the thickness

of these wings. However, near the origin where the crossover of the wings occurs, a three-dimensional embedding is essential. One could imagine constructing a perfectly adequate model that does not use a global embedding but rather uses appropriate local embeddings as the system state varies.

If one uses a variable embedding, then the processes of embedding and modeling are merged into one process with a single optimization goal of finding a compact and accurate model. An example of this, and more generally how variable embedding can be implemented, is the class of *cylindrical basis models*.

3.3.4 Cylindrical Basis Models

The standard radial basis model [8, 2, 7] is pseudo-linear, with each of the nonlinear functions depending only on the radial distance from a certain point called a center. That is, in the standard pseudo-linear model of equation 3.1, the functions have the form $f_i(z) = \phi(|z - c_i|/r_i)$ for suitably chosen centers c_i , radii r_i and function ϕ . If the function ϕ is decreasing, we can think of the action of each f_i localized on a ball. If we ignore some coordinates (as a result of a local nonuniform embedding) then the functions act locally on cylinders, instead of radially symmetric spheres, so a reasonable name is *cylindrical basis models*. To construct such models we must define the various cylinders.

A cylinder is defined by a center c_i , a radius r_i and a lag vector $(\ell_1, \ell_2, \dots, \ell_k)$. The lag vector corresponds to a projection P such that

$$P(v(t)) = (x(t - \ell_1), x(t - \ell_2), \dots, x(t - \ell_k)).$$

The basis functions

$$f_i(z) = \phi(|P_i(z - c_i)|/r_i), \quad (3.7)$$

with decreasing ϕ , have the effect of localizing the embedding in a cylindrical neighborhood of c_i , where the axis of the cylinder is parallel to the components of $v(t)$ that have been projected out, that is, the lags that are missing in $P(v(t))$. To simplify the notation we have introduced a redundancy in (3.7), because the center c_i is unique only up to the projection P_i . A very simple example of a cylindrical basis function given $(x, y, z) \in \mathbb{R}^3$ a Gaussian radial basis function with center $(2, -1, 5)$ and radius 3 would be

$$f(x, y, z) = e^{((x-2)^2 + (y+1)^2 + (z-5)^2)/18}.$$

Whereas given a projection $P(x, y, z) = (x, z)$ there is a cylindrical basis function

$$f(x, y, z) = e^{((x-2)^2 + (z-5)^2)/18}.$$

To construct a cylindrical basis model using the selection methods we have described, one would generate a large set of basis functions (3.7), with

decreasing function ϕ , having different centers c_i , radii r_i and projections P_i (which define the lag vectors). There is an obvious combinatorial explosion here because for each potential center there are $2^d - 1$ possible lag vectors. This could be tackled using genetic algorithms or simulated annealing, but we have found that a relatively simple adaptation of our earlier algorithms appears to avoid the worst of this explosion.

Algorithm 2

1. Let S represent an initial set of candidate basis functions. These functions are likely to be generated randomly but possibly using some additional information to choose likely candidates for selection, for example, a selection weighted by the errors of the current best model. One good way to start is to generate centers and radii as usual and then apply the reduced autoregressive method *locally* in the region around each center to get a local lag vector. This lag vector is only an initial guess; it will typically be thinned later in the process.
2. Apply Algorithm 1 to determine the best model using the basis functions of S . Let S^* be the selected basis functions.
3. If desired, locally optimize S^* by tuning its parameters using some standard method such as the Levenberg-Marquardt algorithm [9].
4. Generate a new set of candidate basis functions S that includes S^* . The new candidate functions might be chosen with additional knowledge gained from the selection of S^* . (For example, we might try simplifying cylinders by applying additional projections, randomly perturbing selected cylinders, or putting new cylinders near where the model makes its worst predictions.)
5. Return to step 2, and continue to do so while S^* is changing (significantly).

The authors have found that cylindrical basis models are much more successful at capturing the dynamics of a system than radial basis models [5, 11, 12].

3.4 Systematic Errors: $\Psi\Phi$ -Models

One of the important reasons for wanting to model the dynamics correctly is for making long-term predictions. One can obtain longer-term predictions from a short-term predictor by simply iterating the predictor. Such longer-term predictions are limited in their accuracy by observational and dynamic noise and the sensitivity to initial conditions of the dynamical system, but iterated predictors are also limited by systematic errors in the short-term

predictor, which can arise from under- or overfitting data, or from the model class not containing the system under study.

There are many reasons why a model can have systematic errors. The model may have been derived from physical or chemical principles but various assumptions may not hold in practice, or a dynamic model may lack sufficient resolution, or the exact nature of the noise may be unknown. For example, even an analogue circuit or mechanical device [13] designed to instantiate the Lorenz equations would not be modeled exactly by the Lorenz equations, because the diodes and capacitors, or mechanical devices, will have slightly different nonlinearities. Indeed it is likely that there is no easily described transcendental functions that model the apparatus exactly. This implies that there would not be any reasonable model class that exactly describes the system; the system is outside all reasonable model classes and any reasonable model will display some systematic error.

In principle, with black box and other models, one could make a model increasingly more complex to account for the systematic errors, but there comes a point when it is more productive to abandon this adaptive approach that merely adds more components of similar type and try something entirely different in character. What we suggest here is not to throw out a reasonably good model, but to apply a method that augments and corrects it cheaply and effectively [6]. The just cited paper discusses in more detail how the method we are about to introduce has the effect of greatly expanding the general linear model class in a way that is useful for modeling dynamical systems. Another way of looking at the method is that it tries to rearrange and present the information in a time series in a more useful or accessible way.

It is an important observation that, from an information theory perspective, iterated predictors have a fundamental flaw. Given a time series x_1, \dots, x_t , we can think of a one-step predictor φ as attempting to add an additional datum \hat{x}_{t+1} to the end of the time series, where

$$\hat{x}_{t+1} = \varphi(x_t, x_{t-1}, \dots, x_{t-d+1}) = x_{t+1} + \text{error}. \quad (3.8)$$

Iterated prediction to obtain further \hat{x}_{t+m} requires shifting the arguments of φ one place to the right (removing x_{t-d+1}) and substituting the predicted \hat{x}_{t+1} into the first position, and so on. Observe that each subsequent iterated prediction uses no more information than was used in the first one-step prediction. From an information theory point of view there is a mistake here, the longer-term predictions ought to use *more* information. The $\Psi\Phi$ -method proposed in this chapter attempts to correct this failure by exploiting information retained in the systemic errors of iterated predictions.




FIGURE 3.2. Schematic of a time series (solid line) and the sections of the time series that comprise the *history* vector h_t and the *future* vector f_t . The dotted line represents an iterated predictions z_t with systematic error.

Referring to Figure 3.2, define

$$\begin{aligned} h_t &= (x_t, x_{t-1}, \dots, x_{t-d+1}) \in \mathbb{R}^d, \\ f_t &= (x_{t+1}, \dots, x_{t+p}) \in \mathbb{R}^p, \\ z_t &= (\hat{x}_{t+1}, \dots, \hat{x}_{t+q}) \in \mathbb{R}^q, \end{aligned}$$

where \hat{x}_{t+m} is the m th iterated prediction of φ , using the predict, shift, substitute method. The vector h_t represents the important recent *history* of the time series at time t , and f_t is the next p steps of *future* after time t . The vector z_t is our predictions of the next q steps obtained by iterated prediction with φ .²

Ideally we want a mapping $h_t \mapsto f_t$, but we only have a map

$$z_t = \Phi(h_t), \quad (3.9)$$

which gives an estimate z_t of f_t ; if $p = q$, then

$$f_t = z_t + \text{errors}.$$

If there is observational or dynamic noise, or if there is overfitting or underfitting, or if the system is not in our model class, then the predictions z_t of the future f_t may be inaccurate or even erroneous. Also the predictions are incomplete when $q < p$.

²How we obtained the predictions is not important: they could be obtained from a physical model, a black-box model, or a multi-step predictor, or they could even be predictions from several different models.

To correct the errors and inaccuracies of the *predictor* Φ we propose finding a *corrector* Ψ ,

$$f_t = \Psi(z_t) + \textit{smaller errors}. \quad (3.10)$$

The essential point here is that if the errors of Φ are *systematic*, then Ψ can exploit this information to produce a better prediction of f_t . We shall see that Φ can be quite a poor predictor, but provided it is consistent, a Ψ mapping can improve it greatly.

We will refer to the method just described as the $\Psi\Phi$ -method³. The surprising thing is that this method works well, even for simple Ψ , yielding consistently good results for both artificial and experimental systems using cubic polynomial maps for Ψ . Radial basis models have also been used to obtain marginally better results than we show here, but on occasions even a linear or quadratic Ψ can work well.

Figure 3.3 shows long-term predictions obtained for the Lorenz system with $\approx 15\%$ observational noise using the $\Psi\Phi$ -method; a detailed discussion of this figure and other experimental results can be found elsewhere [6].

$\Psi(z_t)$ can best be thought of as a prediction of the mean long-term behavior of the system. This can be seen in Figure 3.3 where it can sometimes happen that so much critical information about the system has been lost that it is no longer possible to predict which “wing” of the Lorenz attractor the system will be on, and consequently $\Psi(z_t)$ predicts that the system’s mean behavior is zero. This is, of course, the correct prediction, as disappointing as it may seem. On the other hand, in Figure 3.3 it can be seen that the estimated prediction error (upper and lower solid lines) of this conservative estimate may vary in a way that demonstrates that not all dynamic information has been lost.

Determining the Corrector Ψ

We now describe how to determine a suitable corrector Ψ and how to estimate prediction errors of the $\Psi\Phi$ -map. We will only consider the case where Ψ is linear in its parameters, for example, linear, polynomial, radial basis, pseudo-linear and general linear, because such mappings are easy to work with and appear to be adequate in all cases we have examined.

The length p of the future vector f_t may be greater than the length q of the iterated prediction vector z_t , so that Ψ extrapolates beyond the end of z_t ; indeed Ψ could also interpolate by using an alternative definition

$$z_t = (\hat{x}_{t+s_1}, \hat{x}_{t+s_2}, \dots, \hat{x}_{t+s_k}) \in \mathbb{R}^k,$$

where $0 < s_1 < \dots < s_k \leq q$. This “thinned out” alternative for z_t is

³ $\Psi\Phi$ is pronounced *sci fi*, in the “To boldly go...” sense.

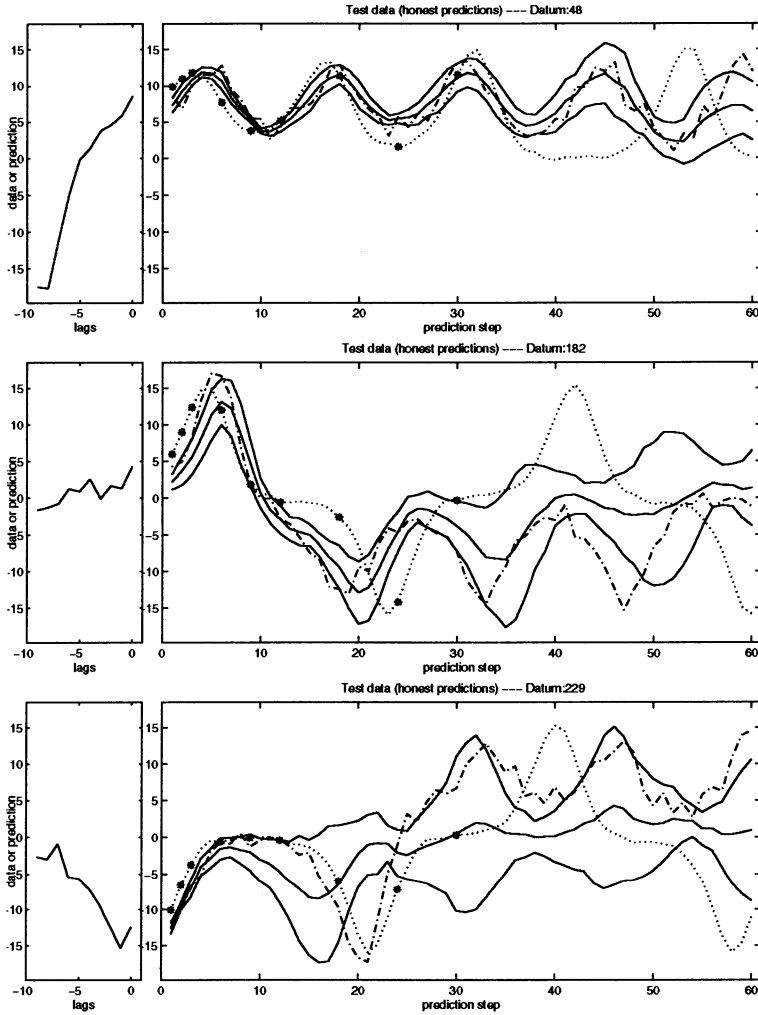


FIGURE 3.3. Three indicative long-term predictions obtained for the Lorenz system with $\approx 15\%$ observational noise using the $\Psi\Phi$ -method. The model was built from 4000 points sampled at 0.05 second. The test data was a distinct trajectory. The left panels show h_t ; the right panels show the future f_t (dash-dotted line); the iterated Φ -model predictions (dotted line) with components of $\Phi(h_t)$ used in z_t indicated by stars, i.e. a thinned prediction vector was used; the $\Psi\Phi$ -model predictions (central solid line) and the Θ predictions of absolute deviations (upper and lower solid lines).

useful to reduce the effort involved in computing Ψ . Assume, without loss of generality that $k \leq q \leq p$, which is the natural thing to do.

Given a time series $x = \{x_t : -d + 1 \leq t \leq n + p\}$ one can construct vec-

tors h_t and f_t for $t = 1, \dots, n$. Then for each h_t one can calculate a z_t that is an *in-sample* long-term prediction of the f_t . Any discrepancies between z_t and the corresponding components of f_t are clues to the systematic errors of Φ at h_t ; see Figure 3.2.

Given n vectors z_t and f_t the construction and estimation of the mapping Ψ is easily achieved when Ψ is linear in its parameters. If, for example, Ψ is also linear in the components of z_t , then Ψ can be defined by a $p \times q$ matrix A (or $p \times k$ for the alternative definition of z_t) such that,

$$\Psi(z_t) = Az_t. \quad (3.11)$$

In general, if Ψ is a linear combination of m *basis* functions $g_i(z)$, $i = 1, \dots, m$, then

$$\Psi(z_t) = A\zeta_t \quad (3.12)$$

where $\zeta_t = (g_1, \dots, g_m)(z_t) \in \mathbb{R}^m$ and A is some $p \times m$ matrix. For example, a polynomial Ψ has as basis functions the constant function, projections onto each component of z_t , and the products of powers of the components of z_t up to the order of the polynomial.

Estimation of the parameter matrix A is straightforward using least squares. If one chooses to minimize the sum of squares of prediction errors $\|f_t - \Psi(z_t)\|$, then this corresponds to solving the following matrix equation in the least squares sense (by singular value decomposition, for example)

$$F = AZ, \quad (3.13)$$

where F is the $p \times n$ matrix having the future vector f_t as the t th column, and Z is the $m \times n$ matrix whose t th column is ζ_t . One could conceivably want to use a weighted norm when calculating $\|f_t - \Psi(z_t)\|$, since the standard Euclidean norm weights all prediction errors equally, regardless of how far into the future; however, we have used the simple length norm.

The procedure is summarized in the following algorithm.

Algorithm 3

1. From a given time series, form the vector time series h_t and f_t .
2. For each “initial condition” h_t iterate the model Φ (whatever that model might be) to form the predictions vector z_t .
3. Extend the vector z_t to include powers and products of the components (optionally, use some other set of basis function).
4. Pack the column vectors f_t and z_t into matrices F and Z .
5. Solve $F = AZ$.
6. To make a long-term prediction, iterate the Φ model from an initial condition and multiply this vector of predictions into A .

Estimating Prediction Errors

It is desirable to predict how good the $\Psi\Phi$ -predictions are, because this essentially estimates the prediction horizon in advance. Because we do not anticipate a necessarily Gaussian distribution of prediction errors, and because it is easier to calculate, we will attempt to predict the *robust* estimator of spread given by the mean absolute differences, rather than attempting to predict the variance of the prediction errors.

Define the absolute prediction error vector e_t by

$$e_t = |f_t - \Psi(z_t)| \in \mathbb{R}^p \quad (3.14)$$

where $|\cdot|$ applies componentwise.

We now propose to estimate a mapping $\Theta: \mathbb{R}^q \rightarrow \mathbb{R}^p$ (or from \mathbb{R}^k for the alternative definition of z_t) such that $\Theta(z_t)$ estimates e_t . There is no reason to estimate this map any differently from how we proposed estimating Ψ , that is, choose Θ to be linear in its parameters and so

$$\Theta(z_t) = B\zeta_t, \quad (3.15)$$

where ζ_t is the basis functions evaluations as before, although possibly a different set of basis functions, and B is some $p \times m$ matrix. Estimating Θ by least squares as before requires solving the matrix equation in the least squares sense

$$E = BZ \quad (3.16)$$

where the t th column of E is e_t the error given in (3.14).

3.5 Conclusions

The focus of this chapter has been the development of reliable methods of modeling nonlinear dynamical systems from time series, so that the model captures the dynamics. This involves using the minimum description length principle to ensure that models neither underfit nor overfit data, using variable embeddings to locally optimize models to the dynamics, and using $\Psi\Phi$ -models to correct any residual systematic errors when making long-term predictions.

Acknowledgments

The work of AIM was partially supported by a grant from the Australian Research Council. AIM also thanks the Department of Systems Engineering at The Chinese University of Hong Kong for hospitality.

References

- [1] A. M. Albano, A. I. Mees, G. C. deGuzman, and P. E. Rapp. Data requirements for reliable estimation of correlation dimensions. In H. Degn, A. V. Holden, and L. F. Olsen, editors, *Chaos in biological systems*, pages 207–220. Plenum, New York, 1987.
- [2] M. Casdagli. Nonlinear prediction of chaotic time series. *Physica D*, 35(3):335–356, 1989.
- [3] A. M. Fraser and H. L. Swinney. Independent coordinates for strange attractors from mutual information. *Physical Review A*, 33(2):1134–1140, 1986.
- [4] K. Judd and A. I. Mees. On selecting models for nonlinear time series. *Physica D*, 82:426–444, 1995.
- [5] K. Judd and A. I. Mees. Embedding as a modeling problem. *Physica D*, 120:273–286, 1998.
- [6] K. Judd and M. Small. Towards long-term prediction. *Physica D*, 136:31–44, 2000.
- [7] A. I. Mees. Parsimonious dynamical reconstruction. *International Journal of Bifurcation and Chaos*, 3(3):669–675, 1993.
- [8] A. I. Mees, M. F. Jackson, and L. O. Chua. Device modeling by radial basis functions. *IEEE Trans CAS/FTA*, 39(1):19–27, 1992.
- [9] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes in C*. Cambridge University Press, Cambridge, 1988.
- [10] J. Rissanen. *Stochastic Complexity in Statistical Inquiry*, volume 15 of *Series in Computer Science*. World Scientific, Singapore, 1989.
- [11] M. Small and K. Judd. Detecting nonlinearity in experimental data. *International Journal of Bifurcation and Chaos*, 8(6):1231–1244, 1997.
- [12] M. Small and K. Judd. Comparisons of new nonlinear modeling techniques with applications to infant respiration. *Physica D*, 117:283–298, 1998.
- [13] S. H. Strogatz. *Nonlinear Dynamics and Chaos*. Addison Wesley, New York, 1994.
- [14] P. Vitanyi and M. Li. Ideal MDL and its relation to Bayesianism. In D. L. Dowe, K. B. Korb, and J. J. Oliver, editors, *Information, Statistics and Induction in Science*, pages 282–291, Melbourne, Australia, 1996. World Scientific.

Chapter 4

Delay Reconstruction: Dynamics versus Statistics

Jaroslav Stark

ABSTRACT Traditionally delay reconstruction has been seen as lying in the realm of dynamics or differential topology. It is thus perceived to be a largely automatic procedure that reconstructs an existing dynamical system. In this chapter we argue that it is as imprecise as all other parts of time series analysis and it should be subject to as much statistical scrutiny as procedures such as modeling, prediction and noise reduction.

4.1 Introduction

Most of the developments in the field of nonlinear dynamics over the past century have assumed that one had a complete description of the dynamical system under consideration. In principle the practical application of these results thus requires the simultaneous measurement of all the state variables. Unfortunately, in many real problems one has only the sketchiest information about what these variables are, and one certainly has no hope of observing them all. Instead, one typically has a time series of one or more observables of the system, whose relationship to the state variables is at best uncertain.

The study of time series has traditionally been within the realm of statistics. This has developed a large body of both theory and practical algorithms for characterizing, modeling, predicting and filtering such data. Such techniques are widely and successfully used in a broad range of applications, including communications, epidemiology and finance. Until very recently, however, this approach largely made use of linear models, and hence it was unable to take advantage of recent developments in nonlinear dynamics. In particular, it is now widely accepted that even simple nonlinear mechanisms can give rise to complex temporal behavior (i.e., chaos) and hence to complex time series. Conventional statistical time series approaches can have considerable problems with such time series and can fail to model or predict them with any degree of accuracy. This is because they have no way of using the fact that the time series was generated by a completely deterministic process, and hence they have to ascribe most of the complexity to random noise. Furthermore, such approaches cannot yield much useful

information about the properties of the original dynamical system.

Fortunately, a remarkable result due to Takens some twenty years ago shows that typically one can reconstruct the dynamics of an unknown deterministic finite-dimensional system from a scalar time series generated by that system ([Takens, 1980]; see also Aeyels [1981]). Over the last two decades this has given rise to a huge range of applications throughout science and engineering, has stimulated the reexamination of old data sets and has motivated the development of a variety of new experiments (for good overviews see Ott et al. [1994], Abarbanel [1995] and Kantz and Schreiber [1999]). Even today, this one theorem underpins all known approaches to chaotic time series (i.e., time series presumed to be generated by chaotic dynamical systems).

Takens' Theorem is actually a result in differential topology which is an extension of a classical theorem due to Whitney in the 1930's. It is thus concerned with purely deterministic autonomous dynamical systems and the framework that it provides for time series analysis is unable to incorporate any notion of random behavior. This means that the process of reconstruction is outside the scope of statistical analysis because any such analysis requires a stochastic model of one kind or another as its starting point. This is reflected in common practice, where reconstruction is seen as a straightforward algorithmic procedure that aims to recover properties of an existing, but hidden, system. It is thus expected that there is a "correct" value for quantities such as the (minimal) embedding dimension (i.e., the number of state variables in the reconstructed system).

On the other hand, no real system and no real data set are entirely free of noise or external influence. Thus strictly speaking we can never appeal to Takens' theorem to justify reconstruction. This raises the question of what exactly delay reconstruction is doing when it is applied to real data and what kind of answers we can expect to obtain from it. This is implicitly acknowledged in most work in this area, where there is an increasing trend to incorporate statistical methods in those procedures that follow reconstruction (such as characterization using fractal dimensions or Lyapunov exponents, modeling, prediction and noise reduction). The situation is thus as shown in Figure 4.1. Note that all the procedures in the shaded region have analogues in the realm of traditional statistical time series analysis. Thus characterization has a close relationship with hypothesis testing, while modeling, prediction and filtering all correspond to various types of estimation. It is thus completely natural that an increasing level of statistical sophistication is incorporated in these tasks in the context of processing delay reconstructed nonlinear time series.

Progress in this direction is, however, hampered by the lack of a proper probabilistic basis for the upper part of the diagram. If we assume that our dynamics and observations are deterministic, so that we can apply Takens' theorem, then we end up with a deterministic reconstructed system. In such a case, from where does the randomness that we require to apply the

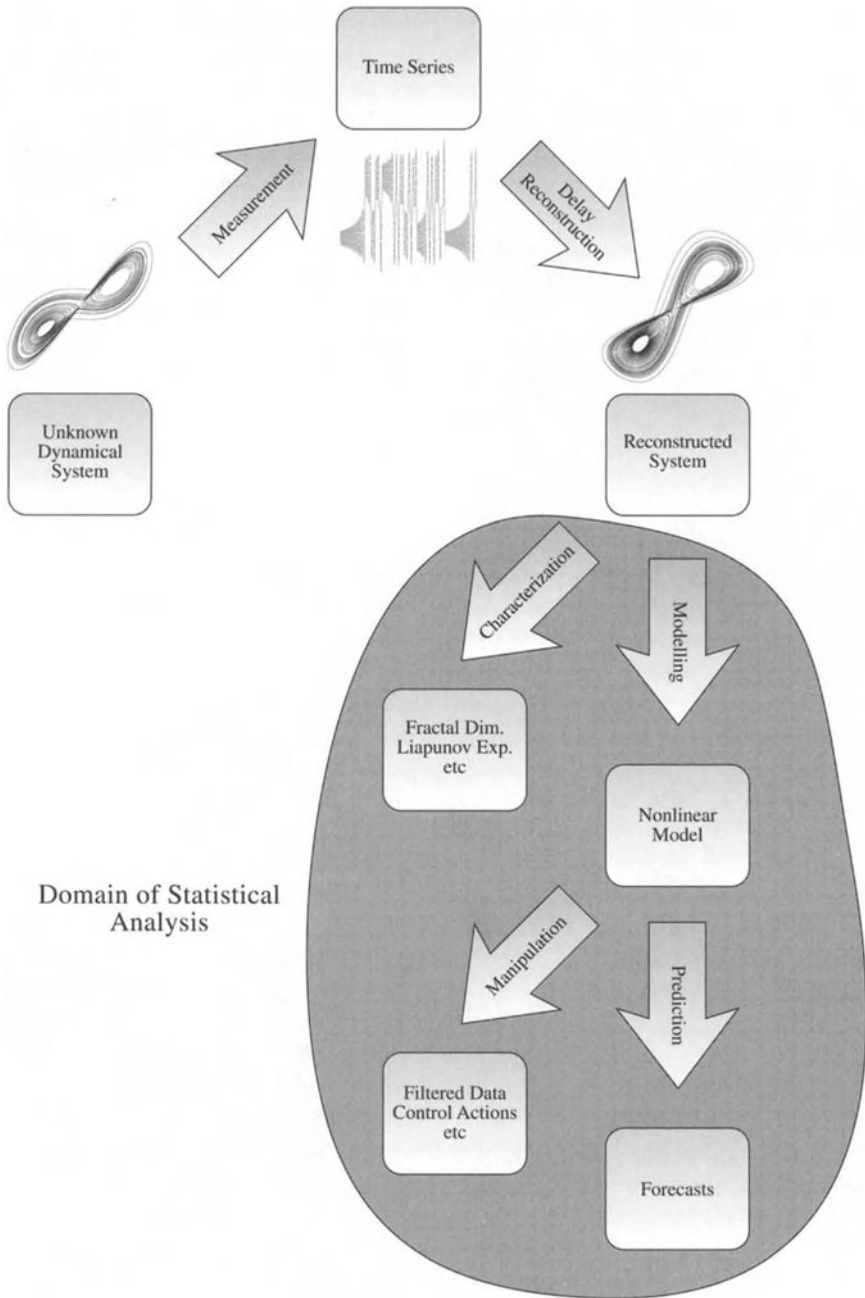


FIGURE 4.1. Overview of delay reconstruction and subsequent analysis of time series generated by nonlinear dynamical systems.

statistical methods in the shaded region come? Alternatively if we accept that our system and observations are noisy, then we have little insight into how much information is preserved by the delay reconstruction process and little indication of how to properly set up statistical models for subsequent analysis. The aim of this chapter is therefore to introduce a framework for the delay reconstruction of random dynamical systems that has been developed by the author and his collaborators over the last few years and to discuss some of its applications and consequences. Like the standard Takens' theorem, this framework is presented from a dynamical systems perspective and in itself has little statistical context. However, because it allows for uncertainty to be incorporated in the reconstruction process, we hope that it will encourage statisticians to explore its implications and develop appropriate statistical tools for this vital part of nonlinear time series analysis.

4.2 Conventional Delay Reconstruction

Although Takens' theorem will be familiar to most readers, we begin by giving a brief summary. This is to both introduce the notation we shall need throughout the chapter, and remind the reader of a number of important aspects of the Theorem that will be relevant to our extension to random systems. We denote the state space of our dynamical system as M , which is assumed to be a finite-dimensional compact manifold. Extensions to non-compact manifolds (e.g., \mathbb{R}^m) or to compact invariant sets that are not manifolds (e.g., attractors) do exist [Takens, 1980; Sauer et al. 1991; Huke, 1993]. However, these all introduce additional technical complications which are irrelevant to the main theme of this chapter and hence we shall not consider them further. We assume that the state of the system at time n , denoted $x_n \in M$ evolves according to $x_{n+1} = f(x_n)$ for some smooth invertible map $f : M \rightarrow M$ (note that the invertibility of f is essential). The system is observed using a smooth measurement function $\varphi : M \rightarrow \mathbb{R}$ giving a scalar time series $\varphi_n = \varphi(x_n)$. The aim of the so-called method of delays is to reconstruct the state space M and the dynamics f from the time series φ_n . because M is high-dimensional and each component of φ_n is only one-dimensional, it is clear that to obtain a suitable state space we need to somehow group different elements of the time series. The most natural (but not the only) way of doing this is to take successive φ_n to create a vector:

$$y_n = (\varphi_n, \varphi_{n+1}, \dots, \varphi_{n+d-1}). \quad (4.1)$$

The number of components d that we use is usually referred to as the *embedding dimension*. There is an obvious generalization to multivariate time series obtained from a measurement function $\varphi : M \rightarrow \mathbb{R}^k$. Although

no published statement of a corresponding extension of Takens' theorem exists, it is clear from the proof of the existing theorem that such an extension holds.

The set of possible $y_n \in \mathbb{R}^d$ obtained from (4.1) provides a candidate for the state space of our reconstructed system. The natural dynamics to impose on this is one that sends y_n to y_{n+1} , i.e., it attempts to define a map F by $F(y_n) = y_{n+1}$. This raises two questions: Is F well defined, and if so, does it have anything to do with the original dynamical system f ? The first question is not as trivial as it might seem because it requires that $y_{n+1} = y_{p+1}$ whenever $y_n = y_p$, which is far from obvious. In terms of the time series this condition amounts to $\varphi_{n+d} = \varphi_{p+d}$ whenever $\varphi_n = \varphi_p, \dots, \varphi_{n+d-1} = \varphi_{p+d-1}$ and is thus equivalent to the time series being perfectly predictable. In this context one might also ask about the regularity of F , i.e., is it continuous, smooth, etc, or if $\varphi_n, \dots, \varphi_{n+d-1}$ is close to $\varphi_p, \dots, \varphi_{p+d-1}$ is the same true for φ_{n+d} and φ_{p+d} ? This will obviously have a profound effect on the sensitivity of our predictions to noise or numerical inaccuracies.

While a positive answer to this question may be sufficient if all we want to do is predict the time series, the second question becomes paramount if we also want to characterize the time series using quantities such as fractal dimensions or Lyapunov exponents or if we want to deduce something about the mechanisms that created the time series. It also arises if we want to compare two time series generated by the same system, but observed through two different measurement functions.

Fortunately, both questions are answered in the affirmative by Takens' so-called *embedding theorem*. Informally, this ensures that for typical systems and observations, and for $d \geq 2m + 1$, the reconstructed dynamics F is well defined and as smooth as the original dynamics f and F is equivalent to f under a smooth (but unknown) coordinate change. Because most of nonlinear dynamics is concerned with coordinate free properties (such as fractal dimensions or Lyapunov exponents) this ensures that these are the same for both the original and reconstructed systems. The coordinate change between f and F is given by the *delay map* $\Phi : M \rightarrow \mathbb{R}^d$ defined by

$$\Phi(x) = (\varphi(x), \varphi(f(x)), \dots, \varphi(f^{d-1}(x))). \quad (4.2)$$

Note that this ensures that $\Phi(x_n) = y_n$. Stated precisely, Takens' Theorem says that for a dense open set of (f, φ) (in the standard topology on the space of all such maps and measurement functions) the delay map Φ is a smooth embedding, that is it is 1 - 1 and an immersion. This means that $\Phi(x) \neq \Phi(x')$ for all $x \neq x'$ and $D_x \Phi.v \neq 0$ for all $v \in T_x M$ such that $v \neq 0$, and implies that F is invertible on its image $\Phi(M)$, and its inverse is smooth. Note that the Sauer et al. [1991] version of the theorem gives merely a topological embedding of the attractor (with Φ smooth by definition, but Φ^{-1} not necessarily so in all directions) but has the benefit that d need only be greater than twice the fractal dimension of the attractor.

Because Φ is invertible on $\Phi(M)$, we may define the smooth dynamical system $\Phi \circ f \circ \Phi^{-1} : \Phi(M) \rightarrow \Phi(M)$ on $\Phi(M) \subset \mathbb{R}^d$. A simple unraveling of the definitions shows that $\Phi \circ f \circ \Phi^{-1}(y_n) = y_{n+1}$ and hence $\Phi \circ f \circ \Phi^{-1}$ is precisely the reconstructed map F described earlier. In the theory of nonlinear dynamics a relation of the form $\Phi \circ f \circ \Phi^{-1} = F$ is called a (smooth) conjugacy and is a formal way of stating that F is equivalent to f under a smooth coordinate change. The whole process of delay reconstruction can thus be represented graphically as in Figure 4.2. Somewhat incorrectly, this




FIGURE 4.2. Overview of delay embedding for deterministic systems

has come to be known as *delay embedding*, even when Φ might not actually be an embedding in the mathematical sense. We prefer to avoid this here and use the expression *delay reconstruction*, thereby retaining the rigorous topological meaning for the term *embedding*.

Finally, observe that F consists of d components, so that $F(y_n) = (F_1(y_n), \dots, F_d(y_n))$, with $F_i : \Phi(M) \rightarrow \mathbb{R}$. Writing F_i in terms of the time series, we obtain $F_i(\varphi_n, \dots, \varphi_{n+d-1}) = \varphi_{n+i}$. Thus the first $d-1$ components F_1, \dots, F_{d-1} are trivial and simply consist of copying an argument. The only nontrivial part of F is the last component F_d , which for convenience we will denote by G , so that

$$\varphi_{n+d} = G(\varphi_n, \dots, \varphi_{n+d-1}). \quad (4.3)$$

This simply formalizes the statement that if F is well defined for a given d then the time series is predictable in terms of its d previous values. Although in principle the relation $F = \Phi \circ f \circ \Phi^{-1}$ gives an explicit formula for G , this is in terms of f and φ , which in most practical applications will not be known. However, the last two decades have seen the development of a

number of algorithms for computing an approximation to G from a sample of the time series (see other chapters in this volume, or Ott et al. [1994], Abarbanel [1995], Kantz and Schreiber, [1999]).

4.3 Delay Reconstruction for Stochastic Systems

4.3.1 *Random Dynamical Systems*

In order to extend the preceding framework to stochastic systems we first need to decide how to incorporate noise in our original unknown system f . A convenient formalism for this, which encompasses a wide class of noisy systems, is that of *random dynamical systems* (see, e.g., Arnold [1998]) which we can think of as deterministic systems driven by a random process ω_n . For simplicity we shall assume that this process is Bernoulli, so that successive ω_n at each time step n are chosen independently with respect to some probability measure μ . The state x_n then evolves according to

$$x_{n+1} = f(x_n, \omega_n). \quad (4.4)$$

If we think of ω_n as a parameter then we can interpret this as a standard dynamical system with noise on the parameters. It can also be helpful to write $f_{\omega_n}(x_n)$ instead of $f(x_n, \omega_n)$. This suggests the interpretation that instead of applying the same function f every time, we choose a different function f_{ω_n} at random at each time step, so that in the most general setting μ can be thought of as a measure on the space of all dynamical systems on M . If this measure consists of a finite number of atoms, so that f_{ω_n} is chosen from a finite set of maps, we obtain a so called *iterated function system* (see, e.g., Norman [1968], Barnsley [1988]). On the other hand, the case of a single deterministic system f subject to additive (dynamical) noise can be included in this formalism by setting $f_{\omega_n}(x_n) = f(x_n) + \omega_n$. It is also possible to add noise to the observations by making the measurement function φ depend on either the same noise process ω_n , or an independent process ξ_n . Because this merely adds technical and notational complexity without introducing any significantly new concepts, we shall not explore this further in this chapter.

From a dynamical systems point of view the best approach to the study of (4.4) is to expand the state space sufficiently to give an autonomous deterministic dynamical system on (x_n, ω_n) . This is analogous to the way that periodically forced differential equations can be treated as autonomous systems by the addition of a dummy variable to represent time, thereby increasing the state space dimension by 1. In our case we end up with an infinite-dimensional state space. Thus suppose that ω_n is drawn from a space N and let $\Sigma = N^{\mathbb{Z}}$ be the space of bi-infinite sequences $\omega = (\dots, \omega_{-1}, \omega_0, \omega_1, \dots)$ of elements of N . Define the usual shift operator σ :




FIGURE 4.3. Graphical representation of a random dynamical system.

$\Sigma \rightarrow \Sigma$ by $[\sigma(\omega)]_n = \omega_{n+1}$. Then the evolution of x_n given by (4.4) can be represented by the skew product $T : M \times \Sigma \rightarrow M \times \Sigma$ (see Figure 4.3) defined by

$$T(x, \omega) = (f(x, \omega_0), \sigma(\omega)). \tag{4.5}$$

Note that the product measure μ_Σ on Σ derived from μ is σ -invariant. We can also consider general σ -invariant measures to take account of correlations in the choice of successive ω_n (so that, for example, ω is a Markov process), but for simplicity we shall restrict ourselves to the Bernoulli case. If we dispense with the measure μ the same formalism can also be used to model a deterministic system driven by an arbitrary input sequence ω . This arises frequently in communications systems where the sequence ω would represent the information being transmitted (see, e.g., Broomhead et al. [1999]). Another application is to irregularly sampled time series where ω_n denotes the time between sample n and $n + 1$ [Martin, 1998].

4.3.2 Conjugacy for Random Dynamical Systems

As before, we assume that the observed time series is generated by $\varphi_n = \varphi(x_n)$ for some measurement function $\varphi : M \rightarrow \mathbb{R}$. The crucial question we then need to address in order to develop an analogue of Takens' theorem is what we mean by a delay reconstruction of T . As we saw in Section 4.2 the fundamental property of a reconstruction is that F and f should be equivalent under a coordinate change, in other words that $F = \Phi \circ f \circ \Phi^{-1}$. We thus need to ask what it means for two random dynamical systems T and T' to be equivalent in this way. The most general concept is simply to require $T' = H \circ T \circ H^{-1}$ for some (invertible) coordinate change $H :$

$M \times \Sigma \rightarrow M \times \Sigma$. Because we are in a probabilistic setting, we only require this relationship to hold for μ_Σ -almost every ω .

It turns out, however, to be convenient to place some further restrictions on H . The space $M \times \Sigma$ is infinite-dimensional, and hence we have little hope of reconstructing it with a finite-dimensional delay space \mathbb{R}^d . Furthermore φ is a function of M only and attempting to reconstruct Σ using φ seems foolhardy (but note that in the case of deterministic forcing, this is in fact possible [Stark, 1999]). In some sense we therefore want to restrict the reconstruction procedure to M . A reasonable interpretation of this idea is to require that the Σ component of H is the identity, that is, $H(x, \omega) = (h(x, \omega), \omega)$ for some map $h : M \times \Sigma \rightarrow M$. If H is to be invertible for μ_Σ -almost every ω then $h_\omega = h(\bullet, \omega) : M \rightarrow M$ has to be invertible for μ_Σ -almost every ω . Coordinate changes of this form are common in the theory of random dynamical systems (see, e.g., Arnold [1998]).

For convenience define $\tilde{f} : M \times \Sigma \rightarrow M$ by $\tilde{f}(x, \omega) = f(x, \omega_0)$ so that $T = (\tilde{f}, \sigma)$. If similarly $T' = (\tilde{f}', \sigma)$ then writing $T' = H \circ T \circ H^{-1}$ in component form gives

$$(\tilde{f}', \sigma) = (h, Id) \circ (\tilde{f}, \sigma) \circ (h, Id)^{-1} \quad (4.6)$$

If we denote $\tilde{f}_\omega = \tilde{f}(\bullet, \omega) : M \rightarrow M$ we have

$$\begin{aligned} (h, Id) \circ (\tilde{f}, \sigma) \circ (h, Id)^{-1}(\bullet, \omega) &= (h, Id) \circ (\tilde{f}_\omega \circ h_\omega^{-1}, \sigma(\omega)) \\ &= (h_{\sigma(\omega)} \circ \tilde{f}_\omega \circ h_\omega^{-1}, \sigma(\omega)) \end{aligned}$$

Thus the first component of equation (4.6) is

$$\tilde{f}'_\omega = h_{\sigma(\omega)} \circ \tilde{f}_\omega \circ h_\omega^{-1} \quad (4.7)$$

where as usual $\tilde{f}'_\omega = \tilde{f}'(\bullet, \omega) : M \rightarrow M$. Note the similarity of this to the deterministic conjugacy $f' = h \circ f \circ h^{-1}$. Essentially all we do in the random case is index both the dynamics and the coordinate change with ω . The only slightly delicate point is the σ appearing in $h_{\sigma(\omega)}$. The reason for this is that by the time we come to apply h in equation (4.7) we have carried out one time step of the dynamics, and hence ω has moved to $\sigma(\omega)$. (See Figure 4.4).

4.3.3 Stochastic Takens Embedding Theorem

Given such a definition it seems reasonable to attempt to generalize Takens' theorem to random dynamical systems by requiring that the delay map be an invertible coordinate transformation in the preceding sense for μ_Σ -almost every ω . To formulate an appropriate theorem it thus only remains to give a rigorous definition of the delay map. From (4.7) it is clear that this will depend on ω , and by analogy with the deterministic case (i.e., equation 4.1) it should satisfy $\Phi(x_n, \sigma^n(\omega)) = y_n = (\varphi_n, \varphi_{n+1}, \dots, \varphi_{n+d-1})$. Note




FIGURE 4.4. Coordinate change (conjugacy) for random dynamical systems.

that $\sigma^n(\omega)$ is the correct argument of Φ because this corresponds to the choice of noise at time n . This yields (*cf* equation 4.2)

$$\Phi_\omega(x) = (\varphi(x), \varphi(f_{\omega_0}(x)), \varphi(f_{\omega_1\omega_0}(x)), \dots, \varphi(f_{\omega_{d-2}\dots\omega_0}(x))), \quad (4.8)$$

where $f_{\omega_i\dots\omega_0} = f_{\omega_i} \circ \dots \circ f_{\omega_0}$. We see that Φ_ω does not just depend on ω_0 , which is the case with f_ω , but on $d-1$ successive values $\omega_0, \dots, \omega_{d-2}$. This is entirely reasonable because the delay reconstruction process requires the dynamical system to evolve over $d-1$ time steps and hence to be subject to the effects of $d-1$ values of the noise. It causes us no difficulty, and in particular equations (4.6) and (4.7) do not need any modification. The reconstructed map is defined by

$$F_\omega = \Phi_{\sigma(\omega)} \circ \tilde{f}_\omega \circ \Phi_\omega^{-1} \quad (4.9)$$

for all those ω such that Φ_ω is invertible. Note that $f_\omega : \Phi_\omega(M) \rightarrow \Phi_{\sigma(\omega)}(M)$ and hence it is not just the map F_ω itself that is random, but its domain and range also. From equation (4.6) we have

$$\begin{aligned} (F, \sigma)(y_n, \sigma^n(\omega)) &= (h, Id) \circ (\tilde{f}, \sigma) \circ (h, Id)^{-1}(y_n, \sigma^n(\omega)) \\ &= (h, Id) \circ (\tilde{f}, \sigma)(x_n, \sigma^n(\omega)) \\ &= (h, Id)(x_{n+1}, \sigma^{n+1}(\omega)) \\ &= (y_{n+1}, \sigma^{n+1}(\omega)) \end{aligned}$$

or in other words,

$$F_{\sigma_n(\omega)}(y_n) = y_{n+1}.$$

4.3.4 The Time Series Model

In terms of the time series this can be written as

$$F_{\sigma_n(\omega)}(\varphi_n, \varphi_{n+1}, \dots, \varphi_{n+d-1}) = (\varphi_{n+1}, \varphi_{n+2}, \dots, \varphi_{n+d}).$$

As in the derivation of equation (4.3), the first $d-1$ components of F_ω are trivial, while the last, which we shall denote G_ω , gives the time series model

$$\varphi_{n+d} = G_{\sigma_n(\omega)}(\varphi_n, \varphi_{n+1}, \dots, \varphi_{n+d-1}). \quad (4.10)$$

At an informal level, therefore, a reconstruction framework for random dynamical systems is virtually identical to the standard deterministic one, apart from the fact that all its constituent objects are random, i.e., indexed by ω . This is a particularly appealing point of view if we regard f_ω as a parameterized family of deterministic systems with noise on the parameters. In such a case the reconstructed map F_ω is an equivalent family,

with the same realization of the same stochastic process on its parameters F_ω . Note however one subtle difference between F_ω and f_ω that is not apparent from the notation. Because Φ_ω depends on $\omega_0, \dots, \omega_{d-2}$ and hence $\Phi_{\sigma(\omega)}$ depends on $\omega_1, \dots, \omega_{d-1}$, equation (4.9) implies that F_ω depends on $\omega_0, \dots, \omega_{d-1}$. Thus $F_{\sigma_n(\omega)}$ and $G_{\sigma_n(\omega)}$ depend on $\omega_n, \dots, \omega_{n+d-1}$ and the time series model in equation (4.10) can also be written as

$$\varphi_{n+d} = G(\varphi_n, \varphi_{n+1}, \dots, \varphi_{n+d-1}; \omega_n, \dots, \omega_{n+d-1}). \quad (4.11)$$

From a time series point of view this equation is the main outcome of our stochastic delay reconstruction framework. It suggests the kinds of models that should be used as starting points for a statistical analysis of a time series believed to have been generated by a stochastic nonlinear system. Although some authors (e.g., Chen and Billings [1989]) have already used models of the form (4.11) (with ω_n restricted to one dimension), more commonly it is assumed that there is only a single additive random component:

$$\varphi_{n+d} = G(\varphi_n, \varphi_{n+1}, \dots, \varphi_{n+d-1}) + \omega_{n+d-1}$$

It is clear that estimating both G and ω in a model of the form (4.11) is a major challenge. Note that $\omega_n, \dots, \omega_{n+d-2}$ have all been determined by time $n + d - 1$, and so there is at least some hope of estimating them from previous values of the time series. By contrast ω_{n+d-1} corresponds to new uncertainty entering the system in the time step from $n + d - 1$ to $n + d$. The special case

$$\varphi_{n+d} = G(\varphi_n - \omega_{n-1}, \varphi_{n+1} - \omega_n, \dots, \varphi_{n+d-1} - \omega_{n+d-2}) + \omega_{n+d-1}$$

is essentially equivalent to so-called *shadowing* or *noise reduction* for chaotic systems (e.g., see Ott et al. [1994]). Note that the appearance of the additional term ω_{n-1} in this equation makes little fundamental difference, it would, for example, arise in (4.11) if in our original dynamics (4.4) the map f depended on ω_{n-1} rather than just ω_n . While a variety of well-developed techniques exist to carry out shadowing, extending these to more general systems of the form (4.11) seems difficult.

4.3.5 Technical Conditions

The whole of this framework relies on Φ_ω being invertible for μ_Σ -almost every ω . Stark et al. [1999] show that if $d \geq 2m + 1$ this is indeed the case for generic random dynamical systems and observations under suitable technical conditions on ω (see also Stark et al. [1997]). These require N (the space from which the ω_n are drawn) to be a compact smooth manifold, for f_ω to depend on ω_n smoothly and for the marginal measure on N^d to be absolutely continuous with respect to Lebesgue measure. In the Bernoulli case, the latter amounts to μ being absolutely continuous with respect to Lebesgue measure on N .

We can motivate such a condition as follows: We know from the deterministic embedding theory that there are maps f for which the delay map Φ is not an embedding. Suppose that such a map corresponded to a $\omega_0 \in N$ and $\mu(\{\omega_0\}) > 0$. The probability of choosing f_{ω_0} on d successive occasions would then be nonzero. The corresponding Φ_{ω} would thus occur with nonzero probability with respect to μ_{Σ} . But such a Φ_{ω} is identical to the deterministic Φ obtained from f_{ω_0} , and hence is not an embedding. The same argument applies if μ assigns a nonzero probability to the set of f_{ω_0} that fails to give a delay embedding. Because the generic part of Takens theorem ultimately relies on Smale's generalization of Sard's theorem it is not difficult to see that this set has Lebesgue measure zero [Stark, 1999]. Hence, if we want Φ_{ω} to be an embedding for μ_{Σ} -almost every ω we need μ to assign measure zero to the set of such f_{ω_0} , and the most elegant hypothesis that achieves this is to assume absolute continuity with respect to Lebesgue. The same underlying idea generalizes to the case where we choose a different f_{ω_0} at each time step. It is thus intuitively reasonable that we need some condition that ensures that μ_{Σ} is not too singular.

By contrast, the conditions on N and the smooth dependence on ω are technical ones, imposed by the method of proof (which uses differential topological methods developed by Stark [1999]). Both conditions are restrictive. In particular, the former requires the noise in our system to be bounded and hence excludes the most common case, namely that of Gaussian noise. On the other hand there are now beginning to be some suggestions that time series with unbounded noise have significantly different properties to those with bounded noise (e.g., the papers by Lalley and by Guegan in this volume). Turning to the issue of smoothness, most of the theory of random dynamical systems assumes far weaker regularity for ω (see, e.g., Arnold [1998]), though in fact most concrete examples of random maps (as opposed to stochastic differential equations) do satisfy this condition.

4.4 Trading-Off Noise and Complexity

Up to now we have treated the unknown system f as a predetermined object in the real world whose properties are fixed. This implies that there is a "correct" outcome of any reconstruction procedure and a "correct" value for quantities such as the minimal embedding dimension. If, however, we allow ourselves to speculatively extend the stochastic reconstruction framework from the last section we arrive at an entirely different point of view which permits us to trade off the complexity of f against the amount of noise in the system. The "correct" unknown system is thus replaced by a spectrum of possible models of increasing complexity and our task becomes choosing the one that is most appropriate for our application. There is

thus, for example, no longer a “correct” embedding dimension, but rather we may choose between a small dimension giving a very noisy system or a large dimension and a much less noisy system. Such a perspective is a much better starting point for a proper statistical analysis of delay reconstruction and is already implicit in much of the recent work of Mees and his group on using complexity measures to choose good nonlinear time series models (see, e.g., Mees [1993], [Judd and Mees, 1995], [Judd and Mees, 1998] and the chapter by Judd et al. in this volume). In the context of spatio-temporal systems, a formal framework of this kind was first described by Ørstavik and Stark [1998], although something similar was hinted at in the last two sentences of Kantz and Olbrich [1997].

We begin with the informal idea that chaotic behavior, particularly if it is high-dimensional, can be virtually indistinguishable from genuine stochasticity. Thus, if we have a system such as (4.4) for which x exhibits complex dynamics, it should be possible to ascribe part, or even all, of its behavior to the random part of the model ω . For simplicity, let us first consider this in a purely deterministic model. Suppose we can decompose the state space of this into a product $M \times L$, so that the state variables are (x, z) with $x \in M$, $z \in L$. If we use standard deterministic delay reconstruction, as in Section 4.2, then Takens theorem requires us to have an embedding dimension of $2m + 2l + 1$, where $m = \dim M$ and $l = \dim L$. On the other hand, if we can somehow regard z as a random variable and perform a reconstruction in the sense of Section 4.3, then an embedding dimension of $2m + 1$ will suffice. The trade-off is of course that we now have a noisy model and have to use (4.11) instead of (4.3).

Let us suppose that the dynamics on $M \times L$ are given by

$$\begin{aligned}x_{n+1} &= f(x_n, z_n) \\z_{n+1} &= g(x_n, z_n).\end{aligned}\tag{4.12}$$

In order to treat z_n as a random variable, we have to replace g by the shift σ on an appropriate space and obtain a map of the form (4.5). This can be done by defining

$$X = \{(x, \eta) \in M \times L^{\mathbb{Z}} : \eta_n = g^{(n)}(x, \eta_0) \text{ for all } n \in \mathbb{Z}\},$$

where $g^{(n)}$ is the second component of $(f, g)^n$. If we define $T : M \times L^{\mathbb{Z}} \rightarrow M \times L^{\mathbb{Z}}$ the same way as in (4.5), that is

$$T(x, \eta) = (f(x, \eta_0), \sigma(\eta)),\tag{4.13}$$

then we claim that X is a T -invariant subset of $M \times L^{\mathbb{Z}}$. To see this, let $h : M \times L \rightarrow M \times L^{\mathbb{Z}}$ be the map $h = (Id, (\dots, g^{(n-1)}, g^{(n)}, g^{(n+1)}, \dots))$, so that $h(x, z) = (x, \eta)$ with $\eta_n = g^{(n)}(x, z)$. Thus $X = h(M \times L)$ and since h is obviously 1 – 1, it is invertible on X . In fact, taking the product topology on $M \times L^{\mathbb{Z}}$, h is a homeomorphism between $M \times L$ and

X , though we do not need to use this here. Now observe that if we denote the first component of $(f, g)^n$ by $f^{(n)}$, so that $(f^{(n)}, g^{(n)}) = (f, g)^n$, then $(f^{(n)}(f(x, z), g(x, z)), g^{(n)}(f(x, z), g(x, z))) = (f, g)^n \circ (f, g)(x, z) = (f, g)^{n+1}(x, z) = (f^{(n+1)}(x, z), g^{(n+1)}(x, z))$. Thus $g^{(n)}(f(x, z), g(x, z)) = g^{(n+1)}(x, z)$ and hence

$$\begin{aligned} h \circ (f, g) \circ h^{-1}(x, \eta) &= h \circ (f, g)(x, \eta_0) \\ &= h(f(x, \eta_0), g(x, \eta_0)) \\ &= (f(x, \eta_0), (\dots, g^{(n)}(f(x, \eta_0), g(x, \eta_0)), \dots)) \\ &= (f(x, \eta_0), (\dots, g^{(n+1)}(x, \eta_0), \dots)) \\ &= (f(x, \eta_0), \sigma(\eta)). \end{aligned}$$

Therefore

$$T = h \circ (f, g) \circ h^{-1},$$

so that not only is $T(X) \subset X$, but in fact T and (f, g) are conjugate, so that there is a 1 – 1 correspondence between their orbits. Thus if we have a measurement function $\varphi : M \times L \rightarrow \mathbb{R}$ on our original system (4.12) and a resulting time series $\varphi_n = \varphi((f, g)^n(x, z))$, then we get an identical time series from an orbit of the skew product (4.13) with initial condition $h(x, z)$. More precisely if we define the observable $\tilde{\varphi} : M \times L^{\mathbb{Z}} \rightarrow \mathbb{R}$ by $\tilde{\varphi}(x, \eta) = \varphi(x, \eta_0)$ then $\tilde{\varphi}(x, \eta) = \varphi(h^{-1}(x, \eta))$ and thus $\tilde{\varphi}(T^n(h(x, z))) = \varphi(h^{-1} \circ T^n \circ h(x, z)) = \varphi((f, g)^n(x, z))$. The only difference compared to Section 4.3 is that $\tilde{\varphi}$ depends on η , but as indicated there, this causes no difficulties. If the original system also has a genuine stochastic component ω we can group η and ω together and regard x as the deterministic part of the system and (η, ω) as the random part.

We thus see that we can take a deterministic system, treat some of its degrees of freedom as a stochastic process and arrive at a framework that is outwardly very similar to the one developed for random systems in Section 4.3. The only significant change is that instead of iterating T on the whole of $M \times L^{\mathbb{Z}}$ we restrict ourselves to a closed invariant set X . Unfortunately, this apparently innocuous modification is potentially rather problematical. This is because for points in X there is an implicit dependence of η on x : once we know x and η_0 , we know the whole of η . Thus, for example, if we again define a delay map Φ_η using (4.8), then $\Phi_\eta(x)$ depends on x both explicitly via its argument, and indirectly through η . The same holds for the reconstructed map (4.9) and the time series model (4.11). Furthermore, because η is no longer an autonomous stochastic process it is not clear what a reconstruction like (4.9) actually means. Nevertheless there does seem to be some numerical evidence, described in the next section, that the speculative framework presented here offers some insight into time series obtained from high-dimensional deterministic systems.

Of course, the ideas presented in this section are purely conjectural; there is no embedding theorem for this case analogous to those in Sections 4.2

and 4.3. The only exception is where g is independent of x , where the Bundle Delay Embedding Theorem (Theorem 3.2) of Stark [1999] gives exactly the required result. However, such a situation is much simpler since η then no longer depends on x . It seems plausible that the techniques developed by Stark [1999] could be used to prove an embedding theorem for general g but the technical calculations required are daunting. It is therefore probably not worth attempting such a proof, until more numerical evidence for the usefulness of the approach outlined in this section has been accumulated.

4.5 Reconstruction for Spatio-Temporal Systems

One class of systems where it is now clear that there is scope for trading off noise against complexity is that of spatio-temporal systems. By this we mean systems in which spatially distinct parts can evolve in time in distinct ways. We can thus think of such systems as consisting of many local dynamical systems, one at each spatial location, coupled into one large system. They arise in many applications throughout science and engineering including most types of fluid flow, pattern formation in chemical and biological systems, dynamics of ecosystems, road traffic, vibration of structures such as beams, plates and shells, and many others. In some cases (e.g., weather prediction) there is a reasonable understanding of the underlying deterministic mechanisms governing the evolution of the system, and we can use this to construct an a priori state space model to which observations can be fitted. In many other situations this is not possible and we need to reconstruct the unknown dynamics from observed data in a similar fashion to that described in Section 4.2.

This immediately leads to a number of problems, both theoretical and practical. Spatio-temporal dynamical systems often exhibit various symmetries; typically they are transitionally invariant. As a result they may fail to satisfy the genericity assumptions of Takens theorem. More seriously, observables often depend on a single spatial location or at best a small spatial neighborhood. They are thus very atypical in the space of all measurement functions on the state space and any appeal to Takens theorem to justify reconstruction is therefore highly suspect. Common sense also suggests that it is unreasonable to expect to reconstruct such a system by using a sufficient number of delays of a single observable. The amount and quality of data obtained by observing a single spatial location will almost certainly be insufficient to reconstruct the behavior of the whole spatially extended system. Any serious approach to the reconstruction of spatio-temporal systems will therefore inevitably use multiple measurements, distributed in an appropriate way around the system. Delay reconstruction techniques are increasingly being applied to such data (see, e.g., Muldoon et al. [1994],

Little et al. [1996], Rand and Wilson [1997], Ørstavik and Stark, [1998]).

However, simply taking a large number of measurement functions is not a complete solution. Unless the dynamics of the system asymptotically collapses down to a low-dimensional attractor (for example, if there is an inertial manifold) we will need to use a high-dimensional reconstruction space, that is a high embedding dimension. This leads to the so-called *curse of dimensionality*. Whatever method of function approximation (e.g., local polynomial, radial basis function, neural network, kernel density *etc*) we use to approximate the reconstructed dynamics, we need to have a reasonable spread of data in the region of interest. Thus, for instance, a local linear or quadratic approach estimates the value of a function at a point $y \in \mathbb{R}^d$ by performing a linear least squares regression in a neighborhood of y . Such a neighborhood must contain a sufficient number of data points to yield a meaningful estimate and yet not be so large that the function is no longer linear (or quadratic respectively). As d grows, more data is necessary to ensure that sufficient numbers of neighbors can be found. Similar problems plague other approximation methods and for d much larger than 10 or so it becomes almost completely impractical to obtain meaningful estimates for the reconstructed dynamics.

Despite this, in many circumstances it turns out to be possible to make useful predictions for time series from genuinely high-dimensional spatio-temporal systems. A particularly relevant example, which ultimately motivated the discussion in Section 4.4, is given by Ørstavik and Stark [1998]. This investigated the predictability, using various reconstruction schemes, of time series from a *coupled map lattice*. This consists of a collection of maps coupled together on a lattice (Figure 4.5a) and provides perhaps the simplest context in which to study complex spatio-temporal dynamics. In this case $N = 100$ lattice sites were used with periodic boundary conditions. The local state space at each site was one-dimensional giving an overall state space dimension $m = 100$. The dynamics at each site was given by the usual logistic map and the coupling was nearest-neighbor. If we denote the state of lattice site i at time n by $x_n^{(i)}$ (this is consistent with the remainder of this chapter but is the exact opposite of the notation used by Ørstavik and Stark [1998]) then the dynamics at each site is given by

$$x_{n+1}^{(i)} = (1 - \varepsilon)f(x_n^{(i)}) + \frac{1}{2}\varepsilon f(x_n^{(i-1)}) + \frac{1}{2}\varepsilon f(x_n^{(i+1)}),$$

where ε is a coupling parameter and $f : [0, 1] \rightarrow [0, 1]$ is the logistic map $f(x) = 1 - ax^2$. This system is known to exhibit spatio-temporal chaos for wide ranges of parameters, and at the particular parameters used for this study the attractor was estimated to have a Lyapunov dimension of approximately 70. Thus Takens' original criterion $d \geq 2m + 1$ gives a minimal embedding dimension of $d = 201$, while even Sauer et al. [1991] require $d \approx 140$. As indicated earlier, embedding dimensions that large are completely impractical.

Rather surprisingly, however, Ørstavik and Stark [1998] found that it was possible to make useful predictions in low embedding dimensions, and in

a




FIGURE 4.5. a) Coupled map lattice; b) truncated coupled map lattice.

fact the best predictions were obtained in $d = 4$. The false nearest neighbor algorithm for determining an embedding dimension (see Abarbanel [1995] or Kantz and Schreiber [1999]) also indicated a value of 4. This is not surprising given that this technique essentially looks for optimal prediction using a locally constant predictor and illustrates the dangers of relying too heavily on such algorithms. On the other hand, the scaling behavior for the prediction error was more reminiscent of a stochastic system than a deterministic one. This suggests that from a time series perspective this system looks more like a low-dimensional deterministic system driven by noise than the high-dimensional noise-free system that it really is. This is consistent with the folklore view that high-dimensional chaos is indistinguishable from “genuinely” random behavior.

Additional evidence for such a point of view is provided by the fact that it is possible to obtain extremely good estimates of the Lyapunov spectrum of the whole lattice by appropriately rescaling the spectrum computed from even quite a small sub-system (Carretero-González et al. [1999a], and references therein). The explanation for this is that the whole lattice is acting as if it were made up of a number of weakly interacting smaller systems. Thus if our observations are local we shall primarily see just one such small system, and all the remaining ones appear as a broad global background.

This is the structure that we had in mind in Section 4.4; if the speculative framework outlined there is indeed correct then we can reconstruct the lattice as a low-dimensional stochastically driven system, using a small embedding dimension. This would explain the prediction results of Ørstavik and Stark [1998].

A more direct test of such a viewpoint can be obtained by studying truncated lattices. This is done by Carretero-González et al. [1999b]), where observations obtained from a small stochastically driven lattice (Figure 4.5b) are compared with those from the full deterministic lattice of size $N = 100$. This shows that even for large coupling, time series from lattices as small




FIGURE 4.6. Invariant probability density at a central site for full ($N = 100$) and truncated lattices; a) $\varepsilon = 0.2$, b) $\varepsilon = 0.8$. The truncated lattice was driven at each boundary by independent i.i.d. processes with uniform density in $[0, 1]$.

as ten sites or so are indistinguishable from those generated by the full lattice. Thus, for example, if we compare invariant probability densities we find that small truncated lattices gives remarkably good approximations to the invariant density for the full lattice (Figure 4.6). This is despite the fact that the processes driving the boundaries were given by simple white noise with uniform distribution in $[0, 1]$ and no attempt to reproduce the correct statistics of these processes was made. If this is done, even better fits can be obtained.

The difference between the densities for the truncated and full lattices decays exponentially with the truncation size (Figure 4.7a), reaching a plateau determined by the discretization and finite data size used to estimate the density. Similar behavior occurs with the power spectrum (Figure 4.7b) and with two-point correlation functions [Carretero-González et al. 1999b]. Of course such measures only evaluate linear relationships in the data and so it is perhaps not that surprising that they fail to distinguish between very high dimensional dynamics and stochastic behavior. A more serious test is provided by predictability (Figure 4.8). Once again, it turns out that as the




FIGURE 4.7. L_1 distance between a) invariant densities and b) power spectra for full and truncated lattices. Densities were calculated using 100 bins and a total of 10^8 data points (100 orbits of length 10^6). Spectra were computed by averaging 10^6 spectra of 1024 points each.




FIGURE 4.8. Comparison of normalized root-mean square one step prediction error for full and truncated lattices; a) $\epsilon = 0.2$, b) $\epsilon = 0.8$. In all cases local linear fits were made to a data set of 20,000 points using a spatio-temporal reconstruction in embedding dimension 4, and errors were calculated for a further out-of-sample 5000 points.

lattice size grows the time series from the stochastically driven truncated lattice rapidly reach the same degree of prediction error as the full deterministic lattice. Indeed, for low coupling ($\varepsilon = 0.2$), a stochastically driven lattice of size 3 is indistinguishable in this context from the full lattice, and even for very high coupling parameters ($\varepsilon = 0.8$) truncated lattices of size 15 are sufficient.

These results strongly suggest that from a time series perspective many large spatio-temporal systems should not be viewed as a single deterministic system, but rather as a collection of loosely coupled low-dimensional systems. When making localized observations the effects of the other unseen local systems can be effectively described by a simple stochastic term. This allows us to build relatively simple effective models, rather than working in impracticably large embedding dimensions. If we are able to observe the system at many spatial locations, then we should attempt to model the resulting multivariate time series by a collection of more or less independent low-dimensional systems, with perhaps some allowance for weak correlations. Where we have good reason to believe that the system is spatially homogeneous, we can fit one such local model using data obtained from many spatially locations, thereby significantly increasing the amount of data available to fit the model.

4.6 Conclusions

In this chapter we have described a framework and an associated embedding theorem, for the delay reconstruction of random dynamical systems. By only attempting to reconstruct the deterministic part of the dynamics, this allows us to obtain a full reconstruction using the same embedding dimension as for a deterministic system, despite the fact that a stochastic system is in a sense infinite-dimensional.

We have then gone on to speculate that we may have a substantial amount of choice in which parts of the original unknown system we regard as deterministic and which parts as stochastic. Emphasizing the deterministic part leads to a less noisy system at the cost of a higher embedding dimension and vice versa. Numerical evidence that such a framework may be useful is provided by recent studies of simple coupled map lattices, where it has been observed that a large-dimensional deterministic lattice is indistinguishable from a small truncated lattice driven by noise at its boundaries.

We therefore believe that rather than being seen as a procedure to recreate a single “correct” system, delay reconstruction should be regarded as a method of choosing from among a spectrum of noisy systems of differing dimensions and noise levels. The choice is a statistical one, trading off complexity against levels of noise. Developing appropriate statistical methods

to embody such a philosophy is a major challenge and will require close co-operation between statisticians and those working in nonlinear dynamics.

Acknowledgments

Many people have contributed to the development of the ideas discussed here. I would particularly like to thank David Broomhead, Jerry Huke, Sakse Ørstavik and Ricardo Carretero-González. The work on spatio-temporal systems described here was partly carried out under grant GR/L42513 from the UK Engineering and Physical Sciences Research Council, and I would also like to express my gratitude to the Royal Commission for the Exhibition of 1851 for additional financial support.

References

- Abarbanel HDI, 1995, *Analysis of Observed Chaotic Data*, Springer-Verlag, New-York.
- Aeyels D, 1981, Generic Observability of Differentiable Systems, *SIAM J. Control and Optimization*, 19, 595–603.
- Arnold, L., 1998, *Random Dynamical Systems*, Springer-Verlag, New-York.
- Barnsley MF, 1988, *Fractals Everywhere*, Academic Press, New-York.
- Broomhead DS, Huke JP and Muldoon MR, 1999, Oversampling Nonlinear Digital Channels, *in preparation*.
- Carretero-González R, Ørstavik S, Huke J, Broomhead DS and Stark J, 1999a, Scaling and Interleaving of Sub-system Lyapunov Exponents for Spatio-temporal Systems, *Chaos*, 9, 466–482.
- Carretero-González R, Ørstavik S, Huke J, Broomhead DS and Stark J, 1999b, Thermodynamic Limit from Small Lattices of Coupled Maps, *Physical Review Letters*. 83, 3633–3636.
- Chen S. and Billings S.A., 1989, Modeling and Analysis of Non-linear Time Series, *Int. J. Control*, 50, 2151–2171.
- Huke J.P., 1993, Embedding Nonlinear Dynamical Systems, *A Guide to Takens Theorem*, Internal Report, DRA Malvern.
- Judd K. and Mees A.I., 1995, On Selecting Models for Nonlinear Time Series, *Physica D*, 82, 426–444.
- Judd K. and Mees A.I., 1998, Embedding as a Modeling Problem, *Physica D*, 120, 273–286.
- Kantz H and Olbrich E, 1997, Scalar Observations From a Class of High Dimensional Chaotic Systems: Limitations of the Time Delay Embedding, *Chaos*, 7, 423–4.
- Kantz H and Schreiber T, 1999, *Nonlinear Time Series Analysis*, 2nd ed., CUP, Cambridge, UK.
- Little S, Ellner S, Pascual M, Neubert M, Kaplan D, Sauer S, Caswell H and Solow A, 1996, Detecting Nonlinear Dynamics in Spatio-Temporal Systems, Examples From Ecological Models, *Physica D*, 96, 321–333.

- Martin R, 1998, *Irregularly Sampled Signals: Theories and Techniques for Analysis*, PhD Thesis, Centre for Nonlinear Dynamics and its Applications, UCL, London, UK.
- Mees AI, 1993, Parsimonious Dynamical Reconstruction, *Int. J. of Bif. and Chaos*, 3, 669–675.
- Muldoon MR, Broomhead DS and Huke JP, 1994, Delay Reconstruction for Multiprobe Signals, *IEE Digest*, 143, 3/1–3/5.
- Norman MF, 1968, Some Convergence Theorems for Stochastic Learning Models with Distance Diminishing Operators, *Journal of Mathematical Psychology*, 5, 61–101.
- Ørstavik S and Stark J, 1998, Reconstruction and Cross-Prediction in Coupled Map Lattices using Spatiotemporal Embeddings, *Phys. Lett. A.*, 247, 145–160.
- Ott E, Sauer T and Yorke JA, 1994, *Coping with Chaos*, Wiley, New-York.
- Rand DA and Wilson HA, 1997, Using Spatio-Temporal Chaos and Intermediate-Scale Determinism to Quantify Spatially Extended Ecosystems, *Proc.R. Soc. Lond. B*, 259, 111–117.
- Sauer T, Yorke JA and Casdagli M, 1991, Embedology, *J. Stat. Phys.*, 65, 579–616.
- Stark J, Broomhead DS, Davies ME and Huke JP, 1997, Takens Embedding Theorems for Forced and Stochastic Systems, *Nonlinear Analysis*, 30, 5303–5314.
- Stark J, 1999, Delay Embeddings of Forced Systems: I Deterministic Forcing, *J. Nonlinear Sci.*, 9, 255–332.
- Stark J, Broomhead DS, Davies ME and Huke JP, 1999, Delay Embeddings of Forced Systems: II Stochastic Forcing, *in preparation*.
- Takens F, 1980, Detecting Strange Attractors in Turbulence, in *Dynamical Systems and Turbulence*, Warwick, 1980, ed. D.A. Rand and L.-S. Young, LNM 898, Springer-Verlag, New-York.

Chapter 5

Some Remarks on the Statistical Modeling of Chaotic Systems

Dominique Guegan

ABSTRACT Many of the issues in modeling nonlinear systems are statistical ones. In this chapter we give a statistician's viewpoint of the problems of studying and understanding data from real-world dynamical systems. We present some new insights into the possibilities of obtaining consistent estimates for the invariants of a dynamical system, and some new results concerning noise-removal.

5.1 Introduction

In this chapter we would like to address some questions that arise when we want to study real data from a statistical modeling point of view.

Assume we observe such real data Y_1, Y_2, \dots, Y_n which come from experimental systems, various kinds of recordings (electrocardiogram, encephalogram, ...), observational phenomena (meteorology, astrophysics, economics, finance, ...) and so on. We are interested in providing some statistical information about this sample set. Of course, this set has a finite size, and we will assume that the size is sufficiently large to avoid technical problems in the following. Our purpose is to have a better understanding of the evolution of this observational trajectory, for instance with a view to making predictions. We are going to assume that there exists a process $(Y_t)_{t \in \mathbb{Z}}$ which explains the data set. We work in discrete time and we suppose that the observations are made with a constant time interval.

One of the first problems we want to solve is to determine in which class this process evolves: *is it a stochastic process or a deterministic process?* This means that we have to decide if the process $(Y_t)_{t \in \mathbb{Z}}$ is generated by one of the two following equations:

$$Y_t = \psi(\underline{X}_t, \varepsilon_t), \quad (5.1)$$

where \underline{X}_t is some deterministic process that we will specify precisely later, or

$$\underline{Y}_t = \varphi(\underline{Y}_{t-1}). \quad (5.2)$$

In these equations, an underlined variable represents a vector. The process $(\varepsilon_t)_{t \in \mathbb{Z}}$ is white noise, i.e., it is a sequence of independent identically distributed random variables (i.i.d.r.v.), and we assume this for (ε_t) throughout the chapter.

Equation 5.1 means that in the evolution process, $(\varepsilon_t)_{t \in \mathbb{Z}}$ is present; thus we say that $(Y_t)_{t \in \mathbb{Z}}$ is a stochastic process. In equation 5.2 no noise appears in the dynamics of the system and we say that $(Y_t)_{t \in \mathbb{Z}}$ is a deterministic process. The functions ψ and φ are assumed to be nonlinear and defined on \mathbb{R}^d . Answering this first question is fundamental to deciding what techniques and what statistics we are going to use to study this observational set of data. We will discuss this question in Section 5.2. In Section 5.3 we will present the main features that we will be able to capture with the stochastic approach 5.2. The presentation of these characteristics could be one possible way to try to answer certain questions arising in Section 5.2. In Section 5.4 we assume that we observe a deterministic system 5.1 and we are interested in ergodic theory and estimation for such models. We will present some new insights into the mixing properties of certain classical well-known chaotic systems, and we will specify how it is possible to obtain consistent estimates for the invariants of a chaotic system. In Section 5.5, we will assume that the chaotic system can be polluted in different ways by different classes of noise; this is generally what we observe in reality. Then we will present new contributions concerning the deconvolution problem and possible new developments on this subject. Section 5.6 will be devoted to further remarks on statistical modeling for chaotic processes in the context of prediction theory. The main results of this chapter are contained in Sections 5.4 and 5.5.

We would like to emphasize that the approach developed in this chapter is neither exhaustive nor complete. It corresponds to a personal point of view on the different problems arising in this type of analysis. This approach is complementary to many others, particularly those developed by physicists, chemists, astrophysicists, controllers and dynamicists. Through this work we hope to make a small contribution to the great developments that we can actually observe in chaos theory.

5.2 How Do We Decide If a Process Is Generated by a Stochastic or a Deterministic System Based on Observed Data?

Let us try to see how this problem can be considered from a statistical approach. One answer lies in the construction of a consistent test. What does this mean? Building a consistent test requires five steps: let a level α be chosen which corresponds to the first order risk of the test; thus we must:

- precisely state the null hypothesis denoted H_0 ,
- similarly state the alternative denoted H_1 (if this is not known, the test is not complete),
- determine the statistic T_n that characterizes the test,
- compute the asymptotic law of T_n under H_0 , and
- compute the corresponding law under H_1 .

The knowledge of the asymptotic distribution of T_n under H_1 is fundamental to obtaining the power of the test. We can only be sure of our decision when using the test (T_n) for an observational sample if this power is quite good. To my knowledge, no such test exists. Some steps are very difficult to solve: the choice of the statistic T_n for which we need to have different behaviors for stochastic or deterministic models and the definition of the alternative. Indeed, if we choose H_0 : the process (Y_t) is stochastic (then we can consider different well-known processes to characterize (Y_t)); against the alternative H_1 : the process (Y_t) is chaotic, then the problem is: How are we going to quantify this assumption?

Until now three approaches to solving this problem could be found in the literature: the predictive approach, the use of Lyapunov exponents and the study of the dimensions.

Concerning the predictive approach, a body of work exists. We refer, for instance, to Farmer and Sidorowich (1987), Casdagli (1989), Smith (1992), Stokbro and Umberger (1992), Michel and Hero (1995), Guégan and Mercier (1997, 1998b) and Mercier (1998) and some earlier references therein. These works propose different approaches to making predictions, most of which use processes such as 5.1, with the main aim of minimizing the noise. No formal test has been proposed with this approach.

Concerning the use of Lyapunov exponents, there is some work which attempts to discriminate between stochastic and chaotic systems: see, for instance, Eckmann and Ruelle (1985), McCaffrey et al. (1992), Abarbanel et al. (1993), Wolff (1992). Nevertheless, the great difficulty is the construction of the statistic T_n . To have discriminating behavior of T_n based on Lyapunov exponents, we need to have a precise definition of Lyapunov exponents for dynamical systems polluted by noise. The recent work of Sauer (1998) presented at the Cambridge workshop seems really promising in this respect; see also Lu and Smith (1998). Thus, from this approach the problem is still open.

Finally, concerning the work using the notion of dimension, there exists the so-called BDS test, see Brock, Dechert and Scheinkman (1996). It uses a statistic T_n based on the correlation dimension. The test assumes that under H_0 the process (Y_t) is white noise and the alternative H_1 is free. Thus we can only detect, in the more favorable cases, whether the process is white noise or not. We are unable to compute the power of this test and thus it

is difficult to make a decision based on observed data. Some later works have studied the empirical power of this test proposing under H_1 some well-known nonlinear stochastic processes. Here the statistic T_n seems to identify the existence of large nonlinearities in the data (see Pesaran and Potter, (1993), Ashley and Patterson, (1998) and Chauveau, Damon and Guégan, (1999)). Nevertheless the test does not permit the detection of chaos; thus the problem is still open. Recent work tries to develop the problem using new insights: see, for instance, Cutler (1998).

Even if this problem is still not solved, some new developments, particularly those presented at the workshop, provide the beginnings of responses to these different questions.

5.3 What Features Are Captured by Stochastic Processes?

In this section we are going to assume that we observe real data generated by a system such as 5.2. We assume that the underlying process $(Y_t)_{t \in \mathbb{Z}}$ is a nonlinear stochastic process and we are interested in describing, in that context, what kinds of characteristics of the data we are able to capture using such stochastic models. This means that we work with parametric models: thus when we reconstruct a dynamic from the data using these specific models, we only consider the properties provided by these models.

The main characteristics that we are able to actually capture with this approach are: some trends, seasonality, different cycles, heteroscedasticity, long memory, intermittenencies, explosions and persistence. Some models permit just one of these properties, others permit several. In any case, the approach can appear very partial in the reconstruction of the dynamics. We briefly describe the most popular models and how we are able to recognize them. We denote in the following the σ -algebra generated by Y_s , $s < t$, by I_{t-1} . First, we will assume when we model stochastic processes for $(Y_t)_{t \in \mathbb{Z}}$ that this process is ergodic; this means that when we retain one model we may verify specific conditions to be sure that the ergodic assumption is verified. We discuss this point at the end of this section. To choose between different classes of models, we mainly study the properties of the empirical moments and the empirical distribution of the observed data.

From 5.2, we concentrate on two classes of stochastic processes that we describe in the following ways:

$$Y_t = h_t \cdot \varepsilon_t \quad (5.3)$$

where h_t is a measurable time-varying function of the information set I_{t-1} , or:

$$Y_t = \psi(Y_{t-1}, \dots, Y_{t-p}, \varepsilon_{t-1}, \dots, \varepsilon_{t-q}) + \varepsilon_t \quad (5.4)$$

Now we will describe in detail some of these functions h and ψ .

From the study of second-order moments and conditional second-order moments, we can detect the existence of trend, cycles and variability in the variance which leads us to choose between different classes of models: the heteroscedastic models and the long memory processes. The heteroscedastic models take into account the great variability of the conditional variance. This interesting feature was first used to model the existence of volatility in financial data. The earliest model is due to Engle (1982), followed by other models: see, for instance, Bollerslev (1986), Taylor (1986), Dingh, Granger and Engle (1993), Diebolt and Guégan (1991, 1993), among others. In the case of a GARCH process (Generalized Heteroscedastic Autoregressive model), the function h_t in 5.3 has the following representation:

$$h_t^2 = a_0 + \sum_{i=1}^p a_i Y_{t-i}^2 + \sum_{j=1}^q b_j h_{t-j}^2 .$$

This representation also permits us to take into account persistence in the data, i.e., the fact that the data stay above a certain value (for instance, the mean), for a long time before decreasing below the threshold. As a matter of fact, this latter behavior is better modeled by long memory processes which are characterized by a slow decay of the autocorrelation function with sometimes seasonality, or by specific behavior of the spectral density exhibiting infinite values in different frequencies. This behavior has been modeled by the FARMA processes introduced by Granger and Joyeux (1980) and Hosking (1981), and by the k -factor GARMA process introduced respectively by Andel (1986), Gray, Zhang and Woodward (1989), Giraitis and Leipus (1995) and Wayne, Woodward, Cheng and Gray (1998). For the k -factor GARMA model, the right hand side of equation 5.3 becomes:

$$Y_t = \prod_{i=1}^k (1 - 2u_i B + B^2)^{-\lambda_i} \varepsilon_t$$

with $0 < u_i \leq 1$, $-\frac{1}{2} \leq \lambda_i < \frac{1}{2}$ and B the backshift operator such that $B\varepsilon_t = \varepsilon_{t-1}$. With the parameters u_i and λ_i , we can model at the same time presence of seasonality and long memory in the data. We can find examples and theory in Ferrara and Guégan (1999, 2000a, b) for these models. If we also observe variability in the variance of observed data, a new model proposed by Guégan (2000) which combines GARCH and k -GARMA models may be appropriate since it permits all the above properties together. Other approaches have been also proposed by Baillie, Bollerslev and Mikkelsen (1996) and Dingh and Granger (1996).

From the study of empirical higher models and empirical distributions, it is also possible to detect specific behavior of the data: existence of extreme values, explosions, fat tails in the empirical distribution. In this case we

can use bilinear models, see Subba Rao and Gabr (1984) or Guégan (1994) for a review. Thus, the function ψ in 5.4 has the following form:

$$\psi(Y, \varepsilon) = \sum_{i=1}^p a_i Y_{t-i} + \sum_{j=1}^q b_j \varepsilon_{t-j} + \sum_{k, \ell=1}^{P, Q} c_{k\ell} Y_{t-k} \varepsilon_{t-\ell}.$$

Here we mainly model the existence of bursts inside the data. We might also prefer to study the graph of the distribution function and use its properties: In that case we can use α -stable models characterized by fat tails; for these models we refer to Samorodnisky and Taqqu (1994). There are also models that allow hidden cycles or different cycles in the data. Recently a new class of models has been proposed which permits the existence of different cycles and passage from one cycle to another. This model is not very far from the behavior of certain chaos in the sense that we are close to the notion of existence of several attractors for the data. Here, this model is defined by a stochastic equation, and Y_t follows the recursive scheme

$$Y_t = \sum_{k=1}^{l_1} \sum_{j=1}^{l_2} [a_0^{(j,k)} + \sum_{i=1}^p a_i^{(j,k)} Y_{t-i} + \varepsilon_t^{(j,k)}] \cdot 1_{(Y_{t-d_1} \in R_j^1)} \cdot 1_{(Y_{t-d_2} \in R_k^2)},$$

for all j, k , $1 \leq j \leq l_1$ and $1 \leq k \leq l_2$; $\{\varepsilon_t^{(j,k)}; n \geq 1\}$ is white noise such that $\varepsilon_t^{(j_1, k_1)}$ is not correlated with $\varepsilon_t^{(j_2, k_2)}$ for all $(j_1, k_1) \neq (j_2, k_2)$. The constants $\{a_i^{(j,k)}; 0 \leq i \leq p; 1 \leq j \leq l_1; 1 \leq k \leq l_2\}$ belong to R . Finally for all $j = 1, \dots, l_1$ and $k = 1, \dots, l_2$, we have $R_j^1 = [r_{j-1}^1, r_j^1]$ and $R_k^2 = [r_{k-1}^2, r_k^2]$ with $r_j^1, r_k^2 \in R$ which verify: $-\infty < r_1^1 < r_2^1 < \dots < r_{l_1}^1 < +\infty$ and $-\infty < r_1^2 < r_2^2 < \dots < r_{l_2}^2 < +\infty$. The constants l_1, l_2, p and $d_1 \leq p, d_2 \leq p$ are positive integer and represent, respectively, *the number of thresholds, the order and the delays* of the model. This model is called COMTAR and is studied and illustrated in two papers by Guégan and Nguyen (1998, 1999). It appears as a generalization of the SETAR model introduced by Lim and Tong (1980).

We recalled briefly some of the most important stochastic models which take into account specific features such as trend, cycles, seasonality, heteroscedasticity, persistence and long memory. The question now is the following: *Do these features also exist in systems generated by chaotic equations?* The investigation of this question can provide an interesting way of discriminating between these two classes of models.

In the beginning of this section we stated that we need to verify, for all the models we use whether, with respect to the observed data, the chosen process is ergodic. We can now give some indication of how to assess this condition for the models we have presented. Generally, we will need to obtain the Markovian representation for these models with a view to using the well-known criteria established by Tweedie, of which a complete description is given in the interesting book of Meyn and Tweedie (1994).

For most of the previous models such criteria have been established; nevertheless this problem is not solved in all contexts and a lot of works remains. For bilinear models some conditions obtained are often difficult to verify, we refer to Guégan (1983) and Pham (1986). For SETAR models, we refer to Petrucelli and Woodward (1984), Chan and Tong (1985, 1991) and Chan (1989), and for the COMTAR models to Guégan and Nguyen (1999). For ARCH models we refer to Diebold and Guégan (1991) and Guégan and Diebolt (1994).

Finally before closing this section we want to emphasize the fact that all the developments using these models are parametric. Some non-parametric techniques have been also used to study some of the previous cited behaviors, particularly semi-parametric methods developed by Robinson (1983, 1991).

5.4 Recent Developments in the Statistical Study of Chaotic Systems

When statisticians work in the stochastic modeling domain, a large part of their research concerns the construction of specific models with a view to taking into account particular features of the data: thus, this approach is mainly parametric as we have seen in the previous section. If we are interested in working with deterministic models, it is not reasonable to try to reconstruct specific chaotic systems from the data. Thus, the methods statisticians are going to use are mainly non-parametric. The choice of methods will depend on the kinds of results we want to obtain. To this end, we might distinguish different purposes: Probabilistic properties, reconstruction of the chaotic map, estimation of several invariants of the chaotic systems such as Lyapunov exponents, dimension of the attractor, embedding dimension, predictions and tests. In this section we present some results based on a specific approach; we do not attempt to be exhaustive in our presentation. This is a particular point of view on the statistical approach for chaotic systems.

Now we assume that we observe real data denoted by X_1, X_2, \dots, X_n , which are generated by the (ideal) following system:

$$\underline{X}_t = \varphi(\underline{X}_{t-1}) \quad (5.5)$$

where φ is a nonlinear function defined on a compact subset D of \mathbb{R}^d to \mathbb{R}^d . We assume that φ is sensitive to initial conditions and that it possesses an attractor in the state space D , see Devaney (1989) and Smale (1967) for precise definitions of these notions. This will be our basic condition for φ to be chaotic. In this section we will discuss some results concerning ergodic theory and estimation theory of chaos. The ergodic property of a process

such as 5.5 is an important consideration. Indeed, we know that the growing interest in deterministic dynamical systems comes from the fact that their trajectories show highly complex behavior, and a common feature of all these models is that they show very different kinds of asymptotic behavior (stable points, limit cycles, sensitivity to initial conditions). Although these concepts are of a topological nature, the behavior of chaotic systems is considered random and it is in that sense that their comparison with stochastic processes is interesting. In a rigorous way this randomness or stochasticity is provided by the ergodic theoretical approach. With respect to the measure associated with the state space of a chaotic system, it is the study of the mixing properties of the dynamics with respect to this measure which specifies the degree of chaos of a dynamical system. Existence or otherwise of an invariant probability measure absolutely continuous with respect to this canonical measure is closely related to this problem. In this sense the study of deterministic dynamical systems is not different to that of stochastic processes. Finally, before describing results concerning specific systems, we will note that few systems, apart from one-dimensional ones will have an invariant measure absolutely continuous with respect to Lebesgue measure. On the other hand, as long as the system has some bounded behavior, it is guaranteed to have an ergodic invariant measure (e.g., Katok and Hasselblatt, (1995)). Typically (in an informal sense) if it is chaotic it will have an infinite number of such measures. The important ergodic property is thus the existence of a so called “natural measure”, and this has been proved to occur in hyperbolic systems, as well as an increasing number of more interesting systems, such as the Henon map, for a set of positive measure of parameters.

5.4.1 Invariance and Mixing Properties in Chaotic Systems

In this part we consider maps defined on $[0, 1]^d \rightarrow [0, 1]^d$. First, if we assume $d = 1$, we can describe two well known classes of families of two branch map models. The type of maps to be studied take the general form 5.5 where

$$\varphi(X_{t-1}) = \begin{cases} \varphi_1(X_{t-1}) & 0 \leq X_{t-1} < c \\ \varphi_2(X_{t-1}) & c \leq X_{t-1} \leq 1. \end{cases}$$

It thus has two branches $\varphi_1(x)$ and $\varphi_2(x)$: each branch will be continuous and monotone, with $0 \leq \varphi_1(x), \varphi_2(x) \leq 1$. We consider two families of maps:

- the generalized tent map defined by

$$X_t = \begin{cases} 1 - (1 - 2X_{t-1})^\nu & 0 \leq X_{t-1} < \frac{1}{2} \\ 1 - (2X_{t-1} - 1)^\nu & \frac{1}{2} \leq X_{t-1} \leq 1 \end{cases} \quad (5.6)$$

where $0 < \nu \leq 2$. The family 5.6 contains the tent map for $\nu = 1$ and the logistic map for $\nu = 2$.

- the generalized binary-shift map defined by

$$X_t = \begin{cases} 1 - (1 - 2X_{t-1})^\nu & 0 \leq X_{t-1} < \frac{1}{2} \\ 1 - (2 - 2X_{t-1})^\nu & \frac{1}{2} \leq X_{t-1} \leq 1 \end{cases} \quad (5.7)$$

where $0 < \nu \leq 2$. The family 5.7 contains the 2-adic map for $\nu = 1$.

No other special cases of maps are known. The invariant distributions of these two families 5.6 and 5.7 are known only in special cases we will discuss and they cannot be obtained analytically, although the methods developed by Kohda and Murao (1990), for instance, should be able to give numerical approximations.

1 - First we consider the generalized binary shift map family.

- When $\nu = 1$, we have the analytical form for the invariant distribution which is the uniform density. The general form of the r -adic map defined by

$$X_t = \varphi(X_{t-1}) = rX_{t-1} \pmod{1}, \quad r > 1, \quad r \in \mathbb{N}$$

also admits this invariant measure. We can obtain this result using the theory of Markov operators (see Lasota and Mackay, (1994)) or from an infinite moving average representation of white noise or from a coarse graining representation. We state here the representation by moving average of white noise that we can use for other systems. We suppose that $r = 2$ (but the method is true for any integer value of r). Fix an initial value X_0 , then $X_n = \varphi^n(X_0)$. The map φ associated with X_0 is a sequence of digits a_0, a_1, a_2, \dots (which is its trajectory) which satisfied the following rule

$$a_i = 0 \text{ if } \varphi^i(x_0) \in \left[0, \frac{1}{2}\right], \quad a_i = 1 \text{ if } \varphi^i(x_0) \in \left[\frac{1}{2}, 1\right].$$

By the way, a_0, a_1, a_2, \dots is the binary expansion of X_0 , thus:

$$X_0 = \sum_{i=0}^{\infty} a_i 2^{-(i+1)}.$$

Endowed with the Lebesgue measure, $[0, 1]$ is a probability space and, through their dependence on X_0 , the a_i become independent identically distributed random variables, with $P[a_i = 0] = P[a_i = 1] = \frac{1}{2}$. Then

$$X_n = \varphi^n(X_0) = \sum_{i=0}^{\infty} a_{n+i} 2^{-(i+1)}$$

and the deterministic sequence $(\varphi^n(X_0))_{n \geq 0}$ is realized as a functional of the I.I.D. process $(a_n)_{n \geq 0}$. Then we can derive that there exists an invariant measure for the process (X_t) and also obtain its density with respect to Lebesgue measure, see Denker and Keller (1986) for other applications.

Another approach consists of mapping the deterministic model in an exact manner into a stochastic process governed by a Master equation. The basic idea is to perform finite coarse graining by choosing a suitable partition in the state space, to project the fine-grained evolution described by the Pearson-Frobenius equation onto the partition and to require that the evolution of the resulting probability vector be generated by a time-independent transition matrix. One obtains in this way a Chapman-Kolmogorov condition imposing constraints on the partition to be chosen and on the type of dynamical systems amenable to such a description. When the Master equation is obtained, on the basis of the Lasota-Yorke theorem (1973), we derive that the transform φ possesses a smooth invariant probability measure. In the case of the r -adic map, the partition has two cells $[0, \frac{1}{2}]$ and $[\frac{1}{2}, 1]$, and the sequence (a_n) permits determination of the transition matrix, and we get density of this distribution with respect to Lebesgue measure.

b. A map called *bimodal*, defined by

$$\begin{aligned} \varphi(x) &= \varphi_1(x) & 0 \leq x \leq \frac{1}{2} \\ \varphi(x) &= 1 - \varphi_1(x) & \frac{1}{2} \leq x \leq 1, \end{aligned}$$

with $\varphi_1(x) = (1 + \delta)x$, $0 \leq x \leq a < \frac{1}{2}$, $\delta \ll 1$ is close in its behavior to the previous map. A coarse graining partition can be obtained, given $0, \varphi(a), \frac{1}{2}, \varphi(1 - a)$, thus we can show the existence of an invariant measure absolutely continuous with respect to Lebesgue measure, see Nicolis, Nicolis and MacKernan (1993).

c. When $d = 2$, we can consider two two-dimensional maps, the Baker map defined by

$$\varphi(x, y) = \begin{cases} (2x, \frac{1}{2}y) & 0 \leq x < \frac{1}{2}, \quad 0 < y < 1 \\ (2x - 1, \frac{1}{2}y + \frac{1}{2}) & \frac{1}{2} \leq x < 1, \quad 0 \leq y \leq 1 \end{cases}$$

and the Anosov - Diffeomorphism defined by

$$\varphi(x, y) = \begin{cases} (\frac{1}{2}x, 2y) & 0 \leq x < 1, \quad 0 \leq y \leq \frac{1}{2} \\ (\frac{1}{2} + \frac{1}{2}x, 2y - 1) & 0 \leq x < 1, \quad \frac{1}{2} \leq y \leq 1. \end{cases}$$

These two transformations are invertible and it can be shown that they admit the Borel measure as an invariant measure, see Lasota and Mackey (1994), and Katok and Hasselblatt (1995).

d. For $d > 1$, we can consider the class of locally expanding maps of Shub (1969), and of Mayer (1984). These classes of maps contain the r -adic ($d = 1$), with $r \in \mathbb{N}$ and $r \in \mathbb{Q}$. They also contain the continued fractional transformation defined by ($d = 1$):

$$\varphi(x) = \frac{1}{x} - \left[\frac{1}{x} \right].$$

The locally expanding maps possess invariant measures, even if the analytical form is not always known and they are characterized by the existence of a Markov partition associated with them, see Krzyzewski and Szlenk (1969) and Mayer (1984).

2. Now we consider the family of generalized tent maps: The analytical form of the invariant distributions is known to be the uniform density for the tent map ($\nu = 1$ in 5.6), and the beta ($\frac{1}{2}, \frac{1}{2}$) density for the logistic map ($\nu = 2$ in 5.6). Proofs can be found in Ulam and von Neumann (1947), Bowen (1978), Hoßbauer and Keller (1982), Lasota and Mackey (1994). If we consider the general representation of the logistic map given by

$$X_t = \alpha X_{t-1}(1 - X_{t-1}) \quad 0 \leq X_{t-1} \leq 1, \quad (5.8)$$

the existence of the beta (1/2, 1/2) density for the logistic map applies to the unique case $\alpha = 4$. For other values of α , the invariant measure is not known. For $3.57 < \alpha < 4$, iterative simulated histograms have been provided by Hall and Wolff (1995) and Ladoucette (1999), but no analytical expression is provided and theoretical results need to be given. Lawrance and Spencer (1996) also give simulated histograms for equation 5.6 for different values of $\nu \neq 1, 2$. An approach for these maps can be also developed by the coarse graining method.

3. Finally, important chaotic systems are those introduced by Smale (1967) which verify the Axiom-A. Ruelle (1977) and Bowen (1978) have been able to show existence of an asymptotic measure for these systems, whose support is often a complicated Cantor-like set. But these measures are not absolutely continuous with respect to ordinary Lebesgue measure. Nevertheless all Axiom A basic sets support an ergodic invariant measure and possess such measures with nice properties, particularly if they have attractors they have the SBR measure, and hence a natural measure. It is interesting to note that for certain non Axiom-A systems invariant measures can be derived, see Meyer and Roepstorff (1983). Standard results show that continuous maps on a compact set always have ergodic invariant measure.

Even if the existence of invariant measure can be ascertained for most of the maps defined from $[0, 1]^d \rightarrow [0, 1]^d$, $d \geq 1$, the properties of mixing for these systems are much more difficult to obtain. Actually most of the results concern the r -adic map. Lasota and Mackey (1994) prove its exactness which implies that this map is strong mixing and geometrically ergodic. See also Denker and Keller (1986) for another approach. No proof has been proposed concerning mixing properties, for instance, for the Anosov and Baker systems.

Recently, Gayraud and Guégan (1999) established the geometric ergodicity of the tent map ($\nu = 2$ in 5.6) using a Master equation. To obtain the latter they follow the work of Nicolis and Nicolis (1988). They first focus

on the two-cell partition of $[0, 1]$ for the tent map:

$$C_1 = \{x : 0 \leq x \leq x_c\}, \quad C_2 = \{x : x_c < x \leq 1\} \quad (5.9)$$

where $x_c = 2/3$ is the position of the unstable fixed point. Using two states denoted by 1 and 2, it is possible to get the transition matrix,

$$W = \begin{pmatrix} 1/2 & 1/2 \\ 1 & 0 \end{pmatrix}$$

for the partition 5.9. The Master equation obtained with this transition matrix describes a Markovian process. It is thus possible to apply to this representation the approach developed by Denker and Keller (1986) for the r -adic map using the Lasota Yorke (1973) theorem. The tent map can be represented as a functional of an I.I.D. process. The stationarity of the function φ is due to the φ invariance of the Lebesgue measure on $[0, 1]$, while the independence of the sequence which defines the transition matrix reflects the mixing property of the tent map φ . This method can be easily extended to another class of generalized tent maps represented by:

$$\begin{aligned} \varphi(x) &= ax \text{ if } 0 \leq x \leq 1/2 \\ \varphi(x) &= a(1-x) \text{ if } 1/2 < x \leq 1, \end{aligned}$$

where a is greater than 2, see MacKernan (1997). Furthermore, the logistic map can be derived from the tent map by a invertible transformation, so the same properties can be obtained in the same way. It is also possible to get the exact Master equation associated with the logistic map using the coarse graining method, see MacKernan (1997). In his thesis MacKernan considered other non-hyperbolic chaotic maps but computing the Markovian matrix is much more difficult. For such maps, for which the dynamics are highly non uniform as in intermittent systems, it is necessary to define partitions containing an infinite but countable number of cells. Another possible extension is a piecewise expression over each cell of a minimal partition as a polynomial of order N . The method was developed by MacKernan and Nicolis (1993) using a generalized Markov coarse graining on a piecewise monomial basis. By this approach, it may be possible to derive mixing conditions for some deterministic systems as soon as the existence of the invariant measure is established, even if we do not know the analytical form of the distribution function.

5.4.2 *Nonparametric Methods for an Estimation Theory in Chaotic Systems*

First we present approaches belonging to the class of kernel methods to construct estimates of the invariant measure of a system and the embedding dimension. This class of methods provides consistent estimates without

using stronger assumptions for the function φ than ergodicity and Lipschitz conditions. For a review of these non-parametric methods, we refer to Silverman (1993) and Guégan (1997).

The first classical method is the kernel method used to obtain a consistent estimate for the invariant measure. If we denote this estimate by f_n , it is defined by:

$$f_n(x) = \frac{1}{nh_n} \sum_{i=1}^n K\left(\frac{x - X_i}{h_n}\right), \quad (5.10)$$

where K is a kernel function on a compact support and (h_n) is a sequence which tends to zero as n tends to infinity. The estimate permits determination of the window around x . With this approach, the rate of convergence for the estimate f_n of the density (when it exists) of the distribution function associated with the system defined in 5.5, is $O(\frac{1}{nh_n^{2h}})$ when $nh_n^{2h} \rightarrow \infty$ as $n \rightarrow \infty$, see Bosq and Guégan (1995). Other approaches have been also considered by Hall and Wolff (1995) and Guégan and Wolff (1999).

Another interesting method is the so-called zero-one explosive method introduced first by Bosq (1989) using the notion of singularity for the density of a vector. We present this method to estimate the chaotic map φ defined in 5.5. The estimate of φ is denoted $\varphi_n(x)$. To obtain rates of convergence for this estimate, we consider the following function

$$g_n(x, y) = \frac{1}{nh_n^2} \sum_{t=1}^{n-1} K\left(\frac{x - X_t}{h_n}\right) K\left(\frac{y - X_{t-1}}{h_n}\right),$$

with the same notation as before and with $(x, y) \in \mathbb{R}^2$, so that $\varphi_n(x)$ is such that:

$$g_n(x, \varphi_n) \geq \sup_{y \in \mathbb{R}^k} g_n(x, y) - \zeta_n$$

where $\zeta_n \rightarrow 0^+$, as $n \rightarrow \infty$. See Bosq and Guégan (1995) for details and theoretical results of this approach.

With this last method it is also possible to obtain a consistent estimate of the embedding dimension with a good rate of convergence under only ergodic assumptions, see Bosq, Guégan and Léorat (1999) for a theoretical approach and Guégan and Léorat (1997) for applications. There exist other works on the estimation of the embedding dimension with more classical approaches, see, for instance, Grassberger and Procaccia (1983) and Cutler (1997).

Another method which can also be used to estimate φ is the regressogram method. Gayraud and Guégan (1999) used this method in a recent paper. The estimate φ_n is defined by

$$\varphi_n(x) = \frac{\sum_{i=1}^n X_i K\left(\frac{x - X_i}{h_n}\right)}{\sum_{i=1}^n K\left(\frac{x - X_i}{h_n}\right)},$$

with same notation as before. Gayraud and Guégan were able to obtain a rate of convergence for φ_n similar to those obtained for mixing processes. The method is close to that developed by Hardle and Tsybakov (1997) for stochastic processes. This estimate is easy to use and gives very nice results for reconstructing the dynamics of a chaotic system such that 5.5.

Besides kernel methods, there exist many non-parametric methods which permit solution of the problems in estimation and prediction theory for chaotic systems; we refer in particular to spline functions, nearest neighbor methods, regression trees, radial basis function methods, neural networks and wavelets. Concerning regression trees we refer in particular to Flandrin and Michel (1993) and Badel, Hero and Michel (1997). For radial basis functions and neural networks a lot of work exists, see, for instance, Girosi and Anzelotti (1993). In his thesis Mercier (1998) gives a rate of convergence for the reconstruction of φ using a radial basis function improving on known results. Concerning nearest neighbor methods, we refer, for instance, to Abarbanel et al. (1993) and Finkenstadt and Kuhbier (1995). One of the problems of interest concerning all of these methods concerns the influence of the initial conditions on the estimates. For new thoughts on this problem we refer to Guégan and Tchernig (2000) and references therein.

5.5 The Deconvolution Problem

When we observe real data $Y_1, Y_2, \dots, Y_n, \dots$, even if we assume that a chaotic behavior characterizes these data, it is reasonable to assume that there exists some noise which pollutes these data. Two situations can arise in that context. The more general situation is the following

$$\begin{cases} Y_t = \psi(X_t, \varepsilon_t) \\ X_t = \varphi(X_{t-1}) \end{cases} \quad (5.11)$$

where $(X_t)_{t \in \mathbb{Z}}$ is assumed to be a chaotic process. The choice of the function ψ is the first problem to solve. Until now, most of the research has assumed that we have the following decomposition for Y_t :

$$Y_t = X_t + \varepsilon_t. \quad (5.12)$$

We know that this decomposition is not unique and we have until now no tools to split (5.11) correctly, so (5.12) is only an arbitrary choice. One of the first problems in this situation is to estimate non-parametrically the density f (when it exists) of the invariant measure μ associated with the process (X_t) using the observations described by (Y_t) . Two particular but important cases can be considered:

- Suppose that we observe a dynamical system such 5.11 but with

errors-in-variables, thus, we obtain the following model:

$$\begin{cases} Y_t &= X_t + \varepsilon_t, & t \geq 0 \\ X_t &= \varphi(X_{t-1}), & t \geq 1. \end{cases} \quad (5.13)$$

Models such 5.13 can appear, for instance, when one wants to simulate systems, in particular aperiodic ones (like chaotic systems, see Guégan and Mercier, 1997). Now, if in 5.13, ε_t is independent of X_t and ε_t has an invariant measure ν , $t \geq 1$, the existence of an invariant measure, say μ for (X_t) implies the same property for (Y_t) and the invariant measure associated with (Y_t) is $\mu * \nu$ where $*$ denotes the convolution product. If the characteristic function of ν does not vanish, then $\mu * \nu$ determines μ : in that case one deals with a deconvolution problem.

- A general model corresponds to propagation of errors. We consider model 5.11 with $\psi(x, y) = \varphi(x) + y$:

$$Y_t = \varphi(Y_{t-1}) + \varepsilon_t. \quad (5.14)$$

Assume Y_0 is observed. For a general t , the relation between Y_t and Y_0 is intricate. However, this representation may be simplified by using successive approximations, so that

$$Y_t = \varphi^{(t)}(Y_0) + \xi_t, \quad t \geq 1 \quad (5.15)$$

where

$$\xi_t = \varepsilon_t + \Gamma_{t-1}\varepsilon_{t-1} + \Gamma_{t-1}\Gamma_{t-2}\varepsilon_{t-2} + \dots + \Gamma_{t-1} \dots \Gamma_1\varepsilon_1,$$

with

$$\Gamma_t = (\varphi' \circ \varphi^{(t)})(Y_0),$$

where φ' , the derivative of φ , is supposed to exist except in a countable set of points and where \circ represents the composition of functions. Notice that if φ is linear, the models given by 5.14 and 5.15 coincide. Models such as 5.15 can be easily found in experimental situations. For example we refer to the Couette-Taylor fluid flow experiment described in Brandstätter and Swinney (1987) and other examples of deterministic noise amplifiers can be found in Deissler and Farmer (1992). In all of these experimental systems, the smallness of the noise is fundamental.

In a recent paper, Blanke, Bosq and Guégan (1999) give the convergence of a kernel estimate for the invariant density (when it exists) of processes such 5.11, for different classes of noise.

In the case of models with errors in variables 5.13, it can be proven that the variance of the estimate f_n is of $O(1/nh_n^{(2\beta+1)d})$ when $nh_n^{(2\beta+1)d} \rightarrow \infty$

as $n \rightarrow \infty$. This result is established for ordinary smooth noises such as Laplacian and Gamma densities. Another class of noise called the super-smooth noises, which includes Gaussian noise or a Cauchy law, have slower rates of convergence for f_n . The nonparametric kernel-type density estimator used here is defined by:

$$f_n(x) = 1/nh_n^d \sum_{i=0}^{n-1} W_h\left(\frac{x - Y_i}{h_n}\right), \quad x \in R^d$$

with

$$W_h(x) = \left(\frac{1}{2\pi}\right)^d \int_{R^d} e^{-i\langle t, x \rangle} \frac{\phi_K(t)}{\phi_\varepsilon(t/h_n)} dt$$

where ϕ_K is the Fourier transform of a kernel K and ϕ_ε is the characteristic function of the noise ε . This is a classical deconvolution kernel developed, for instance, for mixing processes by Masry (1991).

Models with small noise may also be represented in the following way:

$$Y_i = X_i + \sigma\varepsilon_i,$$

with (X_t) satisfying relation 5.5. Hence the variance of the estimate f_n converges to zero in $O(n^{-4/5})$ as n tends to $+\infty$, when we choose $h_n = n^{-1/5}$. Since we no longer follow a deconvolution approach for this last model, the estimate used for f_n is the classical estimate defined in 5.10.

Based on this work, Blanke, Bosq and Guégan (1999) show that it is possible to adapt theoretical results obtained for chaotic systems to systems with noise. Nevertheless, it appears that the presence of Gaussian noise gives poor result, and, for instance, unbounded noise such as beta(1/2, 1/2) noise give the worst. This seems to confirm some of the results presented elsewhere in this workshop, see Sauer (1998) and Stark et al. (1998), for instance.

The reconstruction of the chaotic map φ of 5.5 in the state space D , with only the observational data Y_1, Y_2, \dots, Y_n , belongs also to the deconvolution problem. Some proposed solutions exist already in the literature, see, for instance, Broomhead and King (1986), Farmer and Sidorowitch (1988), Lisi, Nicolis and Sandri (1995). Nevertheless, until now no theoretical work strengthens the consistency of the reconstruction. To solve this problem, we need to establish what kind of noise is possible; we can use the previous developments to try to answer this question.

Some work also exists on estimation of other characteristics for the process $(X_t)_{t \in \mathbb{Z}}$ obtained from $(Y_t)_{t \in \mathbb{Z}}$. They concern the construction of Lyapunov exponents, (see Sauer, (1998)), the estimation of embedding dimension, (see Cutler (1998)) and the extension to Takens' theorem, (see Stark et al. (1998)). All show that the choice of the noise is fundamental. Given this, the way in which the noise pollutes the original system is also very

important. From the published results, we can see that the rates of convergence are completely different observing whether we have errors in variables or propagation errors. Thus, further development is required in order to understand all the effects of noise on a chaotic system.

5.6 Further Developments and Remarks

In order to describe a theory of identification of chaotic systems many problems need to be solved. Different approaches to building such a theory might be considered. We list a selection of these and give some idea of possible developments from a statistical point of view:

- Whenever we discuss modeling of chaos, we need to state the characteristics we want to take into account. In the literature there are different definitions for a chaotic system and before attempting to unify these (which may not be realistic from a modeling point of view), it is important to specify the domain in which we will work.
- Until now, detecting chaos from observational data has been very difficult, unless there has been the opportunity of directly finding an attractor from the data. Most of the time, this is the preference of physicists, but econometricians and people working in the finance area generally cannot achieve this because their data are extremely polluted. This can also be the case in meteorology and astrophysics. Thus, we generally work under the assumption of chaotic behavior within the data or very complex behavior that is not close to that generally known and described by stochastic processes. This means that the construction of the statistic T_n or different statistics permitting discrimination in a consistent way of stochastic behavior and deterministic behavior must be in accordance with the a priori assumption. In my opinion, it seems that the work of Sauer (1998) could be an important contribution in this way.
- The theoretical properties of characteristics of a chaotic system are well known. We need to estimate them in a consistent way with knowledge of the asymptotic laws of their estimates. The non-parametric approaches based on kernel methods developed, for instance, by the works of Bosq and Guégan (1995) and Bosq, Guégan and Léorat (1997), are also promising and give interesting results. It seems that these works need to be extended to other statistics which characterize the deterministic dynamical systems.
- Making predictions with chaotic dynamical system seems a joke because generally we think that such a system is unpredictable—in a certain sense, like random walks. Nevertheless, if we can reconstruct

the chaotic map in a consistent way, then short term prediction can be obtained, see, for instance, Guégan and Lisi (1997) and Guégan and Mercier (1998 a). Indeed, predictions with chaotic systems are better than those obtained with complex stochastic systems because in the latter case we mainly predict noise whereas with the former approach it is the complex structure of the systems which provides the predictions. We can also consider medium and long-term horizons for predictions. In these cases we need to have a better understanding of Lyapunov exponents and of the distribution of the predictions.

- Concerning the deconvolution problem, very little work exists and the problem needs further development. We again postulate that kernel methods and wavelets may provide useful results.

We see by this short review that the problem of statistical modeling of chaotic systems still requires much work, but there has been some progress toward its solution in recent years.

References

- Abarbanel H.D.I., Brown R., Sidorowitch, L.S. Tsimring (1993), The analysis of observed chaotic data in physical systems, *Reviews of modern physics*, **65**, 1331–1392.
- Andel J. (1986), Long Memory Time Series, *Kybernetika*, **22**, 105–123.
- Ashley, R.A., D.M. Patterson (1998), A direct comparison of the BDS, Hinich, and other tests for the presence of nonlinear dependence in time series, *Preprint*.
- Badel A.E., O. Michel, A. Hero (1997), Arbres de régression: modélisation non paramétrique et analyse de séries temporelles, *Traitement du Signal*, to appear.
- Baillie R.T., Bollerslev T. and Mikkelsen H.O. (1996), Fractionally Integrated Generalized Autoregressive Conditional Heteroscedasticity, *Journal of Econometrics*, **31**, 307–327.
- Blanke D., Bosq D. and Guégan D. (1999), Modelization and non-parametric estimation for a dynamical system with noise, *Preprint CREST*, Paris.
- Bollerslev T. (1986) Generalized Autoregressive Conditional Heteroscedasticity, *Journal of Econometrics*, **31**, 307–259.
- Bosq D. (1989) Non-parametric estimation of a nonlinear filter using a density estimator with a zero-one explosive behaviour in \mathbb{R}^d , *Statistics and Decisions*, **7**.
- Bosq D., D. Guégan (1995) Non-parametric estimation of the chaotic function and the invariant measure of a dynamical system, *Stat. Proba. Letters*, **25**, 201–212.
- Bosq D., D. Guégan, G. Léorat (1999), Statistical estimation of the embedding dimension of a dynamical system, To appear in *International Journal of Bifurcations and Chaos*.
- Bowen R. (1978), *Israel J. Math.*, **28**, 298–314.
- Brandstater A., J Swinney (1987), *Phys. Rev. A*, **35**, 2207.

- Brock W.A., W.D. Dechert, J.A. Scheinkman, B. LeBaron (1996), A test for independence based on the correlation dimension, *Econometrics Review*, **15**, 197–235.
- Broomhead D.S., G.P. King (1986), Extracting qualitative dynamics from experimental data, *Physica D*, **20**, 217–236.
- Casdagli M. (1989), Nonlinear prediction of chaotic time series, *Physica D*, **35**, 335–356, 1989.
- Chan K.S. (1989), A note on the ergodicity of a Markov chain, *Adv. in Appl. Prob.*, **21**, 702–704.
- Chan K.S., H. Tong (1985), On the use of deterministic Lyapunov function for the ergodicity of stochastic differences equations, *Adv. in Appl. Prob.*, **17**, 666–678.
- Chauveau T., J. Damon, D. Guégan (1999) Testing for nonlinearity in intra-day financial series: the case of two French stocks, *Preprint Caisse des dépôts*, Paris.
- Cutler C. (1997), A general approach to predictive and fractal scaling dimensions in discrete index time series, *Fields Institute Communications*, **11**, 29–48.
- Cutler C. (1998), Determinism, (bad) Embeddings and scaling structures in time series, *Presentation at the Workshop in Cambridge*, U.K.
- Deissler M., D. Farmer (1992), Deterministic noise amplifiers, *Physica D*, **55**, 155–165.
- Denker M. and Keller G. (1986), Rigorous Statistical Procedures for Data from Dynamical Systems, *Journal of Statistical Physics*, **44**, 67–93.
- Devaney R.L. (1989), An introduction to chaotic dynamical systems, *Addison Wesley Pul. Comp.*, N.Y.
- Diebolt J. and Guégan D. (1991), Le modèle β -ARCH, *CRAS, Série I*, **312**, 625–630.
- Diebolt J. and Guégan D. (1993), Tail Behavior of the Stationary Density of General Nonlinear Autoregressive Process of Order One, *J. Appl. Prob.*, **30**, 315–329.
- Dingh Z. and Granger C.W.J. (1996), Modelling Volatility Persistence of Speculative Returns: A New Approach, *Journal of Econometrics*, **73**, 185–215.
- Dingh Z., Granger C.W.J. and Engle R.F. (1993), A Long Memory Property of Stock Market Return and a New Model, *Journal of Empirical Finance*, **1**, 83–106.
- Eckmann J.P. , Ruelle D. (1985), Ergodic theory of chaos and strange attractors, *Reviews of Modern Physics*, **57**, 617–656.
- Engle R.F. (1982), Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation, *Econometrica*, **50**, 987–1007.
- Farmer J.D., V. Sidorovich (1987), Predicting chaotic time series, *Physical Review Letters*, **59**, 845–848.
- Farmer J.D., V. Sidorovich (1988), Exploiting chaos to predict the future and reduce noise, in *Evolution, Learning and cognition*, eds. Lee, World Scientific press.
- Ferrara L. and Guégan D. (1999), Forecasting with k -factor Gegenbauer Processes, *Preprint CREST*, Paris.
- Ferrara L. and Guégan D. (2000a), Gegenbauer Processes: Estimation and Applications. To appear in *Journal of econometrics*.
- Ferrara L. and Guégan D. (2000b), An approach to financial time series by generalized long memory processes. To appear in *Advances in Asset Management*, ed. C. Dunis, Kluwer Academic Publishers.
- Finkenstadt B., P. Kuhbier (1995) Forecasting nonlinear economic time series: a simple test to accompany the nearest neighbor method, *Empirical economics*, **20**, 243–263.

- Flandrin P., Michel O., P. Ruiz (1993), Chaos et analyse non linéaire du signal, *Rapport Interne Laboratoire de Physique*, E.N.S. Lyon. France.
- Gayraud G. and Guégan D. (1999), Estimation of the chaotic map by regressogram, *Preprint CREST*, Paris.
- Giraitis L. and Leipus R. (1995), A Generalized Fractionally Differencing Approach in Long Memory Modelling, *Lithuanian Math. Journ.*, **35**, 65–81.
- Girosi F., G., Anzelotti (1993), Rates of convergence for radial basis functions and neural networks, in *Artificial Neural Network for speech and Vision*, R.J. Mammone ed., Chapman and Hall, 97–114.
- Granger C.W.J. and Joyeux R. (1980), An Introduction to Long Memory Time Series Models and Fractional Differencing, *J.T.S.A.*, **1**, 15–29.
- Grassberger P. and I. Procaccia (1983), Measuring the strangeness of strange attractors, *Physica D*, **9**, 189–208.
- Gray H.L., Zhang N. and Woodward A. (1989), On Generalized Fractional Processes, *J.T.S.A.*, **10**, 233–257.
- Guégan D. (1983), Une condition d'ergodicité pour des modèles bilinéaires à temps discret,, *CRAS Série I*, 297–300.
- Guégan D. (1994), Séries chronologiques non linéaires à temps discret, *Economica*, Paris.
- Guégan D. (1997) Non-parametric methods for time series and dynamical systems, in *Statistical Challenges in Modern Astronomy II*, G.J. Babu and E.D. Feigelson eds., Springer Verlag.
- Guégan D. (1999), Note on Long Memory Processes with Cyclical Behavior and Heteroscedasticity, *Prepublication n°99 – 08*, 1–21, University of Reims, France.
- Guégan D. (2000), A new model: the k-factor GIGARCH process. To appear in *Journal of signal processing*.
- Guégan D., J. Diebolt (1994), Probabilistic properties of the β -ARCH model, *Statistica Sinica*, **2**, 157–174.
- Guégan D. and Léorat G. (1997), Consistent estimation to determine embedding dimension in financial data,, *The European Journal of Finance*, 231–242.
- Guégan D. and Lisi F. (1997), Predictive dimension: an alternative definition of the embedding dimension,, *Preprint CREST 9749*, Paris.
- Guégan D. and Mercier L. (1997), Prediction in chaotic time series: Methods and comparisons using simulations, in *Signal Analysis and Prediction I*, A. Prochazka, J. Uhler, P. Sovka eds., EURASIP, ICP Press, Prague 1997 , 215–218.
- Guégan D. and Mercier L. (1998a), Stochastic or chaotic dynamics in high frequency financial data, in *Nonlinear modelling high frequency financial time series*, C. Dunis and B. Zhou eds., John Wiley and Sons, 87–107.
- Guégan D. and Mercier L. (1998b), Forecasting and non-parametric techniques for mixing chaotic time series, in *Signal analysis and prediction, Chapter 25*, Birkhauser, Boston, 363–372.
- Guégan D. and Nguyen J.M. (1998), The Multiplicative Threshold Model: An Alternative to Detect Breaks and Hidden Cycles on Real Data, *Preprint CREST n°9839*, 1–13, Paris.
- Guégan D. and Nguyen J.M. (1999), COMTAR Models: Theory and Applications, *Preprint CREST*, Paris.

- Guégan D. and Tchernig R. (2000), Prediction of chaotic time series in presence of measurement error: the importance of initial conditions, to appear in *Statistics and Computing*.
- Guégan D., R. Wolff (1999), On non parametric recursive estimators of invariant densities of chaotic maps, *Preprint, University of Brisbane, Australia*.
- Hall P. and Wolff R.C.L. (1995), Properties of Invariant Distributions and Lyapunov Exponents for Chaotic Logistic Maps, *J.R.S.S., B*, **57**.
- Hofbauer F. and Keller G. (1982), *Math. Zeitschrift*, **180**, 119–140.
- Hardle W., A.B. Tsybakov (1997) Local polynomial estimators of the volatility function in non-parametric autoregression, *J. of Econometrics*, **81**, 223–242.
- Hosking J.R.M. (1981), Fractional Differencing, *Biometrika*, **68**, 165–176.
- Katok A., Hasselblatt B. (1995) Introduction to the Modern Theory of Dynamical Systems, *Encyclopedia of Mathematics and its applications*, **54**, Cambridge University Press.
- Kohda T. and Murao K. (1990), Approach of to Time Series Analysis of One Dimensional Chaos Based on Frobenius-Perron Operator, *Trans. Inst. Electron. Inf. and Comm. Eng.*, E73, **6**, 793–800, Japan.
- Krzyzewski K., Szlenk W. (1969), On invariant measures for expanding differential mappings, *Stud. math.*, **33**, 83–92.
- Ladoucette S. (1999), Etude d'un système dynamique chaotique: approche probabiliste, *Mémoire de DEA*, Université de Reims.
- Lasota, J. and Mac Key, M.C. (1994), Chaos, fractals and noise: stochastic aspects of dynamics, 2nd ed., *Springer Verlag*, N.Y.
- Lasota A., Yorke J. A. (1973), *Trans. Am. Math. Soc.*, **186**, 481.
- Lawrance A.J. and Spencer N.M. (1996), Statistical Aspects of Curved Chaotic Map Models and their Stochastic Reversibility, *Preprint*, University of Birmingham, United Kingdom.
- Lim K.S., H. Tong, (1980), Threshold Autoregression, limit cycles and cyclical data, *J.R.S.S., B* **42**, 245–292.
- Lisi F., Nicolis O., Sandri M. (1995), Combining singular - spectrum analysis and neural networks for time series forecasting, *Neural processing letters*, **2**, 6–10.
- Lu Z.Q., Smith R.L. (1992), Estimating local Lyapunov exponents. To appear in *Fields Institute Communication*.
- Mac Caffrey D.F., S. Ellner, A.R. Gallant, D.W. Nychka (1992), Estimating Lyapunov exponents of a chaotic system with non parametric regression, *J. T.S.A.*, **87**, 682–695.
- MacKernan D. N. (1997), Generalized Markov coarse graining and the observables of chaos, *Thèse de doctorat*, U.L.B., Brussels.
- MacKernan D., Nicolis G. (1993), Generalized Markov Coarse graining and spectral decomposition of chaotic piecewise linear maps, *Preprint ULB*, Brussels.
- Masry E. (1991) Multivariate probability density deconvolution for stationary processes, *IEEE Trans. Inform. Theory*, **37**, 1105–1115.
- Mayer D.H. (1984), Approach to Equilibrium for Locally Expanding Maps in \mathbb{R}^k , *Comm. Math. Phys.*, **95**, 1–15.
- Meyn S.P., R. Tweedie (1994), Markov chains and stochastic stability, *Springer Verlag*, N.Y.

- Meyer D.H. and Roepstorff G. (1983), Strange Attractors and Asymptotic Measures of Discrete-Time Dissipative Systems, *Journal of Statistical Physics*, **31**, 309–326.
- Mercier L. (1998) Séries temporelles chaotiques appliquées à la finance: problèmes statistiques et algorithmiques., *Thèse de Doctorat*, Université Paris IX.
- Michel O., A. Hero (1995), The structured nonlinear signal modelling and prediction, *Preprint*.
- Nicolis G., Nicolis C. (1988), Master equation approach to deterministic chaos, *Physical Review A*, **38**, 427–433.
- Nicolis G., Nicolis C. and MacKernan D. (1993), Stochastic Resonance in Chaotic Dynamics, *Journal of Statistical Physics*, **70**, 127–139.
- Pesaran M.H. , Potter S.M. (1993), Nonlinear dynamics, chaos and econometrics, J. Wiley, New York.
- Petrucci J.D., S.W. Woolford (1984), A threshold AR(1) model, *J. Appl. Prob.*, **21**, 270–286.
- Pham D.T. (1986) The mixing property of bilinear and generalized random coefficient autoregressive models, *Stoch. Proc. and their Appl.*, **23**, 291–300.
- Robinson P.M. (1983), Non-parametric estimation for time series models, *J.T.S.A.* **4**, 185–208.
- Robinson P.M. (1991), Testing for Strong Serial Correlation and Dynamics Conditional Heteroscedasticity in Multiple Regression, *Journal of Econometrics*, **47**, 67–84.
- Ruelle D. (1977), Applications conservant une mesure absolue continue par rapport à dx sur $[0, 1]$, *Commun. Math. Phys.*, **55**, 477–493.
- Samorodnisky V. and Taqqu M. (1994), α -Stable Processes, *Chapman and Hall*, New York.
- Smale S. (1967), *Bull. Am. Math. Soc.*, **73**, 817.
- Sauer T. (1998), Embedology, Eckmann-Ruelle matrices and Lyapunov exponents, *Presentation to the Workshop in Cambridge*, U.K..
- Silverman B.W. (1993), Density estimation, *Chapman and Hall*, N.Y.
- Shub M. (1969), *Am. J. Math.*, **91**, 175–199.
- Smith L.A. (1992), Identification and prediction of low dimensional dynamics, *Physica D*, **58**, 50–76.
- Stark J., D.S. Broomhead, M.E. Davies, J. Huke (1998)? Takens embedding theorems for forced and stochastic systems, *Presentation to the Workshop in Cambridge*, U.K..
- Stockbro K., D.K. Umberger (1992), Forecasting with weighted maps, in *Nonlinear modelling*, eds. Casdagli and Eubank, Addison-Wesley.
- Subba Rao T., M.M. Gabr (1984), Bilinear time series, *Lecture notes in Statistics*, **10**, Springer.
- Taylor S. (1986), Modelling Financial Time Series, *J. Wiley*, Chichester.
- Ulam S.M. , J. Von Neuman (1947), On the combination of stochastic and deterministic properties, *Bull. Am. Math. Soc.*, **53**, 1120.
- Wayne A., Woodward A., Cheng A.C. and Gray H.L. (1998), A k -Factor GARMA Long Memory Model, *J.T.S.A.*, **19**, 485–504.
- R. Wolff (1992), Local Lyapunov exponents : looking closely at chaos, *J.R.S.S B*, **54**, 353–372.

Part II

Fundamentals

Chapter 6

The Identification and Estimation of Nonlinear Stochastic Systems

Peter Young

ABSTRACT This chapter describes what might be called the system theorist's approach to understanding dynamics of nonlinear stochastic systems. The method uses so-called state-dependent parameters, and is able to handle non-stationarity, as long as the state-dependent parameters vary slowly compared to the significant dynamics. One of the main points made here is that most realistic systems have time-varying inputs which can be measured; models must take this into account, and indeed modeling often becomes easier rather than harder when this is done. We describe the methods used, based on recursive fixed-interval smoothing, and present applications to some realistic problems.

6.1 Introduction

Previous publications (e.g., Young, 1978, 1983, 1993a,b, 1996, 1998a,b, 1999a,b; Young and Runkle, 1989; Young and Minchin, 1991; Young *et al.*, 1991; Young and Lees, 1993; Young and Beven, 1994; Young and Pedregal, 1997, 1998, 1999) have discussed an approach to non-stationary and nonlinear signal processing based on the identification and estimation of stochastic models with time variable (TVP) or state dependent (SDP) parameters. Here the term 'non-stationarity' is assumed to mean that the statistical properties of the signal, as defined by the parameters in an associated stochastic model, are changing over time at a rate which is "slow" in relation to the rates of change of the stochastic state variables in the system under study. Although such non-stationary systems exhibit nonlinear behavior, this can often be approximated well by TVP (or piece-wise linear) models, the parameters of which can be estimated using recursive methods of estimation in which the parameters are assumed to evolve in a simple stochastic manner (e.g., Young, 1984, 1999a). On the other hand, if the changes in the parameters are functions of the state or input variables (i.e., they actually constitute stochastic state variables), then the system is truly nonlinear and likely to exhibit severe nonlinear behavior. Normally, this cannot be approximated in a simple TVP manner; in which case, re-

course must be made to the alternative, and more powerful SDP modeling methods that are the main topic of this chapter.

As far as the author is aware, the idea of SDP modeling originated in his 1978 paper on the modeling of badly defined dynamic systems (Young, 1978) and was then taken up by Priestley, in a series of papers and a book on the subject (Priestley, 1988). These earlier publications do not, however, exploit the power of recursive *fixed interval smoothing* (FIS), which provides the main engine for the developments described in this chapter. In particular, the presently proposed approach involves two statistical stages.

- First, the non-parametric *identification* of the state dependency using recursive methods of time-variable parameter estimation based on FIS which allow for rapid (state-dependent) parametric change. As we shall see, the standard methods of TVP estimation developed previously for non-stationary time series analysis need to be modified considerably in this SDP setting to allow for the much more rapid temporal changes that arise from the state dependency.
- Second, the parameterization of the identified non-parametric relationships, followed by the *statistically more efficient estimation* of the (now normally constant) parameters that characterize these SDP nonlinearities.

The first identification stage in this process exploits FIS algorithms, combined with special data re-ordering and back-fitting procedures, to obtain estimates of any state-dependent parameter variations. These state dependencies are estimated in the form of non-parametric relationships (graphs) between the estimated rapid parameter variation and the associated state or input variable(s). Parameterization of these non-parametric relationships can be accomplished in various ways, from simple curve fitting based on weighted least squares methods (Young, 1993a; Young and Beven, 1994) to the use of neural networks or radial basis functions.

Having identified a structural form for the nonlinear model of the system based on the parameterized nonlinear relationships, the (normally constant) parameters of this model are then re-estimated using some form of statistical parameter estimation. This can take various forms: for example, at one extreme, it could be based on simple, nonlinear least squares estimation; while at the other it could exploit a *maximum likelihood* (ML) approach, based on Gaussian assumptions for the stochastic disturbances and the application of *prediction error decomposition* (see later and also Scheppe, 1965). The resulting model should then provide a parametrically efficient representation of the stochastic, nonlinear system that has considerable potential for use in subsequent signal processing, time series analysis and automatic control system design. For example, the methodology described in this chapter exploits recursive estimation in an *off-line* manner, but this sequential processing of the data facilitates the development of related *on-line adaptive* methods of signal processing, forecasting and control.

Although primarily concerned with nonlinear state dependent parameter models, as outlined earlier, this chapter also provides a sequel to a previous paper (Young, 1999a) that discusses the simpler class of “linear” TVP regression relationships. As a prelude to our discussion of SDP estimation, therefore, it is instructive to consider the simpler *Time Varying Parameter Regression* (TVPR) model¹, since the major algorithmic aspects of the estimation are very similar to those used in the SDP situation.

6.2 Time Varying Parameter Regression (TVPR) Models

The TVPR model can be written in the following vector equation form:

$$y_t = \mathbf{u}_t^T \mathbf{p}_t + e \quad e_t = N(0, \sigma^2) \quad t = 1, 2, \dots, N \quad (6.1a)$$

where

$$\begin{aligned} \mathbf{u}_t^T &= [u_{1,t} \ u_{2,t} \ \dots \ u_{n,t}] \\ \mathbf{p}_t &= [a_{1,t} \ a_{2,t} \ \dots \ a_{n,t}]^T \end{aligned} \quad (6.1b)$$

Here, the subscript t denotes the value of the associated variable at the t^{th} instant of time; $u_{i,t}$, $i = 1, 2, \dots, n$ are the regression variables (regressors or input variables), assumed here to be functions of t ; and $a_{i,t}$, $i = 1, 2, \dots, n$ are the associated parameters, *assumed to be slowly variable functions of time*.

In order to estimate the assumed time variable model parameters in \mathbf{p}_t , it is necessary to make some assumptions about the nature of their temporal variability. Reflecting the statistical setting of the analysis and referring to previous research on this topic, it seems desirable if this is characterized in some stochastic manner. Normally, when little is known about the nature of the time variability, this model needs to be both simple but flexible. One of the simplest and most generally useful class of stochastic, state space models involves the assumption the i^{th} parameter, $p_{i,t}$, $i = 1, 2, \dots, n$, in \mathbf{p}_t is defined by a two-dimensional stochastic state vector $\mathbf{x}_{i,t} = [l_{i,t} \ d_{i,t}]^T$, where $l_{i,t}$ and $d_{i,t}$ are, respectively, the changing level and slope of the associated TVP. This selection of a two-dimensional state representation of the TVPs is based on practical experience over a number of years. Initial research by the author and others in the 1960s (e.g., Young, 1969, 1970) tended to use a simple scalar random walk (RW) model for the parameter variations. However, later work in the 1980s (see above references) showed the value of modeling not only the level changes in the TVPs but also their rates of change.

¹For historical reasons, this model has also been called the *dynamic linear regression* (DLR) model (see Young, 1991a and the references therein), which can be misleading since the model is not necessarily dynamic in a systems sense.

The stochastic evolution of each $\mathbf{x}_{i,t}$ (and, therefore, of each of the n parameters in \mathbf{p}_t) is assumed to be described by a generalized random walk (GRW) process defined in the following state space (SS) terms:

$$\mathbf{x}_{i,t} = \mathbf{F}_i \mathbf{x}_{i,t-1} + \mathbf{G}_i \eta_{i,t} \quad i = 1, 2, \dots, m \quad (6.1c)$$

where

$$\mathbf{F}_i = \begin{bmatrix} \alpha_i & \beta_i \\ 0 & \gamma_i \end{bmatrix}, \quad \mathbf{G}_i = \begin{bmatrix} \delta_i & 0 \\ 0 & \varepsilon_i \end{bmatrix}$$

and $\eta_{i,t} = [\eta_{1i,t} \ \eta_{2i,t}]^T$ is a 2×1 , zero mean, white noise vector that allows for stochastic variability in the parameters and is assumed to be characterized by a (normally diagonal) covariance matrix \mathbf{Q}_{η_i} . Omitting the i subscript for convenience, this general model comprises as special cases: the *integrated random walk* (IRW: $\alpha = \beta = \gamma = \varepsilon = 1; \delta = 0$); the scalar *random walk* (scalar but equivalent to (6.1c) with $\beta = \gamma = \varepsilon = 0; \alpha = \delta = 1$, and the first order auto regressive, AR(1) model with $\beta = \gamma = \varepsilon = 0; 0 < \alpha < 1; \delta = 1$: i.e., both just relating to the first equation in (6.1c), see below); the intermediate case of *smoothed random walk* (SRW: $0 < \alpha < 1; \beta = \gamma = \varepsilon = 1$; and $\delta = 0$); and, finally, both the *local linear trend* (LLT: $\alpha = \beta = \gamma = \varepsilon = 1; \delta = 1$)² and *damped trend* (DT: $\alpha = \beta = \delta = \varepsilon = 1; 0 < \gamma < 1$): see Harvey, 1984, 1989. The various, normally constant, parameters in this GRW model ($\alpha, \beta, \gamma, \delta, \varepsilon$, and the elements of \mathbf{Q}_{η_i}) are often referred to as *hyper-parameters*. This is to differentiate them from the TVPs that are the main object of the estimation analysis. However, the hyper-parameters are also assumed to be unknown a priori and need to be estimated from the data, as we shall see in the subsequent discussion.

The full GRW model (6.1c) was introduced in Jakeman and Young (1979, 1984): further discussion and practical examples appear in Young (1988), Young et al. (1989), Young and Ng (1989), Ng and Young (1990). Note that, in the case of the RW model, i.e.,

$$l_{i,t} = l_{i,t-1} + \eta_{1i,t} \ ; \ l_{i,t} = p_{i,t} \quad (6.2)$$

each parameter can be assumed to be time-invariant if the variance of the white noise input $\eta_{1i,t}$ is set to zero. Then the stochastic TVP setting reverts to the more normal, constant parameter regression situation. In other words, the recursive estimation algorithms described later for the general stochastic TVP case will provide *constant parameter* estimates, identical to the normal *en-bloc* regression, if RW models with zero variance white noise inputs are specified. Of course, there is some added value to the recursive solution even in this situation, since the user is provided

²Interestingly, the LLT model can be considered simply as the combination of the simpler RW and IRW models.

with the recursive estimates over the whole interval $t = 1, 2, \dots, N$. These can provide additional useful information on the model: for example, they show how the estimates are converging and can be used (see Brown *et al*, 1975; Young, 1984) to detect both the presence of potential parametric change and possible over-parameterization (i.e., the model contains too many parameters to provide unambiguous estimation results).

Clearly other, more general and higher order stochastic processes could be used to model the stochastic TVPs, provided such models can be identified satisfactorily from the data. For example the higher order IRWs (double and triple integrated random walk (DIRW, TIRW), etc.), the integrated or double integrated auto regressive (IAR, DIAR: see Young, 1994) model, and even more general processes (e.g., Pedregal and Young, 1996, 1998). However, the more complex models introduce additional hyper-parameters that would have to be well identified from the data and optimized, thus introducing potential practical difficulties.

The idea of assuming that the model parameters evolve over time as non-stationary stochastic variables may seem complex at first sight but it is, in fact, just a statistical device to allow for the estimation of parametric change. After all, the assumption of the RW model is simply a means of introducing into the estimation problem the freedom for the associated parameter to vary at each sample in time by a small random amount defined by the variance of the white noise input $\eta_{1i,t}$. And the more complex GRW models in (6.1c) are just a way of refining and adding to this freedom. In fact, it can be shown (Young and Pedregal, 1998) that the GRW assumptions on the parameter variations have an *implicit* but physically interpretable effect: They make the recursive parameter estimates, at any sample time t , depend only on the local data in the vicinity of this sample, with the selected GRW model defining the local weighting effect. In the case of the RW model, for instance, this weighting effect or “kernel” has a Gaussian-like shape that applies maximum weight to the current data with declining weight on each side. And the “bandwidth” of the kernel is defined by the ratio of the variance of the white noise input η_{1it} to the residual variance σ^2 (the noise variance ratio (NVR)). This can be related to the more *explicit* use of localized data weighting in methods such as locally weighted kernel regression (e.g., Holst *et al*, 1996; Young and Pedregal, 1996) and ‘wavelet’ methods (e.g., Daubechies, 1988) that are currently receiving so much attention in the statistical and signal processing literature.

Having introduced the GRW models for the parameter variations, an overall SS model can be constructed straightforwardly by the aggregation of the subsystem matrices defined in (6.1c), with the “observation” equation defined by the model equation (6.1a): i.e.,

$$\begin{array}{ll} \text{Observation equation} & : \quad y_t = \mathbf{H}_t \mathbf{x}_t + e_t & (i) \\ \text{State equations} & : \quad \mathbf{x}_t = \mathbf{F} \mathbf{x}_{t-1} + \mathbf{G} \eta_t. & (ii) \end{array} \quad (6.3a)$$

If $p = 2n$, then \mathbf{F} is a $p \times p$ block diagonal with blocks defined by

the \mathbf{F}_i matrices in (6.1c); \mathbf{G} is a $p \times p$ block diagonal matrix with blocks defined by the corresponding sub-system matrices \mathbf{G}_i in (6.1c); and η_t is a p dimensional vector containing, in appropriate locations, the white noise input vectors $\eta_{i,t}$ ('system disturbances' in normal SS terminology) to each of the GRW models in (6.1c). These white noise inputs, which provide the stochastic stimulus for parametric change in the model, are assumed to be independent of the observation noise e_t and have a covariance matrix \mathbf{Q} formed from the combination of the individual covariance matrices $\mathbf{Q}_{\eta,i}$. Finally, \mathbf{H}_t is a $1 \times p$ vector of the following form,

$$\mathbf{H}_t = [u_{1,t} \quad 0 \quad u_{2,t} \quad 0 \quad \dots \quad u_{n,t} \quad 0] \quad (6.3b)$$

that relates the scalar observation y_t to the state variables which are defined by (6.3a)(ii), so that it represents the TVPR model (6.1a) with each parameter defined as a GRW process. In the case of the scalar RW and AR(1) models for parameter variation, the alternate zeros are simply omitted.

The SS formulation in (6.3a) is particularly well suited to estimation based on optimal, time variable parameter, recursive estimation, in which the time variable parameters (acting as surrogate states in this SS formulation) are estimated sequentially while working through the data in temporal order (usually termed forward-pass filtering). In the off-line situation, where all the time series data are available for analysis, this filtering operation is accompanied by optimal recursive smoothing (see e.g. Young, 1984, 1999a). Here the estimates obtained from the forward-pass filtering algorithm are updated sequentially while working through the data in *reverse* temporal order (usually termed backward-pass smoothing) using a backward-recursive *fixed interval smoothing* (FIS) algorithm, where the fixed interval is the interval covered by the total sample size N .³

The reason for this two-pass approach is easy to understand. The forward pass filtering estimate of \mathbf{x}_t , which defines the estimated TVPs, can be denoted by $\hat{\mathbf{x}}_{t|t}$ (or simply $\hat{\mathbf{x}}_t$, for convenience) since it represents the estimate at sample t given only the data up to and including sampling instant t . However, under our assumption that the parameters evolve stochastically according to the equation (6.1c), a superior smoothed estimate $\hat{\mathbf{x}}_{t|N}$ exists and can be generated by the FIS algorithm, in which the estimate at t is based on *all* the data over the sampling interval $t = 1, 2, \dots, N$. As a result, the phase lag associated with the forward pass filtering estimate (since it cannot anticipate any change until the evidence for change in the series has been processed) is eliminated on the backward smoothing pass. Thus, any variation in the parameters is estimated as it occurs, without any lag effect (indeed, it may even be anticipated if the smoothing effect is

³On-line 'fixed lag smoothing' is also possible, where the recursive estimation works in a forward-pass, filtering mode but with smoothed estimates provided at every sampling instant t over a finite interval of l samples into the past (i.e., over the interval $t - l$ to t), but this is not discussed here.

substantial, as it can be in high noise situations). This proves particularly useful in operations such as interpolation over gaps in the data, estimation and removal of individual components from the data (signal extraction), and seasonal adjustment.

The complete recursive filtering-smoothing algorithm takes the following form:

Forward Pass Filtering Equations

Prediction:

$$\begin{aligned}\hat{\mathbf{x}}_{t|t-1} &= \mathbf{F}\hat{\mathbf{x}}_{t-1} \\ \mathbf{P}_{t|t-1} &= \mathbf{F}\mathbf{P}_{t-1}\mathbf{F}^T + \mathbf{G}\mathbf{Q}_r\mathbf{G}^T\end{aligned}$$

Correction:

$$\begin{aligned}\hat{\mathbf{x}}_t &= \hat{\mathbf{x}}_{t|t-1} + \mathbf{P}_{t|t-1}\mathbf{H}_t^T[1 + \mathbf{H}_t\mathbf{P}_{t|t-1}\mathbf{H}_t^T]^{-1}\{y_t - \mathbf{H}_t\hat{\mathbf{x}}_{t|t-1}\} \\ \hat{\mathbf{P}}_t &= \mathbf{P}_{t|t-1} - \mathbf{P}_{t|t-1}\mathbf{H}_t^T[1 + \mathbf{H}_t\mathbf{P}_{t|t-1}\mathbf{H}_t^T]^{-1}\mathbf{H}_t\mathbf{P}_{t|t-1}\end{aligned}\quad (6.4a)$$

The FIS algorithm is in the form of a backward recursion operating from the end of the sample set to the beginning.

Backward Pass Fixed Interval Smoothing Equations

$$\begin{aligned}\hat{\mathbf{x}}_{t|N} &= \mathbf{F}^{-1}[\hat{\mathbf{x}}_{t+1|N} + \mathbf{G}\mathbf{Q}_r\mathbf{G}^T\mathbf{L}_t] \\ \mathbf{L}_t &= [\mathbf{I} - \mathbf{P}_{t+1}\mathbf{H}_{t+1}^T\mathbf{H}_{t+1}]^T[\mathbf{F}^T\mathbf{L}_{t+1} - \mathbf{H}_{t+1}^T\{y_{t+1} - \mathbf{H}_{t+1}\hat{\mathbf{x}}_{t+1}\}] \\ \hat{\mathbf{P}}_{t|N} &= \mathbf{P}_t + \mathbf{P}_t\mathbf{F}^T\mathbf{P}_{t+1|t}^{-1}[\mathbf{P}_{t+1|N} - \mathbf{P}_{t+1|t}]\mathbf{P}_{t+1|t}^{-1}\mathbf{F}\hat{\mathbf{P}}_t\end{aligned}\quad (6.4b)$$

with $\mathbf{L}_N = \mathbf{0}$. Note that other formulations of FIS are available and can have advantages in certain cases (see e.g. Young, 1984): e.g., the recursive update equation for $\hat{\mathbf{x}}_{t|N}$ can provide an update to the forward pass, filtered estimate $\hat{\mathbf{x}}_t$.

In these algorithms, the $p \times p$ noise variance ratio (NVR) matrix \mathbf{Q}_r and the $p \times p$ matrix $\hat{\mathbf{P}}_t$ are defined as follows:

$$\mathbf{Q}_r = \frac{\mathbf{Q}}{\sigma^2} \quad ; \quad \hat{\mathbf{P}}_t = \frac{\mathbf{P}_t^*}{\sigma^2}\quad (6.4c)$$

where \mathbf{P}_t^* is the error covariance matrix associated with the state estimates which, in the current TVP context, defines the estimated uncertainty in the parameters. For simplicity, it is normally assumed that the NVR matrix \mathbf{Q}_r is diagonal, although this is not essential. The NVR parameters that characterize \mathbf{Q}_r (and any other hyper-parameters in the SS model (6.3a)) are unknown prior to the analysis and clearly need to be estimated on the basis of the time series data y_t and $u_{i,t}, i = 1, 2, \dots, n$ before the filtering and smoothing algorithms can be utilized. The optimization of these 'hyper-parameters' is discussed in the next sub-section.

6.2.1 Maximum Likelihood (ML) Optimization of Hyperparameters

The approach to ML optimization based on prediction error decomposition is derived from the work of Scheppe (1965), who showed how to generate likelihood functions for Gaussian signals using the Kalman filter (see also Bryson and Ho, 1969; page 389). Its importance in the present UC context was probably first recognized by Harvey (1981) and Kitagawa (1981). Since then, it has become one of the two standard approaches to the problem (the other being the *expectation and minimization* (EM) algorithm: Dempster *et al.*, 1977).

If n is the number of TVPs being estimated, then the recursive TVP least squares estimation algorithm (6.4a), with given values for the hyperparameters, will yield the one-step-ahead prediction errors (also termed the innovations or recursive residuals), ε_t , where

$$\varepsilon_t = y_t - \mathbf{H}_t \hat{\mathbf{x}}_{t|t-1} \quad t = 1, 2, \dots, N. \quad (6.5)$$

If the first n observations are regarded as fixed, the log-likelihood function of y_{n+1}, \dots, y_N can be defined in terms of the *prediction error decomposition*, i.e.,

$$\begin{aligned} \log L = & \frac{-(N-n)}{2} \log(2\pi) - \frac{1}{2} \log(\sigma^2) - \frac{1}{2} \sum_{t=n+1}^N \log(1 + \mathbf{H}_t \mathbf{P}_{t|t-1} \mathbf{H}_t^T) \\ & - \frac{1}{2\sigma^2} \sum_{t=n+1}^N \frac{\varepsilon_t^2}{1 + \mathbf{H}_t \mathbf{P}_{t|t-1} \mathbf{H}_t^T} \end{aligned} \quad (6.6)$$

where it can be shown that $\sigma^2(1 + \mathbf{H}_t \mathbf{P}_{t|t-1} \mathbf{H}_t^T)$ is the variance of ε_t , so that the last term in (6.6) is based on the sum of squares of the normalized one-step-ahead prediction errors. Now the ML estimate of σ^2 , conditional on the hyperparameters, is given by

$$\hat{\sigma}^2 = \frac{1}{N-n} \sum_{t=n+1}^t \frac{\varepsilon_t^2}{1 + \mathbf{H}_t \mathbf{P}_{t|t-1} \mathbf{H}_t^T} \quad (6.7)$$

so that it can be estimated in this manner and ‘concentrated out’ of the expression (6.6) by substituting (6.7) into (6.6), to yield the following expression for the ‘concentrated likelihood’:

$$\log(L_c) = -\frac{N-n}{2} \log(2\pi+1) - \frac{1}{2} \sum_{t=n+1}^N \log(1 + \mathbf{H}_t \mathbf{P}_{t|t-1} \mathbf{H}_t^T) - \frac{N-n}{2} \log(\hat{\sigma}^2) \quad (6.8)$$

which needs to be maximized with respect to the unknown hyperparameters in order to obtain their ML estimates.

Since (6.8) is nonlinear in the hyper-parameters, the likelihood maximization needs to be carried out numerically. Consequently, it is more convenient to remove the constant term (since it will play no part in the optimization) and multiply (6.8) by -2 , to yield

$$\log(L_c) = \sum_{t=n+1}^N \log(1 + \mathbf{H}_t \mathbf{P}_{t|t-1} \mathbf{H}_t^T) + (N - n) \log(\hat{\sigma}^2), \quad (6.9)$$

which then needs to be minimized. This minimization is accomplished by initiating the optimization with the hyper-parameter estimates either selected by the user or set to some default values (in both cases, ensuring that the resulting optimization does not converge on a local minimum). The recursive TVP estimation algorithm (6.4a) is used repeatedly to generate the one step ahead prediction errors and, thence, the log-likelihood value in (6.9) associated with the latest selection of hyper-parameter estimates made by the optimization algorithm. The optimization algorithm then adjusts its selection of hyper-parameter estimates in order to converge on those estimates that minimize (6.9). Further details of this and alternative ML optimization procedures are given, for example, in Harvey and Peters (1990). Typical methods that can be used for numerical optimization are the 'fmins' and 'fminu' functions available in the Matlab software system, or their equivalents, although more complex and efficient procedures are available.

6.3 State Dependent Parameter Regression Models

The approach to TVPR estimation discussed in the last section works very well in situations where the parameters $a_{i,t}$ are slowly varying when compared to the observed temporal variation in the measured regression variables $u_{i,t}$ (see e.g., Young, 1999a). If the $a_{i,t}$ are functions of the $u_{i,t}$, however, then the recursive algorithm 4(a)-4(c) is, in general, unable to track the potentially much more rapid variations in the parameters, and a modified approach is necessary. This SDR model takes the following form:

$$y_t = \mathbf{u}_t^T \mathbf{p}_t + e_t, \quad e_t = N(0, \sigma^2) \quad t = 1, 2, \dots, N \quad (6.10a)$$

or

$$y_t = x_t + e_t, \quad \text{where} \quad x_t = \mathbf{u}_t^T \mathbf{p}_t$$

is the noise free output. Here,

$$\begin{aligned} \mathbf{u}_t^T &= [u_{1,t} \quad u_{2,t} \dots u_{n,t}] \\ \mathbf{p}_t &= [a_1(u_{1,t}) \quad a_2(u_{2,t}) \dots a_n(u_{n,t})]^T \end{aligned} \quad (6.10b)$$

and the nomenclature $a_i(u_{i,t})$ denotes that each parameter is potentially a function of its associated regression variable.

Since the parameter vector \mathbf{p}_t is potentially state-dependent, it may vary at a rate commensurate with the temporal variations in the regression variables, and so it cannot be assumed that the simple GRW model (6.1c) is appropriate to describe the parametric variation over time. In particular, it would appear at first sight that the stochastic state space model should include prior information on the nature of the parameter variation if the TVP estimation methodology discussed in previous sections is to work satisfactorily. Fortunately, it is possible to remove this requirement if we resort to the rather unusual procedure, at least within a time series context, of sorting the data in a *non-temporal* order. Then, if the ordering is chosen so that the SDP variations associated with the *sorted* series are smoother and less rapid, it is more likely that a simple GRW process can be utilized to describe their evolution.

For example, if the time series are sorted in some common ascending-order-of-magnitude manner (i.e., the *sort* operation in Matlab), then the rapid natural variations in y_t and the $u_{i,t}$ can be effectively eliminated from the data and replaced, in the sorted data space, by much smoother and less rapid variations. And if the SDPs are, indeed, related to these variables, then they will be similarly affected by the sorting. Following FIS estimation, however, these SDP estimates can be unsorted (a trivial *unsort* operation to reverse Matlab's *sort*) and their true, rapid variation will become apparent. Of course, the nature of the sorting will affect the efficacy of this approach and it seems likely that there will be an optimum sorting which results in minimum variance estimates. However, such optimum sorting will naturally depend on the nature of the state dependency and its definition would require some sort of iterative estimation procedure. In practical terms, therefore, the common ascending order sorting and unsorting operations seem the most straightforward and will be used here.

One obvious requirement with this new approach to SDP estimation is that the sorting of data, prior to FIS estimation, must be *common to all of the variables in the relationship* (6.10a). If an ascending order strategy is selected, therefore, it is necessary to decide upon which variable in the model the sorting should be based. The simplest strategy is to sort according to the ascending order of the dependent variable y_t . Depending on the nature of each SDP in the vector \mathbf{p}_t , however, a single variable sorting strategy such as this does not often produce satisfactory results; indeed, it will only work well if the state dependency is associated with the variations in y_t or some related variable. If this is not the case, then a more complicated, but still straightforward, back-fitting procedure can be exploited. Here, each parameter is estimated *in turn*, based on the partial residual series obtained by subtracting all the other terms on the right hand side of (6.10a) from y_t . At each such back-fitting iteration, the sorting can then be based on the single variable associated with the current SDP being estimated (in the current SDR context, each $u_{i,t}$ in turn).

Since the SDP estimates resulting from this back-fitting algorithm are

themselves time series, it will be noted that the algorithm constitutes a special form of non-parametric estimation and, as such, can be compared with other non-parametric methods, such as the generalized additive modeling (GAM) approach of Hastie and Tibshirani (1996). However, in both conceptual and algorithmic terms, the SDP approach described here is significantly different from this earlier approach and seems more appropriate to the estimation of nonlinear, stochastic, dynamic models. Moreover, the recursive methodology, on which SDP estimation is based, is couched in optimal maximum likelihood terms that seem more elegant and flexible than the ‘scatter-plot smoothing’ procedures suggested by Hastie and Tibshirani.

The back-fitting algorithm for the SDP model (6.10) takes the following form:

1. Assume that FIA estimation⁴ has yielded prior TVP estimates $\hat{a}_{i,t|N}$, $i = 1, 2, \dots, n$ of the SDPs.
2. Iterate: $i = 1, 2, \dots, n$; $k = 1, 2, \dots, k_c$
 - (a) form the partial residual $y_t^i = y_t - \sum_{j \neq i} u_{j,t} \hat{a}_{j,t|N}^k$;
 - (b) sort both y_t^i and $u_{i,t}$ according to the ascending order of $u_{i,t}$;
 - (c) obtain an FIS estimate $\hat{a}_{i,t|N}^k$ of $a_{i,t}$ in the partial residual relationship $y_t^i = a_{i,t} u_{i,t}$.
3. Continue 2 (each time forming the partial residual and then sorting according to the current right-hand side variable $u_{i,t}$, prior to FIS estimation), until iteration k_c , when the individual SDPs (which is each a time-series of length N) have not changed significantly according to some chosen criterion. The smoothing hyperparameters required for FIS estimation at each stage are optimized by maximum likelihood, as explained in Section 6.2 and discussed further later.

Note that the ML optimization can be carried out in various ways: after every complete iteration (each involving n FIS operations) until convergence is achieved; only at the initial complete iteration, with the hyperparameters maintained at these values for the rest of the back-fitting; or just on the first two iterations. The last choice seems most satisfactory in practice, since very little improvement in convergence occurs if optimization is continued after this stage. Normally, convergence is completed after only a few iterations, although it can be more lengthy in some circumstances (see Simulation Example 1 and the discussion on this topic in the

⁴As a default, these can be simply the *constant* least squares parameter estimates, since the convergence of the back-fitting procedure is not too sensitive to the prior estimates, provided they are reasonable: see simulation example 1 below.

Conclusions).

Simulation Example 1

This example used data generated by the following fourth order SDR relationship:

$$y_t = f_1\{u_{1,t}\} + f_2\{u_{2,t}\} + f_3\{u_{3,t}\} + f_4\{u_{4,t}\} + e_t \quad t = 1, 2, \dots, N \quad (6.11)$$

or

$$y_t = a_1\{u_{1,t}\}.u_{1,t} + a_2\{u_{2,t}\}.u_{2,t} + a_3\{u_{3,t}\}.u_{3,t} + a_4\{u_{4,t}\}.u_{4,t} + e_t \quad (6.12)$$

where $N = 1000$ and the $u_{i,t}, i = 1, 2, 3, 4$ are the following random variables

$$u_{1,t} = N(0, 1); \quad u_{2,t} = U(0, 1); \quad u_{3,t} = U(0, 1); \quad u_{4,t} = U(0, 1); \quad (6.13)$$

in which $N(a, b)$ denotes a normal distribution with mean a and variance b ; while $U(a, b)$ denotes a uniform distribution in the range a to b . The nonlinear functions of these variables $f_i, i = 1, 2, 3, 4$ are defined as follows:

$$\begin{aligned} f_1\{u_{1,t}\} &= \begin{cases} u_{1,t} & \text{for } -0.5 \leq u_{1,t} \leq 0.5, \\ -0.5 & \text{for } u_{1,t} < -0.5, \\ 0.5 & \text{for } u_{1,t} > 0.5; \end{cases} \\ f_2\{u_{2,t}\} &= u_{2,t}^2 \\ f_3\{u_{3,t}\} &= u_{3,t}^3 \\ f_4\{u_{4,t}\} &= \begin{cases} x_{4,t}^4 & \text{for } u_{4,t}^4 \leq 0.4, \\ 0.4 & \text{for } u_{4,t}^4 > 0.4. \end{cases} \end{aligned} \quad (6.14)$$

Finally, the observational noise $e_t = g.N(0, 1)$, with the noise gain $g = 0.12$, yields a noise level of approximately 20% by standard deviation.

In the back-fitting iterations, optimization is carried out at the first and second iterations: thereafter, the NVR hyper-parameters are maintained constant for the thirty-five iterations required for convergence in this case, at the following optimized values:

$$\begin{aligned} NVR\{a_1(x_{1,t})\} &= 0.0153; & NVR\{a_2(x_{2,t})\} &= 0.00562 \\ NVR\{a_3(x_{3,t})\} &= 0.0087; & NVR\{a_4(x_{4,t})\} &= 0.00596. \end{aligned}$$

Convergence in terms of the model fit appears virtually complete after only a few iterations; however, this leaves some biases in the estimates and these are only removed after the full thirty-five iterations. The resulting FIS estimates of the four nonlinear functions are shown in Figures. 6.1 to 6.4, where $\hat{f}_i\{u_{i,t}\} = \hat{a}_{i,t|N}.u_{i,t}$ is plotted as a full line against the regression variable $u_{i,t}$ in each case, $i = 1, 2, 3, 4$; the dashed lines are the estimated standard error (*se*) bounds (95% confidence intervals); and the dotted line is the true nonlinear function. It is clear that the nonlinear functions have




FIGURE 6.1. FIS estimated (solid) and actual (dotted) limited linear function. Error bounds dashed.




FIGURE 6.2. FIS estimated (solid) and actual (dotted) quadratic function. Error bounds dashed.

all been estimated well, with the true nonlinearity lying well within the se bounds. As a result, the model also explains the dependent variable y_t very well, as shown in Figure 6.5, which compares the model output \hat{x}_t with y_t ,




FIGURE 6.3. FIS estimated (solid) and actual (dotted) cubic nonlinearity. Error bounds dashed.




FIGURE 6.4. FIS estimated (solid) and actual (dotted) limited quartic nonlinearity. Error bounds dashed.

where, $\hat{x}_1 = \mathbf{u}_t^T \hat{\mathbf{p}}_{t|N}$ and

$$\hat{\mathbf{p}}_{t|N} = [\hat{a}_{1,t|N} \quad \hat{a}_{2,t|N} \quad \dots \quad \hat{a}_{n,t|N}]^T. \quad (6.15)$$

The error $\hat{e}_t = y_t - \hat{x}_t$ (full line, +70) provides an excellent FIS estimate of the actual noise e_t . The coefficient of determination based on \hat{e}_t is $R^2 = 0.968$; and that based on the error between x_t (the noise-free output, i.e.,




FIGURE 6.5. SDPR model output (solid) compared with noisy dependent variable (dotted).




FIGURE 6.6. FIS estimates of SDPs in the sorted space.

$y_t - e_t$) and \hat{x}_t is $R^2 = 0.997$; i.e., the model can reconstruct 99.7% of the unobserved noise-free signal, x_t .

Finally, it is interesting to compare the FIS estimates of the SDPs in the sorted space and the natural observation (original unsorted) space, respectively, as shown in Figures 6.6 and 6.7. Here, the slow movement




FIGURE 6.7. FIS estimates of SDPs in the observation space.

of the estimated SDPs in the sorted space (Fig. 6.6) contrasts with the very rapid, real variations in the observation space (Fig. 6.7). It is clear how effectively the sorting operations have modified the natural changes in the state-dependent parameters so that, in the sorted space, the RW model is able to characterize their much slower variations for the purposes of FIS estimation. Without these sorting operations and the associated back-fitting algorithm, the SDPs change much too rapidly to be estimated, unless prior knowledge of state dependency is available.

Before moving on to the estimation of SDP models for stochastic dynamic systems, it should be noted that, although the SDR model (10a,b) is formulated so that each SDP is considered primarily as a function of its associated regression variable, this does not preclude the investigation of models in which the SDP is a function of another known (state) variable, except that then the sorting will relate to this other variable. This is discussed further in subsequent sections of the chapter.

6.4 State-Dependent Parameter Dynamic Models

Although the TVP regression model has been referred to as a dynamic linear regression (DLR) model (see footnote 1 and Young, 1999a), it is not really dynamic in a systems sense. Similarly, while the SDR model can be used very effectively in the processing and modeling of time series data, it is also basically static in form. In this section, we move on to models that are able to describe truly dynamic system behavior, namely SDP versions of

two well-known stochastic, dynamic models: the *state-dependent parameter auto-regressive exogenous variables* (SDARX) input-output model; and its simpler univariate relative, the *state-dependent parameter auto regressive* (SDAR) univariate model.

6.4.1 The SDARX Model

The SDARX model is a transfer function model with a rather simple signal topology: It relates one or more measured input signals $u_i, i = 1, 2, \dots, p$ to a single measured output signal y_t . In the case of a single input signal (denoted here as u_t for simplicity) and using nomenclature similar to that used in Equation (10a-b), the model equation takes the following form⁵,

$$y_t = \mathbf{z}_t^T \mathbf{p}_t + e_t \quad e_t = N(0, \sigma^2) \quad (6.16a)$$

where,

$$\begin{aligned} \mathbf{z}_t^T &= [y_{t-1} \ y_{t-2} \ \dots \ y_{t-n} \ u_{t-\delta} \ \dots \ u_{t-\delta-m}] \\ \mathbf{p}_t &= [a_1(y_{t-1}) \ a_2(y_{t-2}) \ \dots \ a_n(y_{t-n}) \ b_0(u_{t-\delta}) \ \dots \ b_m(u_{t-\delta-m})]^T \end{aligned} \quad (6.16b)$$

and δ is a pure time delay, measured in sampling intervals, which is introduced to allow for any temporal delay that may occur between the incidence of a change in u_t and its first effect on y_t .

As in the case of its constant parameter progenitor, the SDARX model can be considered in estimation terms from a least squares recursive estimation standpoint. Consequently, the FIS estimation of the SDPs in the model can follow the same algorithmic procedures described in the previous section. Thus, to allow, once again, for the estimation of the potentially rapid variation in the parameters, the data are sorted prior to estimation and, in the case of a SDARX model with more than one state dependency, the iterative back-fitting algorithm has to be introduced. The following simulation example demonstrates the efficacy of this approach for a first order ($n = 1, m = 1, \delta = 0$) SDARX model containing nonlinearities in both the y_{t-1} (feedback) and u_t (input) pathways.

Simulation Example 2

This example used data generated by the following first order SDARX relationship:

$$y_t = a_1(y_{t-1}) \cdot y_{t-1} + b_0(u_t) \cdot u_t + e \quad e_t = N(0, 0.000025) \quad u_t = N(0, 0.0064) \quad (6.17)$$

where

$$a_1(y_{t-1}) = 2.0\{1 - y_{t-1}\}; \quad b_0(u_t) = 10u_t^2.$$

⁵The extension to more than one input variable is obvious and presents no difficulties.

When written in the nonlinear functional form

$$y_t = f_1(y_{t-1}) + f_2(u_t) + e_t$$

or

$$y_t = 2.0y_{t-1} - 2.0y_{t-1}^2 + 10u_t^3 + e_t, \quad (6.18)$$

this is revealed as the SDP formulation of the forced logistic growth equation with a measured input forcing signal in the form of a normally distributed white noise sequence passed through a cubic law nonlinearity. A




FIGURE 6.8. Typical response of first order SDARX nonlinear stochastic system.

typical 500-sample response of this system is shown in Figure 6.8, where the output y_t is in the top panel and the input u_t in the lower panel. Here, the percentage noise/signal, based on standard deviations (i.e., $100 \times \{sd(e_t)/sd(10u_t^3)\}$), is approximately 30%. The results of the SDARX analysis for a total sample size of $N = 2000$ are shown in the two graphs in the top panels of Figures 6.9 and 6.10. As in the previous simulation example, the hyper-parameter optimization is carried out at the first and second iterations; thereafter, the NVR hyper-parameters are maintained constant, for the four iterations required for convergence in this case, at the following optimized values:

$$NVR\{a_1(y_{t-1})\} = 0.67; \quad NVR\{b_0(u_t)\} = 6.24.$$

The top panel of Figure 6.9 shows the estimated SDPs obtained using these optimized NVR values, with $\hat{a}_1(y_{t-1}|N)$ in the left graph and $\hat{b}_0(u_t|N)$ in the right. The dotted lines are the actual SDP relationships and it is clear that the state dependency has been estimated well in both cases.

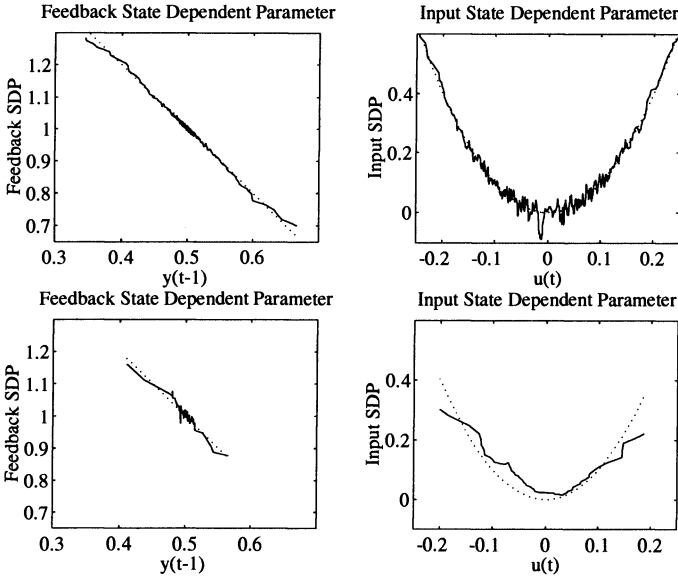


FIGURE 6.9. Feedback state-dependent parameters (left) and input state-dependent parameters (right).

The response of the estimated SDARX model \hat{y}_t can be generated in two ways. First, directly from the equation,

$$\hat{y}_t = \hat{a}_1(y_{t-1}|N).y_{t-1} + \hat{b}_0(u_t|N).u_t \tag{6.19a},$$

in the usual, regression-like manner, where it will be noted that the y_{t-1} on the right-hand side is based on the actual measurements y_{t-1} and not the modeled ones. This suggests that the model explains 95.7% of the output y_t ; i.e., $R^2 = 0.957$. However, since the SDARX model is truly dynamic, this coefficient of determination (COD) is a little misleading. It is more sensible to base the COD on the *simulated* model output generated from the equation,

$$\hat{y}_t = \hat{a}_1(y_{t-1}|N).\hat{y}_{t-1} + \hat{b}_0(u_t|N).u_t. \tag{6.19b}$$

This is easily generated by a Simulink simulation model using lookup tables based on the non-parametric SDP estimation results. Here, the time vector and white noise from the workspace (t, e) are passed through two SDP models; the true model with actual nonlinearities and the estimated model. The middle panel of Figure 6.10 shows the simulated model output \hat{y}_t obtained in this manner. The COD obtained in relation to the actual output y_t , including the effects of the noise e_t , is $R^2_T = 0.935$, where the subscript T is introduced to differentiate this simulation-based COD from the more normal regression-based COD R^2 . However, if this simulation model output is compared with the noise-free output (i.e., $e_t = 0, \forall t$), as shown in the top panel of Figure 6.10, then the error (bottom panel) is very small and the COD is $R^2_T = 0.980$. Clearly, on all measures, the




FIGURE 6.10. Simulink realizations of actual (top) and SDP model (middle) outputs. The bottom graph shows the error.

estimated SDARX model (6.19b) provides an excellent representation of the system (6.18).

The model (6.19b) is non-parametric, in the sense that the nonlinearities are in the form of graphs defined by the FIS-estimated SDPs. However, as pointed out previously, the FIS-based estimation can also be considered an identification stage in the analysis, with the non-parametric estimates of the nonlinearities providing the basis for final parametric estimation, based on some specific parameterization of these identified nonlinearities. Even without our prior knowledge in this simulation example, it is fairly obvious from Figure 6.9 that the input and feedback SDPs are linear and quadratic functions of the associated variables (i.e., the nonlinearities are quadratic and cubic functions, respectively). Thus, it is straightforward to obtain these parametric estimates, either by least squares (LS) or weighted least squares (WLS) estimation based on the SDP estimation results, as discussed in Young (1993a) and Young and Beven (1994); or, preferably, by direct estimation from the data using the identified model structure,

$$y_t = ay_{t-1} - by_{t-1}^2 + cy_t^3 + e_t. \quad (6.20)$$

In this case, the two sets of estimation results are given, respectively, by

$$\hat{a} = 1.991(0.004); \quad \hat{b} = 1.981(0.008); \quad \hat{c} = 10.02(0.028); \quad (6.21a)$$

$$\hat{a} = 1.999(0.011); \quad \hat{b} = 1.998(0.022); \quad \hat{c} = 10.01(0.123). \quad (6.21b)$$

Although the standard errors on the initial SDP-based estimates (6.21a) tend to be too optimistic, the parametric estimates themselves are very

close to the true values, showing the efficacy of the SDP estimation stage in identifying the nature of the nonlinearities. Indeed, the SDP estimates obtained for only $N = 200$ samples, as shown by the two graphs in the lower panel of Figure 6.9, are good enough to identify the form of the nonlinear functions and they yield the following estimates (c.f. (6.21))

$$(i) \hat{a} = 2.046(0.066); \hat{b} = -2.085(0.132); \hat{c} = 8.31(0.433);$$

$$(ii) \hat{a} = 2.044(0.070); \hat{b} = 2.088(0.140); \hat{c} = 9.902(0.721). \quad (6.22)$$

More statistically efficient ML estimates could be obtained in this case but this seems hardly necessary since the results in (6.21) and (6.22) are both satisfactory.

Real Example 1: Modeling Rainfall-Flow Dynamics

One successful practical application of SDARX modeling is in the characterization the nonlinear relationship between rainfall and flow (discharge) in a river catchment (see e.g. Young, 1993a; Young and Beven, 1994; Young et al, 1998; Young, 1999c). Fig. 6.11 is a plot of a typical, daily rainfall-flow data set from a catchment in the United States. It is well known that the primary nonlinearity between the occurrence of rainfall in the catchment area and the subsequent increase of flow in the river occurs because of soil moisture and evapo-transpiration effects. These effects reduce the effective




FIGURE 6.11. Coweeta data: flow (top) and rainfall (bottom).

tive level of the rainfall and the relationship between the measured rainfall and this effective rainfall (or rainfall excess) is quite nonlinear. In the case

of soil moisture, for example, if the catchment is very dry because little rain has fallen for some time, then most new rainfall will be absorbed by the dry soil and little, if any, will be effective in promoting increases in river flow. Subsequently, however, if the soil moisture increases because of further rainfall, the run-off of excess water from the catchment rises and the flow increases because of this. In other words, the effect of rainfall on flow depends on the antecedent conditions in the catchment and a similar rainfall event occurring at different times and under different soil-moisture conditions can yield markedly different changes in river flow.

The previous studies cited have shown that the simplest model that can reproduce these kind of effects and characterize the relationship between rainfall and flow, such as that illustrated in Figure 6.11, is the following first order SDARX model

$$y_t = a_1(y_t) \cdot y_{t-1} + b_0(y_t) \cdot u_{t-\delta} + e_t \quad (6.23)$$

where now u_t is the measured rainfall in mm; y_t is the flow measured in mm equivalent⁶; and δ is introduced to account for any purely advective time delay effects between the occurrence of rainfall and its first effect on flow in the river. In (6.23) it will be noted that, in contrast to previous examples, the state dependency is not in terms of the rainfall u_t but of the flow y_t . This rather paradoxical aspect of the model and its physical explanation, are discussed briefly later. For the moment, let us consider the results obtained from SDARX estimation.

For the data in Figure 6.11, $\delta = 0$ and, since the state dependency being investigated is in terms of the flow variable y_t , all the variables in (6.23) are sorted in relation to the sorted order of y_t (again based on the ascending order of magnitude). Since this is the dependent variable, only a single initial sorting operation is necessary and back-fitting is not required. The NVR hyper-parameters are optimized as follows, under the assumption that $a_1(y_t)$ and $b_0(y_t)$ vary as RW and IRW processes, respectively:

$$NVR\{a_1(y_t)\} = 2.43 \times 10^{-15}; \quad NVR\{b_0(y_t)\} = 7.13 \times 10^{-9}$$

The RW process is selected for $a_1(y_t)$ because initial analysis shows that it does not vary much and may well be constant, in which case the RW model is more appropriate (since it defaults on a constant estimate if the $NVR = 0$). As we see, this is confirmed by the optimized $NVR\{a_1(y_t)\}$, which is very small, resulting in an FIS estimate of a_1 that hardly changes over the observation interval. This is consistent with the hydrological interpretation of the model: the parameter mainly affects the recession in the level of river flow following the completion of a rainfall event and, since

⁶i.e., the equivalent depth is the flow volume was distributed over the whole catchment area.

this recession pattern does not appear to change much over time, it is sensible that the parameter should be constant. Indeed, if the NVR for this parameter is set to zero, thus constraining the estimate of a_1 to be exactly constant, then the other optimized $NVR\{b_0(y_t)\}$ does not change at all, so the FIS estimate $\hat{b}_0(y_t|N)$ of the input parameter $b_0(y_t)$ remains virtually the same.

In this situation, the constant a_1 estimate is obtained as $\hat{a}_1 = 0.769(0.007)$. In contrast to a_1 , however, the SDP estimate $\hat{b}_0(y_t|N)$, as shown in Figure 6.12 as a function of y_t , reveals a significant nonlinearity between the rainfall input and its effect on flow. In particular, $\hat{b}_0(y_t|N)$ increases




FIGURE 6.12. FIS estimate of SDP: SE bounds (dashed); polynomial approximation (dash-dot).

steadily with flow, with this increase tending to get less at higher flows and eventually leveling out, at approximately 0.3, for flows greater than 30mm. Note that the flow dependent steady-state gain of the SDARX model, defined as $G = \hat{b}_0(y_t|N)/(1 - \hat{a}_1)$ is less than unity most of the time; it is only greater than unity for the twenty-two highest flow measures in excess of 19.5mm. In other words, the model suggests that, most of the time, the physical processes between rainfall and river flow involve losses of water to the catchment. During these highest flow events however, the flow output is *transiently* larger than the rainfall input.

Having identified the form of the parametric state dependency and, thus, the nonlinearity in the rainfall-flow process, it is straightforward to define the effective rainfall $u_{e,t}$ in the following manner:

$$\begin{aligned} u_{m,t} &= \hat{b}_0(y_t|N) \cdot u_t \\ u_{e,t} &= c \cdot u_{m,t}, \end{aligned} \tag{6.24a}$$

where $c = \sum y_t / \sum u_{m,t}$ is introduced, for convenience, so that the total effective rainfall is the same as the total flow.

If most of the nonlinearity in the system has been removed by the definition of this effective rainfall measure, then the dynamic characteristics between $u_{e,t}$ and y_t should be linear. This is indeed the case, and the following linear TF model is identified and estimated well from the $u_{e,t} \sim y_t$ data,

$$y_t = \frac{\hat{b}_0 + \hat{b}_1 z^{-1} + \hat{b}_2 z^{-2}}{1 + \hat{a}_1 z^{-1} + \hat{a}_2 z^{-2}} u_{e,t} + \xi_t \quad (6.24b)$$

where ξ is a colored noise process (fourth-order autoregressive) and the estimates of the parameters in the TF polynomials are

$$\hat{a}_1 = -1.545(0.036); \quad \hat{a}_2 = 0.560(0.034);$$

$$\hat{b}_1 = 0.162(0.005); \quad \hat{b}_2 = -0.142(0.012); \quad \hat{b}_3 = -0.0053(0.008);$$

with the standard errors shown in parentheses. Note that once the nonlinearity has been identified and accounted for in the model, the remaining transient dynamics are revealed to be of second order, rather than the first-order form used in the initial SDARX modeling phase. The physical significance of this is discussed later.

The model (6.24b) provides a rather good explanation the rainfall-flow data: the deterministic model output \hat{x}_t , defined as

$$\hat{x}_t = \frac{\hat{b}_0 + \hat{b}_1 z^{-1} + \hat{b}_2 z^{-2}}{1 + \hat{a}_1 z^{-1} + \hat{a}_2 z^{-2}} u_{e,t}$$

explains 92.3% of the flow y_t (i.e., $R_T^2 = 0.923$), and the comparison of \hat{x}_t and y_t is presented in the upper panel of Figure 6.13, with the error shown above (+60mm). The estimation method used to obtain these results is the SRIV algorithm (see Young, 1984, 1985 and the references therein). This is an algorithm based on the optimal instrumental variable (IV) approach to stochastic estimation and it is particularly effective in TF modeling of this kind. The superiority of the nonlinear model (6.24b), in comparison with a purely linear TF model, is illustrated in the lower panel of Figure 6.13, where the linear TF model output is compared with y_t , again with the error shown above (+60mm). The R_T^2 in this case is only 0.76.

The model (24a-b) is based on the effective rainfall $u_{e,t}$ obtained by reference to the non-parametric FIS estimate $\hat{b}_0(y_t|N)$. An alternative, constant parameter model can be obtained by parameterizing the $\hat{b}_0(y_t|N) \sim y_t$ relationship in some manner. Here, two possibilities have been examined: first, approximating the relationship by a simple polynomial in y_t ; and second, using a more powerful approximation based on radial basis functions. Both of these approaches work well but a simple, second-order, polynomial representation of the form

$$\hat{b}_0(y_t|N) = p_2 y_t^2 + p_1 y_t + p_0$$




FIGURE 6.13. SDARX model (top) and linear model (bottom) outputs compared with flow data.

produces the approximation shown as the dash-dot line in Figure 6.12, with a COD of $R^2 = 0.992$. This is perfectly adequate in this case and yields the following, fully parametric model:

$$y_t = \frac{\hat{b}_0 + \hat{b}_1 z^{-1} + \hat{b}_2 z^{-2}}{1 + \hat{a}_1 z^{-1} + \hat{a}_2 z^{-2}} u_{e,t} + \xi_t \quad u_{e,t} = \{p_2 y_t^2 + p_1 y_t + p_0\} \cdot u_t \quad (6.25)$$

with the parameter estimates

$$\begin{aligned} \hat{a}_1 &= -1.552(0.037); & \hat{a}_2 &= 0.467(0.035); & \hat{b}_3 &= -0.0032(0.008); \\ \hat{b}_1 &= 0.162(0.005); & \hat{b}_2 &= -0.145(0.013); & \hat{p}_0 &= 0.00745(2.7 \cdot 10^{-4}). \\ \hat{p}_2 &= -0.000169(2.8 \cdot 10^{-6}); & \hat{p}_1 &= 0.0145(7.1 \cdot 10^{-5}); & & \end{aligned}$$

This model explains 92.2% of the flow y_t (i.e., $R_T^2 = 0.922$), only marginally less than the non-parametric SDARX model (6.24b). Note, however, that this polynomial representation, while very good for the measured flows up to 37mm, is not adequate if higher flows are encountered, since the polynomial curve reduces the effective rainfall for these higher flows. Consequently, it would be best to introduce a superior parameterization that levels out for higher flows if the model is to be used for forecasting purposes where higher flows might be encountered. Nevertheless, the simple polynomial law suffices for the present illustrative purposes.

At this point, the model (6.25) could form the basis for a final stage of estimation in which all of the parameters, including the parameters in the AR(4) model for the noise ξ , are estimated concurrently by some more efficient estimation method (e.g., ML by prediction error decomposition, as mentioned in the Introduction: see also the blow-fly modeling example

in Young, 1999d). However, the model (6.25) is adequate for the present illustrative purposes and, in any case, it performs rather well in a validation sense. For example, when this same model is applied to the later rainfall-flow data from samples 13001-14000, without any re-estimation, it explains 89.1% of the flow ($R_T^2 = 0.891$), even though this data set occurs some thirty three years later. This is only increased to $R_T^2 = 0.903$ if the model is re-estimated against the new data. In this case, the purely linear model has an $R_T^2 = 0.761$, so the nonlinear model remains clearly superior.

An important aspect of modeling real systems is the physical interpretation of the model in terms that are acceptable to other scientists working in the area of study (here, hydrology). In previous publications (e.g., Young, 1998a and the references therein), the author has stressed the need for such a data-based mechanistic (DBM) method of modeling in which the model is not only satisfactory in statistical terms but has a credible and clear physical interpretation. This is well illustrated by the present example, where the hydrological interpretation of the model (6.25) is straightforward.

1. First, the linear TF part of the model conforms with the classical unit hydrograph theory of rainfall-flow dynamics. Indeed, its unit impulse response is, by definition, the unit hydrograph; and the TF model itself can be seen as a parametrically efficient method of quantifying this unit hydrograph.
2. Second, the TF model can be decomposed by partial fraction expansion into a parallel pathway form which has a clear hydrological interpretation; the three parallel pathways in this case (as defined by the instantaneous response and eigenvalues of the TF model) have residence times of 0 (the instantaneous effect), 1.9 and 26.4 days. It is reasonable to assume that these represent the dominant physical pathways (above and below ground) in which the rainfall eventually reaches the river (for a more detailed explanation and other examples, see Young, 1992, 1993a, 1998a, 1999b,c; Young and Beven, 1994; Young et al, 1997, 1998).
3. Third, the SDP relationship suggests that the parameter is a function of flow. Of course, this is physically impossible but the analysis produces such a clearly defined relationship of this sort that it must have some physical connotations. The most hydrologically reasonable explanation is that the flow is acting as a surrogate for soil moisture. Of course, it would be better to investigate this relationship directly by measuring the soil moisture and incorporating these measurements in the SDP analysis. Unfortunately, it is much more difficult to obtain soil moisture measures and these were not available in the present example. However, Fawcett (1999) has used the SDP approach recently to analyze data from another, smaller catchment, including soil moisture measurements, and this has confirmed the close relationship between the variations in soil moisture variables and flow.

Finally, it should be noted that the present example has been analyzed more fully, but a little differently, in Young (1998a, 1999c), where the effects of temperature are also considered and where the explanation of the data by the SDP model is better than that provided by the present, simpler model.

6.4.2 The SDAR Model

The state-dependent parameter autoregressive (SDAR) model is simply a special example of the SDARX model without deterministic inputs; and it represents an SDP (i.e., nonlinear) form of the well-known linear AR model mentioned in the last section. The model takes the form:

$$y_t = \mathbf{z}_t^T \mathbf{p}_t + e_t, \quad (6.26a)$$

where

$$\begin{aligned} \mathbf{z}_t^T &= [y_{t-1} \quad y_{t-2} \quad \dots \quad y_{t-n}] \\ \mathbf{p}_t &= [a_1(y_{t-1}) \quad a_2(y_{t-2}) \quad \dots \quad a_n(y_{t-n})]^T \end{aligned} \quad (6.26b)$$

Clearly, the same estimation methods discussed for the SDARX model can be applied to this model, a simple, first order simulation example of which the chaotic version of the logistic equation where $a_1(y_{t-1}) = -4(1 - y_{t-1})$, as discussed in Young (1999d). An example based on the SDP analysis of real data is discussed later.

Real Example 1: Analysis of Signals from the Axon of a Squid

This example is based on the analysis of the signal shown in Figure 6.14, which was obtained by Kazu Aihara and Gen Matsumoto from experiments on the giant axon of a squid (Mees et al. 1992). The signal comprises voltage measurements made from a micro-pipette inserted into a giant axon. Squid are used for such experiments because they have large diameter axons; this is because the nerves are not myelinated (insulated), so they need to have large diameter to reduce ion leakage and so maintain the transmission speed. The experiment is done *in vitro* (i.e., the nerves are chopped out of the squid), with the membrane voltage clamped. This is normally referred to as the forced response. However, the forcing is periodic and the response is not; in effect, Figure 6.14 shows a putative chaotic response to a periodic signal. It is analyzed here, however, as a purely stochastic, unforced dynamic process.

First-Order SDAR Model

The simplest SDAR model that could produce the dynamics shown in Figure 6.14 is the following first order SDAR

$$y_t = a_1(y_{t-1})y_{t-1} + e_t \quad (6.27)$$

The FIS estimate of the SDP is shown in Figure 6.15 and the associated nonlinear function in Figure 6.16. These results were obtained with the




FIGURE 6.14. Electrical signal obtained from experimental measurements from a squid giant axon.




FIGURE 6.15. The SDP as a function of the delayed output.

single NVR hyper-parameter optimized at $NVR\{a_1(y_{t-1})\} = 5.01 \times 10^{-6}$, under the assumption of an RW model for the SDP parameter variation (an IRW assumption provides a slightly more smoothed estimate but does not make a significant difference). Once again, the standard error band is shown as the dashed lines on both plots. This model explains 92.2% of the




FIGURE 6.16. Nonlinearity in first order SDAR model of the squid data.




FIGURE 6.17. SDAR model output compared with measured output: error-70 at top.

squid data ($R^2 = 0.922$) and the model output is compared with the data in Figure 6.17.

Figure 6.18 provides a more qualitative but more discerning comparison of the model behavior and the squid data (shown in the lower panel): here, the upper panel shows a random realization of the model (6.27) generated by a Simulink simulation model using a look-up table for the nonlinearity based on the SDP estimation results in Figure 6.16 (cf. Simulation Exam-




FIGURE 6.18. Simulated random realization of SDAR model (top); original squid data (bottom).

ple 2). It is clear that the general nature of the response visually matches the actual squid signal rather well and this is confirmed by the similarity in the statistical properties of the two signals. As in the rainfall-flow example, the estimated non-parametric nonlinearity can be parameterized in various ways. Here, a nine function RBF model, designed using the Matlab Neural Network Toolbox function *solverb*, can reproduce the nonlinearity with an $R^2 = 0.997$, and the Simulink-simulated behavior of the resulting parameterized model closely resembles that of both the non-parametric model and the squid data.

The dots shown in Figure 6.16 represent the phase-plane or embedding graph of the squid data, with y_t plotted versus y_{t-1} . This is interesting because it draws a comparison between the present SDP approach to modeling nonlinear systems and existing methods used by nonlinear systems theorists, where smooth curves are often fitted to the embedding graphs in order to model the nonlinear system (see e.g. Mees, 1991, 1993). The advantage of the SDP approach is that the smoothing is optimal in a ML sense and is carried out within a stochastic, dynamic systems setting. Also, the smoothing is applied to the estimation of the SDP parameter, rather than the nonlinear function as a whole, and so it provides a more flexible and informative result. For instance, the SDP in the model (6.27) can be considered, in an approximate sense, as the changing eigenvalue of a first order, discrete-time dynamic system (we might refer to this as a *virtual-eigenvalue*) and this gives us some additional insight into the nature of the system.

By reference to Figure 6.15, for example, we see that for y_t less than

about -120, this eigenvalue is approximately constant at a value of about 0.75, suggesting that the underlying behavior over this part of the state space is quite stable, with a (virtual) time constant ≈ 3.5 sampling intervals. Between y_t values of about -120 and -105, however, the eigenvalue increases steadily to unity, where the system is at a point of neutral stability (i.e., the system is acting transiently as an integrator). Thereafter, for y_t values greater than -105, the eigenvalue rises sharply to around 1.5 and the underlying system is clearly exponentially unstable. However, if we now consider the temporal variation of the SDP in Figure 6.19, it is clear that every time that the SDP exceeds unity, the instability of the system drives it immediately back to a location where the SDP is in the region where the eigenvalue is about 0.75 and the system is stabilized. This is also well




FIGURE 6.19. SDP variation as a function of time.

illustrated by a stacked plot of the changing impulse response associated with the AR model defined by the SDAR model at each instant of time.

Second-Order SDAR Model

Although the first order SDAR model provides a good explanation of the data and produces simulated behavior that closely resembles the actual squid data, it is interesting to consider whether worthwhile improvement is obtained with a second-order model, that is

$$y_t = a_1(y_{t-1}) \cdot y_{t-1} + a_2(y_{t-2}) \cdot y_{t-2} + e_t \quad (6.28)$$

In this case, the NVR hyper-parameters optimized at:

$$NVR\{a_1(y_{t-1})\} = 5.73 \times 10^{-6}; \quad NVR\{a_2(y_{t-2})\} = 3.05 \times 10^{-6}$$




FIGURE 6.20. Estimated state dependent parameters in the 2nd order SDAR model.




FIGURE 6.21. Estimated nonlinearities in 2nd order SDAR model.

again under the assumption of RW models for the SDP parameter variations.

The resulting FIS estimates of the SDPs are shown in Figure 6.20, with the associated nonlinearities in Figure 6.21. The dots shown in Fig. 6.21 are the partial residuals, as mentioned in Section 6.2. Although this SDAR model explains 93.6% of the squid data ($R_T^2 = 0.936$) and so is a little better, in this regard, than the first-order model, the improvement is not all

that great and the back-fitting takes rather long (twenty-six iterations) to converge, suggesting that the SDPs are not all that well defined. Moreover, as we see in Figure 6.22, the contribution to the explanation of the output y_t by the second term in the equation is very small when compared with that contributed by the first term; in Figure 6.21, the second nonlinearity and partial residual are not showing very significant changes with y_{t-2} .




FIGURE 6.22. Contributions of the two nonlinear terms in the second-order SDAR model.

Finally, the interpretation of the state-dependent changes in the virtual eigenvalues in this second-order model is also very similar to that in the first-order case. Although there are now two such eigenvalues, the first behaves very similarly to that in the first-order model. The second, on the other hand, varies rather insignificantly between -0.4287 and -0.6040 , so contributing a small amount of higher-frequency, oscillatory behavior which does not radically affect the primary chaotic mechanism contributed by the first virtual eigenvalue. Once again, this is confirmed by the stacked impulse response plot mentioned earlier this changes only a little (minor oscillations are introduced) in comparison with the same plot for the first order model. To conclude, it would appear that the first-order SDAR model provides a reasonable mechanism for characterizing the behavior of the electrical activity in the axon of the squid and that, although the second-order SDAR model provides a slightly better explanation of the squid data, this does not seem sufficient to justify the increased complexity of the model.

6.4.3 SDTF Estimation

In the SDTF model, in contrast to the SDARX model, it is assumed that noise can enter as either system or measurement noise, or both. For exam-

ple, in the case of the following forced logistic equation,

$$y_t = 2.0y_{t-1} - 2.0y_{t-1}^2 + u_t + e_t \quad u_t = N(0, 0.08) \quad e_t = N(0, 0.008) \quad (6.29a)$$

it could take the form:

$$\begin{aligned} x_t &= 2.0x_{t-1} - 2.0x_{t-1}^2 + u_t + e_t, & u_t &= N(0, 0.08), & e_t &= N(0, 0.008) \\ y_t &= x_t + \xi_t, & \xi_t &= N(0, 0.08) \end{aligned} \quad (6.29b)$$

or,

$$y_t = 2.0y_{t-1} - 2.0y_{t-1}^2 + u_t + \zeta_t \quad (6.29c)$$

where the noise ζ_t is a complex nonlinear function of e_t , ξ_t and y_t . In this situation, estimates obtained under the assumption that the model is of the simpler SDARX form are nominally biased to a level dependent on the noise/signal ratio. Fortunately, however, this bias is often fairly small, even for quite high noise levels and, in consequence, it does not interfere substantially with the identification of any state dependency. For example, in the above example, the estimation results for a measurement noise level of 69% by standard deviation (48% by variance) are quite acceptable. It should be remembered also that the SDP analysis is aimed at identifying the *form* of the nonlinearity and more efficient statistical estimation follows this identification step. Consequently, provided the biased estimation is sufficient to identify the nature of the nonlinearity, it can be acceptable.

Nevertheless, it would be advantageous if a truly bias free estimation method was available in the SDTF model case and research continues on the development of an IV back-fitting algorithm which exploits the methodology discussed above and, as in the constant parameter case (see e.g. Young, 1984), is asymptotically bias-free. Such an algorithm has been developed in the slowly variable parameter (TVP) case (Young, 1999d) but has not yet been extended, in a completely successful manner, to the more difficult SDP situation considered here, although promising results have been obtained. The main problems with such an approach are two-fold. First, the IV approach is largely restricted to input-output systems since the auxiliary model that generates the instrumental variables is normally generated from the input signal(s). Second, the stability of the adaptive auxiliary model has to be maintained and this can be difficult in the nonlinear situation. In the case of nonlinear models with chaotic properties, for example, only small uncertainties can lead to wide differences in response and possible instability. Consequently, other approaches that extend the model to include concurrently estimated (SDP) noise terms are also being investigated.

6.5 Conclusions

This chapter has introduced a new approach to the statistical identification and estimation of nonlinear static and dynamic stochastic systems based

on *state-dependent parameter* (SDP) estimation. In particular, the SDPs in a nonlinear, state-dependent parameter *transfer function* model are estimated by a recursive *fixed interval smoothing* algorithm, involving special data sorting and back-fitting operations. The data have to be sorted in some defined (normally ascending magnitude) order so that the state dependent parameter variations in the sorted space are sufficiently slow to allow for statistical estimation using existing FIS algorithms developed for *time-variable parameter* estimation. The sorting operations normally relate to the state variable primarily responsible for the state dependency, so estimation is straightforward if all the SDPs in the model are dependent on the same state variable, as in the rainfall-flow example discussed in Section 6.4.1. Normally, however, this is not the case and each SDP will be functionally dependent on a different state variable (normally the variable associated with the parameter in the SDP model). In this situation, back-fitting iterations are required so that appropriate sorting can take place prior to the estimation of each SDP in turn.

Within the overall *data-based mechanistic* modeling procedure proposed by the author in previous publications, SDP estimation can be considered as the initial, non-parametric *identification* stage in the modeling process (although, even in this non-parametric form, the SDP model is often sufficient for many practical purposes). Having identified the nature of the principal nonlinearities in the system, however, they can then be parameterized in a finite parametric form; and the (normally constant) parameters of this parametric model can then be re-estimated using some more statistically efficient, form of parameter estimation, such as *maximum likelihood* optimization. Whatever final estimation approach is used, however, the resulting model should provide a parametrically efficient representation of the stochastic, nonlinear system that has considerable potential for use in various application areas, such as signal processing, time series analysis, forecasting and automatic control system design. For example, the SDP estimation methodology exploits recursive estimation in an *off-line manner*, but this sequential processing of the data facilitates the development of related *on-line adaptive* methods of signal processing, forecasting and control.

The simulated and real examples presented in the chapter, combined with those discussed in other cited references, demonstrate the efficacy of the proposed SDP approach to modeling a fairly wide and practically useful class of nonlinear stochastic systems. However, the proposed technique is new and it raises a variety of interesting theoretical questions and possibilities for extending the approach to an even richer class of nonlinear stochastic systems. For example:

- How can the approach be extended to handle multivariable state dependencies, where the SDPs may be functions of several state variables?

- What is the best method of handling the errors-in-variables problem and the estimation bias that occurs when the proposed SDP modeling approach is applied to errors-in-variables TF models? It seems likely that an *instrumental variable* method, such as that used successfully in the case of TVP models (Young, 1999d), can be devised to handle this problem in the case of well-behaved nonlinear models, but alternative approaches will be required in the case of sensitive chaotic models.
- Although no convergence problems have been encountered so far in the evaluation of the proposed SDP estimation procedure, what conditions are required for convergence of the back-fitting procedure? Hastie and Tibshirani (1996) use a similar back-fitting procedure for estimation of their *generalized additive model*. We need to establish whether their conclusions regarding convergence (which are not entirely persuasive, in any case) are applicable to the models and back-fitting procedure described in this chapter. Unlike the GAM, for instance, the nonlinear functions in the SDP models are factorized into the product of the SDP and the model variable; and the SDP is estimated by optimal FIS smoothing (rather than the more conventional scatter-plot smoothing used by Hastie and Tibshirani).
- The back-fitting procedure does not provide complete covariance information on the SDP estimates. In more general terms, therefore, what are full *theoretical* statistical properties of the SDP estimates obtained by back-fitting?
- What are the identifiability conditions on the SDP models? It is clear that problems analogous to collinearity in constant parameter model estimation can occur and that back-fitting convergence will be affected by such problems. Also, in the case of input-output models, the nature of the input signals will affect the identifiability of the model parameters. It is necessary to explore these factors further and establish what other factors may affect the identifiability of the model.

Regardless of the answers to these questions, however, the SDP approach to the identification of nonlinearities in stochastic systems appears to hold great promise. In contrast to other approaches, such as neural networks and NARMAX models, it attempts to identify the type of nonlinearity and, therefore, the form of the nonlinear model *prior* to the estimation of the parameters in the finally identified model. This helps to ensure that the final nonlinear model is efficiently parameterized (parsimonious) and it should avoid the over-parameterization that normally accompanies neural network and, to a lesser extent, NARMAX models. It also provides a non-parametric model that can be useful in its own right. As we have seen,

the SDP model can be simulated easily in programs such as Simulink, thus removing the need for the final parametric estimation in some applications, such as simulation, forecasting and automatic control. In the latter case, for instance, it is clearly possible to develop state estimation and control system design methods based on this new class of nonlinear models. Research on such developments is continuing and has so far led to encouraging initial results.

Acknowledgments

Part of this research was supported by the EPSRC under grant GR/K77884, and the BBSRC under grant EO6813. The author is also grateful to the EPSRC and the Newton Institute for Mathematical Sciences in Cambridge for support during his stay at the Newton Institute, where most of this chapter was formulated. He is also grateful to Professors Kazu Aihara and Gen Matsumoto for providing the squid data and to Professor Alistair Mees for the stimulating discussions on these data and many other topics relating to this chapter during our joint sojourn at the Newton Institute.

References

- Brown, R.L., Durbin, J. and Evans, J. M. (1975) 'Techniques for testing the constancy of regression relationships over time', *Jnl. Royal Stat. Soc., Series B*, 37, 141-192.
- Bryson, A. E. and Ho, Y. C. (1969) 'Applied Optimal Control', Blaisdell Publishing Company, Waltham (Mass).
- Daubechies, I. (1988) 'Orthonormal bases of compactly supported wavelets', *Communications on Pure Applied Mathematics*, 41, 906-966.
- Dempster, A. P., Laird, N. M. and Rubin, D. B. (1977) 'Maximum likelihood from incomplete data via the EM algorithm', *Jnl. Royal Stat. Soc., Series B*, 39,1-38.
- Fawcett, C. (1999) 'Data-based mechanistic modelling of nonlinear environmental systems', Ph.D Thesis, Centre for Research on Environmental Systems and Statistics, Lancaster University.
- Harvey, A. C. (1981) 'Time Series Models', Phillip Allen, Oxford.
- Harvey, A. C. (1984) 'A unified view of statistical forecasting procedures' (with comments), *Jnl. of Forecasting*, 3, 245-283.
- Harvey, A. C., 'Forecasting Structural Time Series Models and the Kalman Filter', Cambridge, Cambridge University Press, 1989.
- Harvey, A. C. and Peters, S. (1990) 'Estimation procedures for Structural Time Series Models', *Journal of Forecasting*, 9, 173-204.
- Hastie, T. J. and Tibshirani, R. J. (1996) 'Generalized Additive Models', Chapman and Hall, London.
- Holst, U., Hssjer, O., Bjrkklund, C., Ragnarsson, P., and Edner, H. (1996) 'Locally weighted least squares kernel regression and statistical evaluation of LIDAR measurements', *Environmetrics*, 7, 410-416.

- Jakeman, A. J. and Young, P. C. (1979) 'Recursive filtering and the inversion of ill-posed causal problems', CRES Report No. AS/R28/1979, Centre for Resource and Environmental Studies, Australian National University.
- Jakeman, A. J. and Young, P. C. (1984) 'Recursive filtering and the inversion of ill-posed causal problems', *Utilitas Mathematica*, 35, 351-376.
- Kitagawa, G. (1981) 'A non-stationary time series model and its fitting by a recursive filter', *Jnl. of Time Series Anal.*, 2, 103-116.
- Mees, A. I. (1993) 'Parsimonious dynamical reconstruction', *International Journal of Bifurcation and Chaos*, 3(3), 669-675.
- Mees, A. I. (1991) 'Dynamical systems and tessellations: detecting determinism in data', *International Journal of Bifurcation and Chaos*, 1(4), 777-794.
- Mees, A. I., Aihara, K., Adachi, M., Judd, K., Ikeguchi, T. and Matsumoto, G. (1992) 'Deterministic prediction and chaos in squid axon response', *Physics Letters A* 169, 41-45.
- Ng, C. N. and Young, P. C. (1990) 'Recursive estimation and forecasting of non-stationary time-series', *Jnl. of Forecasting*, 9, 173-204.
- Pedregal, D. J. and Young, P. C. (1996) 'Modulated cycles, a new approach to modelling seasonal/cyclical behaviour in unobserved component models', Centre for Research on Environmental Systems and Statistics (CRES), Tech. Note No. TR/145.
- Pedregal, D. J. and Young, P. C. (1998) 'Extensions of trend models in unobserved component models', Centre for Research on Environmental Systems and Statistics (CRES), Tech. Note No. TR/156 .
- Schweppe, F. (1965) 'Evaluation of likelihood function for Gaussian signals', *I.E.E.E. Transactions on Information Theory*, 11, 61-70.
- Young, P. C. (1969) 'Applying parameter estimation to dynamic systems: Part I, Theory', *Control Engineering*, 16, 10, 119-125; 'Part II Applications', 16, 11, 118-124.
- Young, P. C. (1970) 'An instrumental variable method for real-time Identification of a noisy process', *Automatica*, 6, 271-287.
- Young, P. C. (1978) 'A general theory of modeling of badly defined dynamic systems', in G.C. Vansteenkiste (ed.), *Modeling, Identification and Control in Environmental Systems*, North Holland, Amsterdam, 103-135.
- Young, P. C. (1983) 'The validity and credibility of models for badly defined systems', in M.B. Beck and G. Van Straten (eds.), *Uncertainty and Forecasting of Water Quality*, Springer Verlag, Berlin, 69-100.
- Young, P. C. (1984) 'Recursive Estimation and Time-Series Analysis', Springer-Verlag, Berlin.
- Young, P.C. (1985) 'The instrumental variable method: a practical approach to identification and system parameter estimation', in H.A.Barker and P.C.Young (eds.), *Identification and System Parameter Estimation*, Pergamon Press, Oxford, 1-16
- Young, P. C. (1988) 'Recursive extrapolation, interpolation and smoothing of non-stationary time series', in Chen, C.F. (ed.), *Identification and System Parameter Estimation*, Pergamon Press, Oxford, 33-44.
- Young, P.C. (1992) 'Parallel processes in hydrology and water quality: a unified time series approach', *Jnl. Inst. of Water and Env. Man.*, 6, 598-612,.
- Young, P. C. (1993a) 'Time variable and state dependent modelling of nonstationary and nonlinear time series', in SubbaRao, T. (ed.), *Developments in Time Series Analysis*, Chapman and Hall, London, 374-413.
- Young, P. C. (1993b) 'Concise Encyclopedia of Environmental Systems', Oxford: Pergamon Press.

- Young, P. C. (1994) 'Time-variable parameter and trend estimation in non-stationary economic time series', *Journal of Forecasting*, 13, 179-210.
- Young, P.C. (1996) 'A general approach to identification, estimation and control for a class of nonlinear dynamic systems', in M. I. Friswell and J. E. Mottershead (eds.) *Identification in Engineering Systems*, University of Wales, Swansea, 436-445.
- Young, P. C. (1998a) 'Data-based mechanistic modelling of environmental, ecological, economic and engineering systems', *Environmental Modelling and Software*, 13, 105-122.
- Young, P. C. (1998b) 'Data-based mechanistic modelling of engineering systems', *Journal of Vibration and Control*, 4, 5-28.
- Young, P. C. (1999a) 'Nonstationary time series analysis and forecasting', *Progress in Environmental Science*, 1, 3-48.
- Young, P. C. (1999b) 'Data-based mechanistic modelling, generalised sensitivity and dominant mode analysis', *Computer Physics Communications*, 115, 1-17.
- Young, P. C. (1999c) 'Data-based mechanistic modelling and validation of rainfall-flow processes', in M. G. Anderson (ed.), *Model Validation in Hydrological Science*, J. Wiley, Chichester, in press.
- Young, P. C. (1999d) 'Stochastic, dynamic modelling and signal processing: time variable and state dependent parameter estimation', in W. J. Fitzgerald, A. Walden, R. Smith and P. C. Young (eds.), *Nonstationary and Nonlinear Signal Processing*, Cambridge University Press, Cambridge, in press.
- Young, P. C. and Beven, K.J. (1994) 'Data-based mechanistic modelling and the rainfall-flow nonlinearity', *Environmetrics*, 5, 335-363.
- Young, P. C. and Lees, M. J. (1993) 'The Active Mixing Volume (AMV): a new concept in modelling environmental systems', Chapter 1 in V. Barnett and K.F. Turkman (eds.), *Statistics for the Environment*. J. Wiley, Chichester, 3-44 .
- Young, P. C. and Minchin, P. (1991) 'Environmetric time-series analysis: modelling natural systems from experimental time-series data', *Int. Jnl. Biol. Macromol.*, 13, 190-201.
- Young, P. C. and Ng, C. N. (1989) 'Variance intervention', *Journal of Forecasting*, 8, 399-416.
- Young, P. C. and Pedregal, D. J. (1996) 'Recursive fixed interval smoothing and the evaluation of LIDAR measurements', *Environmetrics*, 7, 417-427.
- Young, P. C. and Pedregal, D. J. (1997) 'Data-based mechanistic modelling', in Heij, C., Hanzon, B. and Paggman, K. (eds.), *System Dynamics in Economic and Financial Models*, J. Wiley, Chichester, 169-213.
- Young, P. C. and Pedregal, D. J. (1998) 'Recursive and en-bloc approaches to signal extraction', *Journal of Applied Statistics*, 26, 103-128.
- Young, P. C. and Pedregal, D. J. (1999) 'Macro-economic relativity: government spending, private investment and unemployment in the USA 1948-1998', *Structural Change and Economic Dynamics*, 10, 359-380.
- Young, P.C. and Runkle, D. E. (1989) 'Recursive estimation and modelling of non-stationary and nonlinear time series', *Adaptive Systems in Control and Signal Processing*, Vol. 1, IFAC/Inst. Measurement and Control, London, 49-64.
- Young, P.C., Schreider, S. Yu., and Jakeman, A.J. (1997) 'A streamflow forecasting algorithm and results for the Upper Murray Basin', *Proc. MODSIM 97 Congress on Modelling and Simulation*, Hobart, Tasmania.

- Young, P.C., Jakeman, A.J. and Post, D. (1998) 'Recent advances in the data-based modelling and analysis of hydrological systems', *Water Science and Technology*, 36, 99- 116.
- Young, P. C., Lane, K., Ng, C. N. and Parker, D. (1991) 'Recursive forecasting, smoothing and seasonal adjustment of non-stationary environmental data', *Journal of Forecasting*, 10, 57-89.
- Young, P. C., Ng, C. N. and Armitage, P.(1989) 'A systems approach to economic forecasting and seasonal adjustment', *International Journal on Computers and Mathematics with Applications*, 18, 481-501.
- Young, P. C., Pedregal, D. J. and Tych, W. (1999) 'Dynamic harmonic regression', *Journal of Forecasting*, 18, 369-394.

Chapter 7

An Introduction to Monte Carlo Methods for Bayesian Data Analysis

Christophe Andrieu
Arnaud Doucet
William J. Fitzgerald¹

ABSTRACT Often it is natural to describe a signal processing or dynamical modeling problem in terms of probability distributions, and in particular in Bayesian terms, where the unknown parameters are taken to be random variables and their distributions are updated by applying Bayes' theorem to give the distributions of the parameters conditional on the data. In the past, it was not possible to handle many non-trivial problems in this way because the distributions seldom took tractable forms. Considerable progress has been made in recent years in applying Monte Carlo methods to overcome this, and in this chapter we describe some of the new results that have made a full Bayesian approach to signal processing tractable as well as powerful.

7.1 Introduction

In many problems encountered in statistical signal processing, it is possible to describe accurately the underlying statistical model using probability distributions. A natural framework that allows one to take into account both the information given by the observations and prior information is the Bayesian framework. We will adopt here a (fully) Bayesian approach which consists of considering all unknown parameters to be random variables. This approach is now widespread in the applied statistics community but not common in many other fields related to data analysis. In this chapter, we will not enter into the debate between orthodox statistics and Bayesian statistics. We simply point out that the Bayesian framework is a unifying statistical framework; see [9] and [50] for many reasons to be Bayesian.

Apart from these considerations, the Bayesian approach suffered for a long time from severe practical limitations. That is, except for a few simple

¹ Author for correspondence.

cases, Bayesian inference cannot be performed analytically as it requires integration and/or maximization of complex multidimensional functions.

Most problems encountered in applied research (speech processing, communications, spectral analysis, target tracking etc.) require non-Gaussianity and nonlinearity to correctly account for the observed data. In such situations it is not possible to develop closed-form estimators based on the standard criteria of maximum *a posteriori* (MAP) or minimum mean square error (MMSE). One approach to solve this problem is either to make model simplifications or analytic approximations to obtain algorithms that can be implemented in closed form (this is the basis of the extended Kalman filter, for example). However, with the recent availability of high power computers, numerical-simulation based approaches can now be considered and the full complexity of real problems can be addressed.

Before the 1990s, these integration and optimization problems were often tackled using analytic approximation techniques or deterministic numerical integration or optimization methods. These classical methods are either not precise and robust enough or are too complex to implement. An alternative set of methods is simulation-based Monte Carlo methods. The basic idea of these methods is to draw a large number of samples distributed according to the posterior distributions of interest or weighted such that it is possible to estimate simulation-based consistent estimates. We underline that the basic ideas of these methods were introduced in the 1950s but that they became popular in applied statistics only at the beginning of the 1990s mainly because of the great increases in computational power. The development of these methods is at the origin of the Bayesian revolution in applied statistics [13], [55], [57], [58] and related fields including econometrics [19] and biometrics [31]. The methods are not yet well-known in many other fields related to data analysis, such as signal processing, despite their ability to allow statistical estimation to be performed for highly complex models, which may be non-Gaussian, non-linear, non-stationary or any combination of the three. They can also be used for Bayesian model selection and model mixing, which account for the model structure uncertainty inherent in any real-world problem.

This chapter attempts to provide a simple tutorial review of these Monte Carlo methods. We illustrate the potential applications of these methods by applying them to various complex nonlinear or non-Gaussian data analysis (target tracking, blind deconvolution of impulsive processes and robust spectral analysis) that appear difficult to tackle using other methods.

In Section 7.2, a basic introduction to Bayesian inference is given. In Section 7.3, the most popular Monte Carlo approaches, importance sampling (IS) and Markov chain Monte Carlo (MCMC) methods, are described. In Section 7.5, we present the application of these simulation-based methods to several data analysis problems. Various other applications of these methods are also discussed. Finally, in Section 7.6, we briefly discuss the methods themselves. Notation is summarised in an addendum.

7.2 The Bayesian Approach

7.2.1 Bayesian Model and Posterior Distribution

A Bayesian model is described by the prior distribution² $p(\theta)$ of a random parameter $\theta \in \Theta$ and by the likelihood $p(\mathbf{y}|\theta)$ of the observations \mathbf{y} .

In this framework, all information on θ based on the observations \mathbf{y} is included in the posterior distribution $p(\theta|\mathbf{y})$ which one can obtain using Bayes' theorem

$$p(\theta|\mathbf{y}) = \frac{p(\mathbf{y}|\theta)p(\theta)}{p(\mathbf{y})} \quad (7.1)$$

where the normalizing constant $p(\mathbf{y})$ is obtained by integration

$$p(\mathbf{y}) = \int_{\Theta} p(\mathbf{y}|\theta)p(\theta) d\theta. \quad (7.2)$$

The information given by the posterior distribution (7.1) might be too complicated to analyze directly when one is confronted with a decision problem. This can, however, be handled naturally in a Bayesian framework.

7.2.2 Bayesian Decision

When one wants to make a decision, the main goal is to minimize the risk of being wrong. In order to quantify the degree of error committed one introduces a cost function that penalizes the values of θ which we think are not satisfactory. This technique, well known in the orthodox statistical community, takes the following form in a Bayesian framework.

Definition 7.1. Let $L(\cdot, \cdot) : \Theta^2 \rightarrow \mathbb{R}^+$ be a cost function and $p(\theta|\mathbf{y})$ be the posterior distribution of θ . The expected posterior cost function is defined as

$$\rho(\theta_*) \triangleq \int_{\Theta} L(\theta_*, \theta)p(\theta|\mathbf{y}) d\theta = \mathbb{E}_{p(\theta|\mathbf{y})}(L(\theta_*, \theta)). \quad (7.3)$$

Definition 7.2. Given a cost function $L(\cdot, \cdot) : \Theta^2 \rightarrow \mathbb{R}^+$, the associated Bayesian estimator is defined as

$$\hat{\theta}(\mathbf{y}) \triangleq \arg \min_{\theta_* \in \Theta} \rho(\theta_*). \quad (7.4)$$

Obtaining a Bayesian estimator is thus in general a problem simultaneously

²Many readers may be tempted to stop reading at this point, arguing that they do not have any prior distribution for the unknown parameters. We argue that when this is the case, it is in most cases possible to use uninformative priors that do not affect the peaks of the likelihood; see [9] or [50] for examples. This will be illustrated in the Application section.

involving integration and optimization, which takes simpler forms in the following two very important cases:

- the quadratic cost function, $L_1(\theta_*, \theta) = (\theta_* - \theta)^T \mathbf{Q}(\theta_* - \theta)$, for any positive definite matrix \mathbf{Q} leads to the posterior mean, also known as the MMSE estimate for θ

$$\hat{\theta}_{MMSE}(\mathbf{y}) = \int_{\Theta} \theta p(\theta | \mathbf{y}) d\theta \tag{7.5}$$

which is an integration problem.

- the cost function $L_2(\theta_* - \theta; \delta) = 1 - \mathbb{I}_{\{\theta; \|\theta_* - \theta\|_2 \leq \delta\}}$ ($\theta_* - \theta$) which gives as $\delta \rightarrow 0$ the MAP estimator, that is

$$\hat{\theta}_{MAP}(\mathbf{y}) = \arg \max_{\theta \in \Theta} p(\theta | \mathbf{y}) p(\theta) \tag{7.6}$$

which is an optimization problem.

Assume here³ that $\theta = (\theta_1, \dots, \theta_{n_\theta}) \in \Theta \subset \mathbb{R}^{n_\theta}$. Then the evaluation of marginal distributions, such as

$$p(\theta_i | \mathbf{y}) = \int p(\theta | \mathbf{y}) d\theta_1 \dots d\theta_{i-1} d\theta_{i+1} \dots, d\theta_{n_\theta} \tag{7.7}$$

for $i = 1, \dots, n_\theta$, also requires integration. Similarly the evaluation of any marginal estimator (e.g., $\mathbb{E}[\theta_i | \mathbf{y}]$ or $\arg \max_{\theta_i} p(\theta_i | \mathbf{y})$ for $i = 1, \dots, n_\theta$) involves extra-integration steps over the parameters that one wants to integrate out. Other quantities can be of interest to qualify the estimator; these include the conditional covariance

$$\mathbb{E}_{p(\theta | \mathbf{y})}(\theta \theta^T) - \mathbb{E}_{p(\theta | \mathbf{y})}[\theta] \mathbb{E}_{p(\theta | \mathbf{y})}[\theta^T] \tag{7.8}$$

which again involves integrations.

7.2.3 Model Choice

Assume that we are analyzing data \mathbf{y} and we believe that the data arise from one of a set of possible models $\mathcal{M}_0, \dots, \mathcal{M}_{k_{\max}}$, where under model \mathcal{M}_i , \mathbf{y} has density $p_i(\mathbf{y} | \theta_i)$, conditional on $\theta_i \in \Theta_i$. The parameter vectors θ_i are unknown and are typically of different dimension. Let $p_i(\theta_i)$ denote the prior density for θ_i , and let p_i denote the prior probability of the model \mathcal{M}_i . For the sake of convenience, we introduce a random variable $k \in \{0, \dots, k_{\max}\}$ such that $\Pr(k = i) = \Pr(\mathcal{M}_i) = p_i$. The prior probability

³In the case of model selection typically $\Theta \subset \bigcup_i \{i\} \times \Theta_i$.

distribution for the random parameters (k, θ) is defined on a space of the form $\Theta \triangleq \bigcup_{i=0}^{k_{\max}} \{i\} \times \Theta_i$ and can be written

$$p(k, d\theta) = \sum_{i=0}^{k_{\max}} p_i(i, d\theta_i) \mathbb{I}_{\{i\} \times \Theta_i}(k, \theta), \quad (7.9)$$

where

$$p_i(i, d\theta_i) = p_i(\theta_i) d\theta_i p_i \quad (7.10)$$

(we assume that $p_i(d\theta_i)$ admits a dominating measure $d\theta_i$, usually the Lebesgue measure) and

$$\mathbb{I}_{\{i\} \times \Theta_i}(k, \theta) = \begin{cases} 1, & \text{if } (k, \theta) \in \{i\} \times \Theta_i \\ 0, & \text{otherwise} \end{cases} \quad (7.11)$$

that is, (k, θ) is in one of the spaces $\{i\} \times \Theta_i$, and the prior probability of k being equal to i and for θ being in an infinitesimal set centered around θ_i is $p_i(i, \theta_i) d\theta_i$.

After observing \mathbf{y} , one obtains the posterior distribution using Bayes' theorem

$$p(k, d\theta | \mathbf{y}) = \sum_{i=0}^{k_{\max}} p(i | \mathbf{y}) p_i(d\theta_i | \mathbf{y}) \mathbb{I}_{\{i\} \times \Theta_i}(k, \theta) \quad (7.12)$$

where $p_i(d\theta_i | \mathbf{y}) p(i | \mathbf{y})$ is the posterior probability of model \mathcal{M}_i and is given by

$$p(i | \mathbf{y}) \triangleq p(\mathcal{M}_i | \mathbf{y}) = \frac{m_i(\mathbf{y}) p_i}{\sum_{i=0}^{k_{\max}} m_i(\mathbf{y}) p_i}, \quad (7.13)$$

where

$$m_i(\mathbf{y}) \triangleq p(\mathbf{y} | i) = \int_{\Theta_i} p_i(\mathbf{y} | \theta_i) p_i(\theta_i) d\theta_i \quad (7.14)$$

is called the *marginal* distribution of \mathbf{y} under model \mathcal{M}_i . Assuming \mathcal{M}_i is the *true* model, $p(\mathbf{y} | i)$ is the density according to which \mathbf{y} will actually occur. For this reason, $m_i(\cdot)$ is also called the *predictive* density of \mathbf{y} . Under a 0-1 loss function, the optimal model is that \mathcal{M}_i which maximizes the posterior model probability $p(i | \mathbf{y})$, $i = 1, \dots, k_{\max}$.

Note that $p(i | \mathbf{y})$ can be written as

$$p(i | \mathbf{y}) = \left(1 + \sum_{j \neq i} \frac{p_j}{p_i} B_{ji} \right)^{-1}, \quad (7.15)$$

where the factor

$$B_{ji} = \frac{m_j(\mathbf{y})}{m_i(\mathbf{y})} \quad (7.16)$$

is called the *Bayes factor* of model \mathcal{M}_j against \mathcal{M}_i . Intuitively, the Bayes factor can be interpreted as the odds of \mathcal{M}_j against \mathcal{M}_i given by the

observations. Note that Bayes factors can be used to summarize the analysis independently of the model prior beliefs, p_i .

The Bayesian approach to model selection can be applied to a wide variety of problems, including multiple comparisons and the testing of non-nested hypotheses. The results are easily interpreted (as opposed to frequentist P-values) and automatically penalize over-parametrizations [10], [56]. For a detailed discussion of the advantages and applications of Bayes factors see [11] and [40].

7.2.4 Discussion

Bayesian statistics involves integration and/or optimization steps, see expressions (7.2), (7.4), (7.7) and (7.14) for example. Except in certain special cases, Bayesian inference cannot be performed analytically, and this will be illustrated on several applications in Section 7.5. The ability to integrate or maximize complex multidimensional functions is thus extremely important in Bayesian statistics. This problem has severely limited the development of the Bayesian approach in statistics and related fields. Monte Carlo methods are a set of powerful numerical methods which allow to partly solve it.

7.3 Basics of Monte Carlo Methods

With the exception of certain cases it is impossible to analytically evaluate the integrals involving $p(\theta|\mathbf{y})$ (or marginals of $p(\theta|\mathbf{y})$) or to estimate the maxima of $p(\theta|\mathbf{y})$ (or marginals of $p(\theta|\mathbf{y})$). As soon as the dimension of the space, n_θ if $\Theta \subset \mathbb{R}^{n_\theta}$, is large, classical numerical integration methods are difficult to implement and require a huge computational burden (typically the complexity increases exponentially with n_θ). Moreover, taking into account constraints of integration can be quite complex. Similarly, classical optimization methods (gradient, Newton-Raphson) need good initialization and are very sensitive to local maxima. An attractive approach to solve, at least partially, these problems consists of using Monte Carlo methods for integration and optimization.

In this section, we first show that, using a large number of random samples distributed according to $p(\theta|\mathbf{y})$, it is possible to obtain simple estimates of $\mathbb{E}_{p(\theta|\mathbf{y})}[f(\theta)]$ and of the maxima of $p(\theta|\mathbf{y})$. Thus if we are able to simulate samples according to $p(\theta|\mathbf{y})$, then we can solve the problems of integration and optimization. These simulation problems are quite complex to solve, since $p(\theta|\mathbf{y})$ is typically a complex non-standard multivariate probability distribution which is known only up to a proportionality constant. We will describe two classes of methods to solve them. The first is the importance sampling method and the other is Markov chain Monte

Carlo methods.

In Subsection 7.3.5, the IS method is presented and some results using this simple method are reviewed. In Subsection 7.3.6, the basic idea of MCMC methods is described. These methods are based on the simulation of a Markov chain whose limiting distribution is the posterior distribution of interest, $p(\theta|\mathbf{y})$. After giving a few definitions and results related to Markov chains, we will present the classical algorithms. Important practical considerations will be briefly discussed.

7.3.1 Monte Carlo Integration

Let us assume that $N \gg 1$ samples $\{\theta^{(i)}\}_{i=1, \dots, N}$ distributed according to the posterior distribution $p(\theta|\mathbf{y})$ are available. Then a Monte Carlo approximation $\hat{P}_N(d\theta|\mathbf{y})$ of this posterior distribution is given by the empirical estimate

$$\hat{P}_N(d\theta|\mathbf{y}) = \frac{1}{N} \sum_{i=1}^N \delta_{\theta^{(i)}}(d\theta). \quad (7.17)$$

That is, the concentration of the samples in a given zone of the space Θ is assumed to be representative of the probability of this zone under the distribution $p(\theta|\mathbf{y})$.

Using this approximation for $p(\theta|\mathbf{y})$, one can propose the following estimate f_N of $\mathbb{E}_{p(\theta|\mathbf{y})}[f(\theta)]$ where $f: \Theta \rightarrow \mathbb{R}$ is a $p(\theta|\mathbf{y})$ -integrable function

$$f_N = \int f(\theta) \hat{P}_N(d\theta|\mathbf{y}) = \frac{1}{N} \sum_{i=1}^N f(\theta^{(i)}). \quad (7.18)$$

This estimate is unbiased and if the samples $\{\theta^{(i)}; i = 1, \dots, N\}$ are statistically independent, then

$$\lim_{N \rightarrow +\infty} f_N \stackrel{a.s.}{\rightarrow} \mathbb{E}_{p(\theta|\mathbf{y})}(f(\theta)) \quad (7.19)$$

from the strong law of large numbers. Moreover if

$$\begin{aligned} \text{var}[f_N] &= \frac{1}{N} \left[\mathbb{E}_{p(\theta|\mathbf{y})}(f^2(\theta)) - \mathbb{E}_{p(\theta|\mathbf{y})}^2(f(\theta)) \right] \\ &\triangleq \frac{\sigma_f^2}{N} < \infty, \end{aligned} \quad (7.20)$$

then the central limit theorem yields

$$\lim_{N \rightarrow +\infty} \sqrt{N} (f_N - \mathbb{E}_{p(\theta|\mathbf{y})}(f(\theta))) \stackrel{distr.}{\Rightarrow} \mathcal{N}(0, \sigma_f^2). \quad (7.21)$$

In the case where the samples $\{\theta^{(i)}\}_{i=1, \dots, N}$ are not statistically independent, it remains possible to obtain a law of large numbers under weak

(mixing) conditions and under stronger conditions a central limit theorem [58].

A good approximation of the integrals will require a large number N of samples. However, this approach has two major advantages. Contrary to classical numerical integration methods, the empirical distribution (7.17) enables estimates of $\mathbb{E}_{p(\theta|\mathbf{y})} [f(\theta)]$ to be obtained easily for any function f . Moreover, when the samples are statistically independent, the *dimension of the integration space does not appear in the convergence rate of the estimate* towards its theoretical value. This stems from the fact that contrary to deterministic methods, which mainly rely on a regular discretization and exploration of the space Θ , Monte Carlo methods cleverly explore the space, according to the importance of the different areas. This principle of concentration on relevant areas of the posterior distribution can be applied to the optimization problem.

7.3.2 Monte Carlo Optimization for MAP Estimation

To obtain the MAP estimate

$$\theta_{MAP} \triangleq \arg \max_{\theta} p(\theta | \mathbf{y}) \quad (7.22)$$

one can adopt the following estimate

$$\hat{\theta}_{MAP} = \arg \max_{\theta^{(i)}; i=1, \dots, N} p(\mathbf{y} | \theta^{(i)}) p(\theta^{(i)}). \quad (7.23)$$

The samples $\{\theta^{(i)}\}_{i=1, \dots, N}$ being marginally distributed according to $p(\theta | \mathbf{y})$, have a high probability of being in the areas where $p(\theta | \mathbf{y})$ is large. Consequently, it can be a good strategy to estimate $\hat{\theta}_{MAP}$ based on $\{\theta^{(i)}\}_{i=1, \dots, N}$. However it would be preferable to sample directly from a distribution whose support is the set of global maxima of the distribution to maximize. This is what is done approximately by the simulated annealing (SA) method which uses the same principle, but reinforces this concentration effect at each iteration of the algorithm.

In order to obtain marginal MAP (MMAP) estimates, for example, the MMAP of the component θ_j , $1 \leq j \leq n_{\theta}$

$$\theta_{j,MMAP} \triangleq \arg \max_{\theta_j} p(\theta_j | \mathbf{y}) \quad (7.24)$$

we can adopt

$$\hat{\theta}_{j,MMAP} = \arg \max_{\theta_j^{(i)}; i=1, \dots, N} p(\mathbf{y} | \theta_j^{(i)}) p(\theta_j^{(i)}). \quad (7.25)$$

Practically, these estimates cannot always be computed as it is necessary to be able to evaluate, up to a normalizing constant, the functions to maximize. In cases where it is not possible, one can use a graphical method

by monitoring the Monte Carlo approximation of the distribution to maximize, i.e., the histogram of the simulated samples. However, it involves discretizing Θ and this is inefficient for high-dimensional parameter space. A Monte Carlo method related to SA, beyond the scope of this tutorial, has been proposed to solve this problem [53].

7.3.3 Simulation Problems

We have shown that based on a large set of samples $\{\theta^{(i)}\}_{i=1,\dots,N}$ distributed according to the posterior distribution $p(\theta|\mathbf{y})$, one can usually easily estimate the distribution $p(\theta|\mathbf{y})$, its maxima and integrals with respect to it. Until now, we have assumed that this set of samples is available. In practice, it is necessary to simulate these samples. Simulating samples from nonstandard probability distributions known only up to a proportionality factor is a complex problem but there are ways to solve it. We first describe some classical and universal methods that can be used in simple cases and then describe the importance sampling method and Markov chain Monte Carlo methods.

In what follows we present the general case where $\pi(\theta)$ is a probability distribution known up to a normalizing constant. In our applications, $\pi(\theta)$ is a posterior distribution, i.e., $\pi(\theta) = p(\theta|\mathbf{y})$, but of course these methods can also be applied in a non-Bayesian framework, where $\pi(\theta)$ is any distribution from which we wish to sample.

7.3.4 Classical Methods

Procedures for simulation according to standard distributions (e.g., Gaussian, gamma, student) are given in [21] and [49]. Most classical methods are based on the inversion of the cumulative distribution function or on mixtures and compositions of basic distributions.

To sample according to a distribution which is known only up to a proportionality constant, the most well-known method is the accept/reject procedure detailed later. Let us assume that $\pi(\theta) \leq Mq(\theta)$ ($M < +\infty$) for any $\theta \in \Theta$, where $q(\theta)$ is any probability distribution termed the candidate distribution selected so that it is possible and easy to sample from. Then the accept/reject procedure works as follows:

Accept/Reject procedure

1. Sample $\theta \sim q(\cdot)$ and $u \sim \mathcal{U}_{(0,1)}$.
 2. If $u < \frac{\pi(\theta)}{Mq(\theta)}$ then return θ ; otherwise return to step 1.
-

The returned value θ is distributed according to $\pi(\theta)$. Indeed for $\tilde{\theta} \in \Theta$ (we present the scalar case for the sake of simplicity, the extension is straightforward)

$$\begin{aligned} \Pr(\theta \leq \tilde{\theta} \text{ and } \theta \text{ accepted}) &= \int_{-\infty}^{\tilde{\theta}} \frac{\pi(\theta)}{Mq(\theta)} q(\theta) d\theta \quad (7.26) \\ &= \frac{\int_{-\infty}^{\tilde{\theta}} \pi(\theta) d\theta}{M} \end{aligned}$$

and the probability for the “candidate” to be accepted is equal to

$$\begin{aligned} \Pr(\theta \text{ accepted}) &= \Pr\left(u < \frac{\pi(\theta)}{Mq(\theta)}\right) = \int_{\Theta} \frac{\pi(\theta)}{Mq(\theta)} q(\theta) d\theta \quad (7.27) \\ &= \frac{1}{M} \end{aligned}$$

so

$$\Pr(\theta \leq \tilde{\theta} | \theta \text{ accepted}) = \int_{-\infty}^{\tilde{\theta}} \pi(\theta) d\theta, \quad (7.28)$$

i.e., θ is exactly drawn from $\pi(\theta)$. This method suffers from several severe limitations. It is necessary to be able to bound $\pi(\theta)/q(\theta)$ from above by a constant M over the whole space Θ . This is not always possible or the constant M we obtain is such that the acceptance probability $\frac{1}{M} \ll 1$ and the algorithm is impractical. Moreover in a Bayesian framework, one can often not find this acceptance probability, since $\pi(\theta)$ is known only up to a constant of proportionality.

Example 7.1. Let us assume that one wants to simulate samples from $\pi(\theta) \triangleq p(\theta|y) \propto p(y|\theta)p(\theta)$. We assume that $p(y|\theta)$ is known analytically and $p(y|\theta) \leq C$ for any θ , where C is known. We also assume that we can simulate from $p(\theta)$. Thus one can choose $q(\theta) = p(\theta)$ and use the accept/reject procedure to sample from $p(\theta|y)$. Indeed

$$\frac{p(\theta|y)}{p(\theta)} = \frac{p(y|\theta)}{p(y)} \leq \frac{C}{p(y)} = M \quad (7.29)$$

is bounded and

$$\frac{\pi(\theta)}{Mq(\theta)} = \frac{p(\theta|y)}{\frac{C}{p(y)}p(\theta)} = \frac{p(y|\theta)}{C} \quad (7.30)$$

can be evaluated analytically. However, the acceptance rate $1/M$ is usually unknown as it involves $p(y)$ which is itself usually unknown.

7.3.5 Importance Sampling

Classical Case

The importance sampling method is based on the following simple remark. Let us consider any probability density $q(\theta)$ such that the support of $q(\theta)$

includes that of $\pi(\theta)$. $q(\theta)$ is called the importance function. It is chosen such that it is easy to obtain $N \gg 1$ statistically independent samples $\{\theta^{(i)}\}_{i=1, \dots, N}$ distributed according to $q(\theta)$. Then one can propose the following Monte Carlo approximation $\hat{\pi}_N(d\theta)$ of this distribution

$$\hat{\pi}_N(d\theta) = \frac{1}{N} \sum_{i=1}^N \bar{w}(\theta^{(i)}) \delta_{\theta^{(i)}}(d\theta) \quad (7.31)$$

if the so-called importance weights

$$\bar{w}(\theta^{(i)}) = \frac{\pi(\theta^{(i)})}{q(\theta^{(i)})} \quad (7.32)$$

can be evaluated analytically.

Using this approximation of $\pi(\theta)$, one can propose the following estimate f_N of $\mathbb{E}_{\pi(\theta)}[f(\theta)]$ where $f: \Theta \rightarrow \mathbb{R}$

$$\begin{aligned} f_N &= \int f(\theta) \hat{\pi}_N(d\theta) \\ &= \frac{1}{N} \sum_{i=1}^N \bar{w}(\theta^{(i)}) f(\theta^{(i)}). \end{aligned} \quad (7.33)$$

Clearly this estimate is unbiased and by the strong law of large numbers

$$\lim_{N \rightarrow +\infty} f_N \xrightarrow{\text{a.s.}} \mathbb{E}_{\pi(\theta)}(f(\theta)). \quad (7.34)$$

Moreover if

$$\begin{aligned} \text{var}_{q(\theta)}[f_N] &= \frac{1}{N} \left[\mathbb{E}_{q(\theta)}(\bar{w}^2(\theta) f^2(\theta)) - \mathbb{E}_{\pi(\theta)}^2(f(\theta)) \right] \\ &= \frac{\sigma_f^2}{N} < +\infty, \end{aligned}$$

then the central limit theorem yields

$$\lim_{N \rightarrow +\infty} \sqrt{N} (f_N - \mathbb{E}_{\pi(\theta)}(f(\theta))) \xrightarrow{\text{dist.}} \mathcal{N}(0, \sigma_f^2). \quad (7.35)$$

In IS, the simulated samples $\{\theta^{(i)}; i = 1, \dots, N\}$ are independent but they must be weighted to produce a simulation-consistent approximation of $\mathbb{E}_{\pi(\theta)}(f(\theta))$.

Remark 7.1. Note that for a given function $f(\theta)$, this method can be super-efficient in the sense that it is possible to find a distribution $q(\theta)$ that yields an estimate with a lower variance than using a perfect Monte Carlo method, i.e., $q(\theta) = \pi(\theta)$. This method is often used to evaluate the probability of rare events

for communication networks. In our applications, the aim is usually different in the sense that we want to have a good approximation of $\pi(\theta)$ and not of a particular integral with respect to $\pi(\theta)$, so we seek to have $q(\theta) \simeq \pi(\theta)$.

This method requires the ability to evaluate the importance weights $w(\theta^{(i)})$ for $i = 1, \dots, N$. In a Bayesian framework where $\pi(\theta) = p(\theta|\mathbf{y})$, we have

$$w(\theta) = \frac{p(\theta|\mathbf{y})}{q(\theta)} = \frac{p(\mathbf{y}|\theta)p(\theta)}{q(\theta)p(\mathbf{y})}. \quad (7.36)$$

It is usually **impossible** to evaluate $w(\theta)$ as the normalizing constant $p(\mathbf{y})$ is unknown. It is necessary to use another method.

Bayesian Importance Sampling

Using $N \gg 1$ statistically independent samples $\{\theta^{(i)}\}_{i=1, \dots, N}$ distributed according to $q(\theta)$, Bayesian importance sampling constructs the following Monte Carlo approximation of $\hat{\pi}_N(d\theta)$

$$\hat{\pi}_N(d\theta) = \frac{\sum_{i=1}^N w(\theta^{(i)}) \delta_{\theta^{(i)}}(d\theta)}{\sum_{j=1}^N w(\theta^{(j)})}, \quad (7.37)$$

where the importance weights are evaluated only up to a normalizing constant

$$w(\theta^{(i)}) \propto \frac{\pi(\theta^{(i)})}{q(\theta^{(i)})}. \quad (7.38)$$

This general method is called Bayesian IS as it is mostly used in applications in a Bayesian framework.

Using this approximation of $\pi(\theta)$, one can propose the following estimate f_N of $\mathbb{E}_{\pi(\cdot)}[f(\theta)]$ where $f: \Theta \rightarrow \mathbb{R}$

$$f_N = \int f(\theta) \hat{\pi}_N(d\theta) = \frac{\sum_{i=1}^N w(\theta^{(i)}) f(\theta^{(i)})}{\sum_{i=1}^N w(\theta^{(i)})}. \quad (7.39)$$

When N is finite, this estimate is biased. However, it is asymptotically unbiased from the strong law of large numbers:

$$\lim_{N \rightarrow +\infty} f_N \xrightarrow{\text{a.s.}} \mathbb{E}_{\pi(\theta)}(f(\theta)). \quad (7.40)$$

If $w(\theta)$ is bounded above on Θ and $\text{var}_{\pi(\theta)}(f(\theta)) < +\infty$, then it is easily shown that a central limit theorem is valid, the asymptotic variance being given by the delta method, see [33], for example.

Remark 7.2. In a Bayesian framework where $\pi(\theta) = p(\theta|\mathbf{y})$, one can select, for example, $q(\theta) = p(\theta)$. Then we can compute the unnormalized importance weights

$$w(\theta) \propto \frac{p(\theta|\mathbf{y})}{p(\theta)} \propto p(\mathbf{y}|\theta),$$

and so evaluate the estimates given by this IS method. Of course, $q(\theta)$ can depend on \mathbf{y} . For example, $q(\theta)$ could be a multivariate distribution of mean \mathbf{y} and fixed covariance.

This non-iterative method is typically very easy to implement and one can easily derive a version which is well adapted to online processing; see Subsection 7.5.1 and [24], [27] for more details. Unfortunately it is also often ineffective, i.e., the practical rate of convergence can be extremely slow if the importance function is not well chosen. In practice, it is usually never used if the dimension of the state space is high, say $n_\theta > 10$.

Sampling Importance Resampling

We conclude this part by presenting the sampling importance resampling (SIR) procedure introduced by Rubin [54], [57]. The IS or Bayesian IS methods yield the following weighted approximation of $\pi(d\theta)$

$$\hat{\pi}_N(d\theta) = \sum_{i=1}^N \alpha_i \delta_{\theta^{(i)}}(d\theta), \quad (7.41)$$

where $\alpha_i = w(\theta^{(i)})$ in the case of IS and $\alpha_i = \left[\sum_{j=1}^N w(\theta^{(j)}) \right]^{-1} w(\theta^{(i)})$ in the case of Bayesian IS. If one is interested in obtaining say $M(N)$ I.I.D. samples from $\hat{\pi}_N(d\theta)$, then an asymptotically ($N/M(N) \rightarrow +\infty$) valid method consists of resampling $M(N)$ times according to the discrete distribution $\hat{\pi}_N(d\theta)$. After this resampling step, we obtain $M(N)$ samples $\tilde{\theta}^{(i)}$, $i = 1, \dots, M(N)$ (with the possibility that $\tilde{\theta}^{(i)} = \tilde{\theta}^{(j)}$ for $i \neq j$), such that an alternative expression of the posterior distribution is

$$\tilde{\pi}_{M(N)}(d\theta) = \frac{1}{M(N)} \sum_{i=1}^{M(N)} \delta_{\tilde{\theta}^{(i)}}(d\theta). \quad (7.42)$$

The resampling scheme introduces some additional Monte Carlo variations with respect to (7.41), and it is then not clear to see if the SIR procedure is of any practical interest. It appears actually as a key step when one is interested in estimating a sequence of posterior distributions evolving over time using sequential Monte Carlo methods, see Subsection 7.5.1 and [35], [24], [27] for further details.

7.3.6 Markov Chain Monte Carlo Methods

The previous IS-based methods we have presented are typically inefficient when the dimension of the parameter space is large. In this section, we present a class of iterative simulation-based methods: Markov chain Monte Carlo methods. MCMC methods are a set of procedures that allow successful solution of simulation problems for much more complex models and

are at the origin of the Bayesian revolution in applied statistics [13], [29], [12], [55], [52].

The basic idea of MCMC methods is to simulate an ergodic Markov chain whose samples are asymptotically distributed according to $\pi(d\theta)$. In this section, we recall a few results in general, i.e., continuous or discrete, state-space Markov chains that are sufficient to understand the main convergence results. For a much more complete and rigorous treatment of this theory, one can consult the excellent monograph by Meyn and Tweedie [46] and for articles or books devoted to Markov chains theory applied to MCMC see [34, 52] and [58]. The most classical MCMC methods are presented in 7.3.6, as well as non-homogeneous MCMC methods for optimization, i.e., simulated annealing.

A few results on general state-space Markov chains follow.

First, let us introduce a few definitions concerning Markov chains.

Definition 7.3. *A Markov chain is a sequence of random variables $\{\theta_i; i \in \mathbb{N}\}$ defined in the same space (Θ, \mathcal{E}) which satisfies the following property, for any $A \in \mathcal{E}$*

$$\Pr(\theta_i \in A | \theta_{i-1}, \dots, \theta_0) = \Pr(\theta_i \in A | \theta_{i-1}). \quad (7.43)$$

In many applications, we have $\Theta = \mathbb{R}^{n_\theta}$ and $\mathcal{E} = \mathcal{B}(\mathbb{R}^{n_\theta})$ (Borel sets of \mathbb{R}^{n_θ}). We will consider here only time-homogeneous Markov chains, i.e., the transition kernel of the Markov chain is fixed over the time. One can define for any $i \in \mathbb{N}$ and $(\theta, A) \in \Theta \times \mathcal{E}$:

$$P(\theta, A) \triangleq \Pr(\theta_{i+1} \in A | \theta_i = \theta). \quad (7.44)$$

$P(\theta, A)$ is called the transition kernel of the Markov chain. We have $P(\theta, A) = \int_A P(\theta, d\theta')$ where $P(\theta, d\theta')$ is the probability to going to a “small” set $d\theta' \in \mathcal{E}$, starting from θ .

A homogeneous Markov chain is defined by the probability distribution of its initial state $\mu(d\theta_0)$ (which can be deterministic, so that it is a delta-Dirac measure) and its transition kernel. We obtain for the probability of any set $A_0 \times A_1 \times \dots \times A_n$:

$$P((\theta_0, \theta_1, \dots, \theta_n) \in A_0 \times \dots \times A_n) = \int_{A_0} \dots \int_{A_n} \prod_{i=1}^n P(\theta_{i-1}, d\theta_i) \mu(d\theta_0). \quad (7.45)$$

By marginalizing, we obtain the distribution of θ_i given the initial condition. The distribution $P^n(\theta, A) \triangleq \Pr(\theta_n \in A | \theta_0 = \theta)$ is also defined by

the following recursion:

$$P^1(\theta, A) = P(\theta, A) \quad (7.46)$$

$$P^n(\theta, A) = \int_{\Theta} P^{n-1}(\theta, d\theta^*) P(\theta^*, A) \text{ for } n > 1. \quad (7.47)$$

We now give some other definitions.

Definition 7.4. *Invariance.* Let $\pi(d\theta)$ be a probability distribution. We say that $\pi(d\theta)$ is an invariant or stationary distribution for the transition kernel P if for any $A \in \mathcal{E}$

$$\pi(A) = \int_{\Theta} \pi(d\theta) P(\theta, A) = \int_{\Theta} \pi(d\theta) \int_A P(\theta, d\theta^*). \quad (7.48)$$

This notion is very important. It implies that if a state of the Markov chain θ_i is distributed according to $\pi(d\theta)$ then θ_{i+1} and all the following states are distributed marginally according to $\pi(d\theta)$.

Remark 7.3. *All MCMC algorithms are constructed to satisfy this fundamental property and thus it is not necessary to verify it in practice.*

To ensure that π is the invariant distribution of a Markov chain, a sufficient condition is the so-called reversibility condition. This notion is very often used in the framework of MCMC algorithms.

Definition 7.5. *Reversibility.* A transition kernel P is π -reversible if it satisfies for any $(A, B) \in \mathcal{E} \times \mathcal{E}$:

$$\begin{aligned} \int_A \pi(d\theta) P(\theta, B) &= \int_B \pi(d\theta) P(\theta, A) \\ \int_A \pi(d\theta) \int_B P(\theta, d\theta^*) &= \int_B \pi(d\theta) \int_A P(\theta, d\theta^*). \end{aligned} \quad (7.49)$$

This means that when θ is distributed according to $\pi(d\theta)$ then the probability of going to B from A is equal to the probability of going to A starting from B . This notion is important as it implies invariance of the distribution $\pi(d\theta)$ for the transition kernel P .

Let us assume that we have a transition kernel P with invariant distribution π . Then we sample a realization of the Markov chain $\{\theta_0, \theta_1, \dots\}$. One can ask the following questions:

- In practice, the initial state is usually not distributed according to π and thus, if the transition kernel is not trivial, i.e., $P(\theta, d\theta') \neq \pi(d\theta')$, the following states of the Markov chain are not either. Under which conditions does the simulated Markov chain converge asymptotically towards its invariant distribution π ?

- The simulated samples are not statistically independent. Under which conditions is it possible to obtain estimates $\frac{1}{N} \sum_{i=1}^N f(\theta_i)$ which converge asymptotically towards $\int_{\Theta} f(\theta) \pi(d\theta)$?

We will see that the conditions required on the Markov transition kernel to ensure asymptotic convergence are very weak: irreducibility and aperiodicity.

Definition 7.6. *Irreducibility.* Let φ be a probability measure on (Θ, \mathcal{E}) . A Markov chain is φ -irreducible if for any $\theta \in \Theta$ and for any $A \in \mathcal{E}$

$$\varphi(A) > 0 \Rightarrow \exists n \in \mathbb{N}^* \text{ such that } P^n(\theta, A) > 0. \tag{7.50}$$

This means that all sets of non-null φ measure can be reached with a non-null probability in a finite number of iterations. A sufficient condition to verify the irreducibility of P with respect to φ is that there exists $n \in \mathbb{N}^*$ such that P^n has a strictly positive density $p^n(\theta, \theta^*) > 0$ with respect to φ , i.e., $P^n(\theta, A) = \int_A p^n(\theta, d\theta^*) \varphi(d\theta^*)$ [34, p. 63].

Definition 7.7. *Aperiodicity.* A Markov chain is aperiodic if there does not exist a partition of $\Theta = (A_1, \dots, A_d)$ for $d \geq 2$ such that $P(\theta, A_{i+1 \bmod d}) = 1$ for any $\theta \in A_i$.

The aperiodicity condition eliminates kernels which induce a periodic behavior in the trajectories of the Markov chain. If the transition kernel $P(\theta, d\theta^*)$ is such that the probability of staying at θ is non-null or if this kernel has a positive density in the neighborhood of θ , then the chain is aperiodic [34, p. 63], [46], [52].

Before stating the main result of convergence, it is necessary to introduce a distance between two probability distributions μ_1 and μ_2 on (Θ, \mathcal{E}) .

Definition 7.8. *The total variation norm is defined by*

$$\|\mu_1(\cdot) - \mu_2(\cdot)\|_{TV} \triangleq \sup_{A \in \mathcal{E}} |\mu_1(A) - \mu_2(A)|. \tag{7.51}$$

If $\mu_1(d\theta) = \mu_1(\theta) d\theta$ and $\mu_2(d\theta) = \mu_2(\theta) d\theta$, then

$$\|\mu_1 - \mu_2\|_{TV} = \frac{1}{2} \int_{\Theta} |\mu_1(\theta) - \mu_2(\theta)| d\theta. \tag{7.52}$$

We define $\mathbb{E}_{\pi(\theta)}[f] \triangleq \int_{\Theta} f(\theta) \pi(\theta) d\theta$. The estimate of this quantity obtained by averaging the N first simulated values of the Markov chain will be denoted f_N :

$$f_N \triangleq \frac{1}{N} \sum_{i=0}^{N-1} f(\theta^{(i)}). \tag{7.53}$$

We are now able to state the main convergence result.

Theorem 7.1. [58, theorem 1*, p. 1758] Let $\{\theta_i; i \in \mathbb{N}\}$ be a Markov chain with transition kernel P and invariant distribution π . Assume that the Markov chain is φ -irreducible, where φ is a probability distribution on (Θ, \mathcal{E}) . Let $f : \Theta \rightarrow \mathbb{R}$ be such that $\mathbb{E}_\pi [|f|] < +\infty$; then

$$P_{\theta_0} \left[f_N \xrightarrow{N \rightarrow +\infty} \mathbb{E}_\pi [f] \right] = 1 \quad (7.54)$$

for π -almost all initialization points θ_0 . Moreover, if P is aperiodic

$$\lim_{N \rightarrow +\infty} \left\| P^N(\theta_0, \cdot) - \pi(\cdot) \right\|_{TV} = 0 \quad (7.55)$$

for π -almost all θ_0 .

Remark 7.4. Typically, we set $\varphi = \pi$ and verify the π -irreducibility of the kernel.

Remark 7.5. The aperiodicity condition is not necessary if one is only interested in estimates of $\mathbb{E}_\pi [f]$.

The conditions on the transition kernel which lead to this theorem are weak and intuitive; the Markov chain must be able to explore the support of the target distribution and this exploration must not be periodic. Then the estimates obtained by averaging the simulated samples are asymptotically consistent and these samples are asymptotically distributed according to $\pi(d\theta)$. However, this theorem allows some starting points (of null π measure) for which convergence is not ensured. This theorem is valid for all starting points θ_0 if the Markov chain is also Harris recurrent [46, 58, 52].

This theorem validates theoretically under weak assumptions the basic principle of MCMC algorithms, that is, the idea of constructing a Markov chain of invariant distribution $\pi(d\theta)$ when simulating directly from $\pi(d\theta)$ is impossible. In fact, one must be careful because this theorem does not give any information on the rate of convergence of the Markov chain towards its invariant distribution. Practically, this convergence can be so slow that the approach is of no practical interest. It is possible to obtain under additional assumptions on the transition kernels better convergence results that ensure a geometric or even a uniform geometric convergence rate of the Markov chain [46, 58, 52]; we do not detail these conditions here.

MCMC Algorithms

In this section, we first present some classical methods for constructing a Markov chain that admits as invariant distribution $\pi(d\theta) = \pi(\theta) d\theta$, when $\theta \triangleq (\theta_1, \dots, \theta_{n_\theta}) \in \Theta \subset \mathbb{R}^{n_\theta}$. The more complex case, where θ belongs to a union of subspaces of various dimensions, is discussed in Subsection 7.3.6.

The Gibbs Sampler

Given θ , the Gibbs sampler consists of first defining a partition of the components of θ , say $\theta_1, \dots, \theta_p$ ($p \leq n_\theta$) and to sample successively from the so-called full conditional distributions $\pi(\theta_k | \theta_{-k})$, where $\theta_{-k} \triangleq (\theta_1, \dots, \theta_{k-1}, \theta_{k+1}, \dots, \theta_p)$. The algorithm proceeds as follows.

Gibbs Sampling

1. Initialization, $i = 0$. Set randomly or deterministically $\theta^{(0)} = \theta_0$.
2. Iteration i , $i \geq 1$.
 - Sample $\theta_1^{(i)} \sim \pi(\theta_1 | \theta_{-1}^{(i)})$.
 - \vdots
 - Sample $\theta_k^{(i)} \sim \pi(\theta_k | \theta_{-k}^{(i)})$.
 - \vdots
 - Sample $\theta_p^{(i)} \sim \pi(\theta_p | \theta_{-p}^{(i)})$,

where $\theta_{-k}^{(i)} \triangleq (\theta_1^{(i)}, \dots, \theta_{k-1}^{(i)}, \theta_{k+1}^{(i-1)}, \dots, \theta_p^{(i-1)})$.

This algorithm defines the following transition kernel:

$$K(\theta^{(i-1)}, \theta^{(i)}) = \prod_{k=1}^p \pi(\theta_k^{(i)} | \theta_{-k}^{(i)}). \quad (7.56)$$

It is straightforward to verify that this kernel admits $\pi(\theta)$ as invariant density. The algorithm is very general. In particular, there are many possible ways to partition the vector θ . In practice, it is recommended whenever possible to sample highly correlated variables together. This can improve dramatically the rate of convergence of the sampler, see [15] for an application to time series.

Remark 7.6. *The version of the Gibbs sampler we have presented is the so-called deterministic scan Gibbs sampler in which the order that the components are updated is fixed. One can also update these components in a random order.*

Remark 7.7. *When $p = 2$, one obtains a special case of the so-called data augmentation algorithm introduced by Tanner and Wong in 1987 in the statistical literature. Data augmentation is often presented as a stochastic Bayesian alternative to the EM algorithm [57].*

Remark 7.8. An (obvious) necessary condition to generate an ergodic Markov chain is that each component θ_j has to appear in a subset θ_l . However, one does not actually need $(\theta_1, \dots, \theta_p)$ to define a partition of θ , and a parameter θ_j can appear in two subsets θ_l and θ_m with $l \neq m$. This remark is of interest to build some so-called partial samplers [12], [16]; see Subsection 7.5.2 for an application.

Metropolis-Hastings Algorithm

Another very popular MCMC algorithm is the Metropolis-Hastings (MH) algorithm [44], [38] which uses a candidate distribution $q(\theta, \theta')$ to sample from $\pi(\theta')$. This algorithm proceeds as follows.

Metropolis-Hastings Algorithm

1. Initialization, $i = 0$. Set randomly or deterministically $\theta^{(0)} = \theta_0$.
2. Iteration i , $i \geq 1$.

- Sample a candidate $\theta \sim q(\theta^{(i-1)}, \cdot)$.
- Evaluate the acceptance probability

$$\alpha(\theta^{(i-1)}, \theta) = \min \left\{ \frac{\pi(\theta) / q(\theta^{(i-1)}, \theta)}{\pi(\theta^{(i-1)}) / q(\theta, \theta^{(i-1)})}, 1 \right\} \quad (7.57)$$

- Sample $u \sim \mathcal{U}_{[0,1]}$, if $u \leq \alpha(\theta^{(i-1)}, \theta)$ then $\theta^{(i)} = \theta$ otherwise $\theta^{(i)} = \theta^{(i-1)}$.

This algorithm defines the following transition kernel:

$$K_{MH}(\theta, d\theta') = \alpha(\theta, \theta') q(\theta, \theta') d\theta' + \delta_\theta(d\theta') \int [1 - \alpha(\theta, \mathbf{u})] q(\theta, \mathbf{u}) d\mathbf{u}. \quad (7.58)$$

To check that it has $\pi(d\theta)$ as its invariant distribution, one can easily verify that $K_{MH}(\cdot, \cdot)$ is π -reversible.

Remark 7.9. Note that this algorithm does not require knowledge of the normalizing constant of $\pi(d\theta)$ as only the ratio $\pi(\theta) / \pi(\theta^{(i-1)})$ appears in the acceptance probability.

A simple condition which ensures the irreducibility and the aperiodicity of the MH algorithm is that $q(\theta, \theta')$ is continuous and strictly positive on the support of $\pi(\theta')$ for any θ .

Example 7.2. Let us assume that we want to simulate a set of samples from $p(\theta|\mathbf{y})$. Using Bayes' theorem we have $p(\theta|\mathbf{y}) \propto p(\mathbf{y}|\theta)p(\theta)$. A MH procedure consists of simulating some candidates θ' according to $q(\theta, \theta')$, evaluating some quantities $\alpha(\theta, \theta') = \min \left\{ 1, \frac{p(\mathbf{y}|\theta')p(\theta')q(\theta, \theta')}{p(\mathbf{y}|\theta)p(\theta)q(\theta, \theta')} \right\}$, and accepting these candidates with probability $\alpha(\theta, \theta')$.

For the candidate distribution $q(\cdot, \cdot)$, there is an infinity of choices.

Random Walk. A simple choice consists of proposing as candidate a perturbation of the current state, i.e., $\theta' = \theta + \mathbf{z}$ where \mathbf{z} is a random increment of density $\varphi(\mathbf{z})$.

- This algorithm corresponds to the particular case $q(\theta, \theta') = \varphi(\theta' - \theta)$. We obtain the following acceptance probability:

$$\alpha(\theta, \theta') = \min \left\{ \frac{\pi(\theta') \varphi(\theta - \theta')}{\pi(\theta) \varphi(\theta' - \theta)}, 1 \right\}. \quad (7.59)$$

- If $q(\theta, \theta') = \varphi(\theta - \theta') = \varphi(\theta' - \theta)$ then we obtain

$$\alpha(\theta, \theta') = \min \left\{ \frac{\pi(\theta')}{\pi(\theta)}, 1 \right\}. \quad (7.60)$$

This algorithm is called the Metropolis algorithm [44].

Independent Metropolis-Hastings. In this case, we select the candidate independently of the current state according to a distribution $\varphi(\theta')$. Thus $q(\theta, \theta') = \varphi(\theta')$ and we obtain the following acceptance probability:

$$\alpha(\theta, \theta') = \min \left\{ \frac{\pi(\theta') \varphi(\theta)}{\pi(\theta) \varphi(\theta')}, 1 \right\}. \quad (7.61)$$

In the case where $\pi(\theta)/\varphi(\theta)$ is bounded, i.e., we could also apply the accept/reject procedure, this procedure shows (fortunately) better asymptotic performance in terms of variance of ergodic averages [43].

Example 7.3. In a Bayesian framework, if we want to sample from $p(\theta|\mathbf{y}) \propto p(\mathbf{y}|\theta)p(\theta)$ then one can take $p(\theta)$ as candidate distribution. Then the acceptance reduces to

$$\alpha(\theta, \theta') = \min \left\{ \frac{p(\mathbf{y}|\theta')}{p(\mathbf{y}|\theta)}, 1 \right\}. \quad (7.62)$$

One can also couple the MH algorithm with an accept/reject procedure or proposing a candidate distribution which is, for example, a discretization of the target distribution $\pi(\theta)$ [58].

Metropolis-Hastings One-at-a-Time

In the case where θ is high-dimensional, it becomes very difficult to select a good proposal distribution: either the acceptance probability is very low and thus the chain does not mix, or the chain explores only one mode of the distribution. To solve this problem one can define a partition of θ , similarly to the Gibbs sampling algorithm. Then each component θ_k is updated according to a MH step of proposal distributions admitting as invariant distribution the full conditional distribution $\pi(\theta_k | \theta_{-k})$, using a set of proposal distribution $q_k(\cdot, \cdot)$, $k = 1, \dots, p$.

MH One-at-a-Time

1. Initialization, $i = 0$. Set randomly or deterministically $\theta^{(0)} = \theta_0$.
2. Iteration i , $i \geq 1$.
 - – For $k = 1$ to p
 - 1. • – Sample $\theta_k^{(i)}$ according to a MH step with proposal distribution

$$q_k \left(\left(\theta_{-k}^{(i)}, \theta_k^{(i-1)} \right), \theta_k \right) \tag{7.63}$$

and invariant distribution $\pi \left(\theta_k | \theta_{-k}^{(i)} \right)$.

End For.

where $\theta_{-k}^{(i)} \triangleq \left(\theta_1^{(i)}, \dots, \theta_{k-1}^{(i)}, \theta_{k+1}^{(i-1)}, \dots, \theta_p^{(i-1)} \right)$.

This algorithm includes the Gibbs sampler as a special case. Indeed, this corresponds to the particular case where the proposal distributions of the MH steps are equal to the full conditional distributions, i.e.,

$$q_k \left(\left(\theta_{-k}^{(i)}, \theta_k^{(i-1)} \right), \theta_k \right) = \pi \left(\theta_k | \theta_{-k}^{(i)} \right),$$

so that the acceptance probabilities are equal to 1 and no candidate is rejected.

Reversible Jump MCMC

This section describes how to build MCMC algorithms for model selection, i.e., how to construct ergodic Markov chains admitting $p_k(k, d\theta_k | \mathbf{y})$, defined in equation (7.12), as their invariant distribution. In this case, $\Theta \triangleq \bigcup_{i=0}^{k_{\max}} \{i\} \times \Theta_i$. To ease notation in this section, we write $p_k(k, d\theta_k)$ for $p_k(k, d\theta_k | \mathbf{y})$. In the case of model selection, the main difficulty for the

Markov chain is to be able to jump from one subspace Θ_n to another subspace Θ_m while preserving the correct invariant distribution. Green [36] has developed a general methodology that addresses this problem.

We recall here that the MH algorithm is a method of constructing a reversible transition kernel for a Markov chain with a specified invariant distribution $p_k(k, d\theta_k)$. The algorithm requires a proposal distribution $q_{k,k^*}(k^*, d\theta_{k^*}^* | k, \theta_k)$ and an acceptance probability $\alpha(k, \theta; k^*, \theta_{k^*}^*)$. If the current state is $(k, d\theta_k)$, a candidate $(k^*, \theta_{k^*}^*)$ for the next state is proposed from $q_{k,k^*}(k^*, d\theta_{k^*}^* | k, \theta_k)$ and accepted with probability $\alpha(k, \theta; k^*, \theta_{k^*}^*)$. Otherwise, the candidate is rejected and the process remains in the state (k, θ) . The transition kernel of the Markov chain can be written

$$K(k, \theta_k; k^*, d\theta_{k^*}^*) = q_{k,k^*}(k^*, d\theta_{k^*}^* | k, \theta_k) \alpha(k, \theta_k; k^*, \theta_{k^*}^*) + \mathbb{I}_{\{k\}}(k^*) \delta_{\theta_k}(d\theta_{k^*}^*) \int_{\Theta} (1 - \alpha(k, \theta_k; l, \gamma)) q_{k,l}(l, d\gamma | k, \theta_k). \tag{7.64}$$

Ensuring reversibility when the problem dimension is fixed, $p_{k^*}(k^*, d\theta_{k^*}^*)$ and $q_{k,k^*}(k^*, d\theta_{k^*}^* | k, \theta_k)$ admit densities $p_{k^*}(k^*, \theta_{k^*}^*)$ and $q_{k,k^*}(k^*, \theta_{k^*}^* | k, \theta_k)$ with respect to the same dominating measure, which is well known, and the following choice for $\alpha(k, \theta_k; k^*, \theta_{k^*}^*)$

$$\alpha(k, \theta_k; k^*, \theta_{k^*}^*) = \min \left\{ 1, \frac{p_{k^*}(k^*, \theta_{k^*}^*) q_{k^*,k}(k, \theta_k | k^*, \theta_{k^*}^*)}{p_k(k, \theta_k) q_{k,k^*}(k^*, \theta_{k^*}^* | k, \theta_k)} \right\} \tag{7.65}$$

is satisfactory.

Returning to the model selection problem, a Markov chain can be constructed using this methodology, where each state of the chain is proposed from the model spaces Θ_m . However, any information common to the different subspaces might be discarded, leading to an inefficient algorithm. A scenario for which this seems to be crucial is that of nested models in which components need to be added or removed more or less independently. This type of approach raises measure theoretic problems, and Green [36] has described a very general methodology for solving this problem. It requires, for each pair of communicating spaces Θ_m and Θ_n ; firstly the definition of extended versions of these subspaces into $\overline{\Theta}_{m,n} \triangleq \Theta_m \times \Psi_{m,n}$ and $\overline{\Theta}_{n,m} \triangleq \Theta_n \times \Psi_{n,m}$; secondly the definition of a deterministic invertible mapping $f_{n,m}$ everywhere differentiable between $\overline{\Theta}_{m,n}$ and $\overline{\Theta}_{n,m}$

$$\begin{aligned} (\theta_m, \varphi_{m,n}) &= f_{n,m}(\theta_n, \varphi_{n,m}) \\ &= (f_{n,m}^\theta(\theta_n, \varphi_{n,m}), f_{n,m}^\varphi(\theta_n, \varphi_{n,m})) \end{aligned} \tag{7.66}$$

(we define $f_{m,n}$ such that $f_{m,n}(f_{n,m}(\theta_n, \varphi_{n,m})) = (\theta_n, \varphi_{n,m})$); and thirdly the choice of proposal densities for $\varphi_{n,m}$ and $\varphi_{m,n}^*$, respectively $q_{n,m}(\cdot | n, \theta_n)$ and $q_{m,n}(\cdot | m, \theta_m)$. The choice of the extended spaces, the deterministic transformation $f_{m,n}$, and the proposal distributions for $q_{n,m}(\cdot | n, \theta_n)$ and $q_{m,n}(\cdot | m, \theta_m)$ is problem dependent and needs to be addressed on a case-by-case basis.

The proposal distribution can now be written

$$q(k^*, d\theta_{k^*}^* | n, \theta_n) = \sum_{m=1}^N \delta_{f_{n,m}^\theta(\theta_n, \varphi_{n,m})} (d\theta_m) \times \quad (7.67)$$

$$\times q_{n,m}(m, d\varphi_{n,m} | k, \theta_{kk}) \mathbb{I}_{\{m\} \times \Theta_m}(k^*, \theta_{k^*}^*).$$

With the assumption that all probability distributions admit a density with respect to the Lebesgue measure, the acceptance probability of a move from Θ_n to Θ_m satisfies

$$r((n, \theta_n), (m, \theta_m^*)) = \quad (7.68)$$

$$\frac{p_m(m, f_{n,m}^\theta(\theta_n, \varphi_{n,m})) q_{m,n}(f_{n,m}^\varphi(\theta_n, \varphi_{n,m}) | m, \theta_m^*)}{p_n(n, d\theta_n) q_{n,m}(\varphi_{n,m} | n, \theta_n)} \times \mathcal{J}_{f_{n,m}},$$

where $\mathcal{J}_{f_{n,m}}$ is the Jacobian of the transformation $f_{n,m}$, when only continuous variables are involved in the transformation, of the invertible mapping $f_{n,m}$ between the spaces Θ_n and Θ_m ,

$$\mathcal{J}_{f_{m,n}} = \left| \det \frac{\partial f_{n,m}(\theta_m, \varphi_{m,n})}{\partial(\theta_m, \varphi_{m,n})} \right|. \quad (7.69)$$

The procedure can be summarized with the following pseudo-code

Reversible Jump MCMC Algorithm

1. Initialization, $i = 0$. Set randomly or deterministically $(k^{(0)}, \theta_{k^{(0)}})$.
2. Iteration $i, i \geq 1$.
 - Sample m from the discrete distribution $q_{k^{(i-1)},m}(m | k^{(i-1)}, \theta_{k^{(i-1)}})$.
 - Sample $\varphi_{k^{(i-1)},m} \sim q_{k^{(i-1)},m}(d\varphi_{k^{(i-1)},m} | m, k^{(i-1)}, \theta_{k^{(i-1)}})$ and perform the invertible transformation

$$(\theta_m^*, \varphi_{m,k^{(i-1)}}^*) = f_{k^{(i-1)},m}(\theta_{k^{(i-1)}}, \varphi_{k^{(i-1)},m}). \quad (7.70)$$

- Accept the move with probability

$$\alpha\left(\left(k^{(i-1)}, \theta_{k^{(i-1)}}\right), (m, \theta_m^*)\right) = \min\left\{1, r\left(\left(k^{(i-1)}, \theta_{k^{(i-1)}}\right), (m, \theta_m^*)\right)\right\} \quad (7.71)$$

i.e., $(k^{(i)}, \theta_{k^{(i)}}) = (m, \theta_m^*)$, otherwise stay at $(k^{(i-1)}, \theta_{k^{(i-1)}})$.

Note that this algorithm is not guaranteed to produce, even asymptotically, samples from the correct distributions, as the two extra properties, namely, irreducibility and aperiodicity (see Subsection 7.3.6), of the Markov chain need to be checked. ■

Simulated Annealing for Global Optimization

Let us assume that we want to find the global maximum of a complex multivariate probability distribution $\pi(d\theta) = \pi(\theta)d\theta$ that we can evaluate pointwise up to a normalizing constant. In a Bayesian framework $\pi(\theta) = p(\theta|\mathbf{y})$ and global maximization allows us to find the MAP estimate. As discussed briefly in 7.3.2, a solution consists of running a (homogeneous) ergodic Markov chain of invariant distribution $\pi(\theta)$ and estimate the global mode by

$$\arg \max_{\theta^{(i)}; i=1, \dots, N} \pi(\theta^{(i)}). \quad (7.72)$$

This method is not efficient in the sense that random samples approximately distributed from $\pi(\theta)$ only rarely explore the vicinity of the mode, unless the distribution has large probability mass around the mode; much computation is thus wasted exploring areas of no interest for global mode estimation. SA methods are a non-homogeneous variant of MCMC to perform global optimization where the invariant distribution at iteration i of the algorithm is the distribution proportional to $\pi^{\gamma(i)}(\theta)$, $\gamma(i)$ being a positive increasing function diverging at infinity. The basic idea is that as $\gamma(i)$ goes to infinity then $\pi^{\gamma(i)}(\theta)$ concentrates itself upon the set of global modes [52], [59] of $\pi(\theta)$. So all the previous MCMC algorithms discussed earlier can be adapted to perform global optimization; the only necessary modification consists of replacing, at iteration i , the invariant distribution $\pi(\theta)$ by

$$\bar{\pi}_{\gamma(i)}(\theta) \propto \pi^{\gamma(i)}(\theta), \quad (7.73)$$

where $\gamma(i) \geq 0$, $\gamma(i+1) \geq \gamma(i)$ and $\lim_{i \rightarrow +\infty} \gamma(i) = +\infty$.

SA algorithms simulate non-homogeneous Markov chains. Convergence results available in the literature mainly state that, if for a given $\gamma(i)$ the homogeneous Markov transition kernel mixes quickly enough, then convergence to the set of global maxima of $\pi(\theta)$ is ensured for a sequence $\gamma(i) = C \ln(i + \gamma_0)$, where C and γ_0 are problem-dependent. Most of the results have been obtained for finite spaces [28], [32], [52], [59] or compact continuous spaces [37]. Some results for non-compact spaces can be found in [4], [8].

7.4 An Example: The Autoregressive Model

In this section we address the problem of Bayesian estimation and model order selection of autoregressive (AR) models. The aim of this section is to illustrate in a simple example how MCMC methods can be used; for the sake of brevity the straightforward implementation of the IS methods is omitted here. We first address the Bayesian estimation of an autoregressive model. In Subsection 7.4.1 we propose several samplers. We emphasize that the

usefulness of these samplers to solve the problem addressed is questionable; the aim of this subsection is to help newcomers to the field understand how to construct MCMC algorithms for a simple problem. Then, in Subsection 7.4.2, we address the more complex problem of model selection and present a reversible jump MCMC algorithm to perform Bayesian computation.

7.4.1 Bayesian Estimation of AR Process

Model and Estimation Objectives

Let us consider an autoregressive model of order k

$$x_t + \sum_{j=1}^k a_j x_{t-j} = \sigma_k v_t, \quad (7.74)$$

where $v_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$. We assume that $T > k$ and that the initial conditions $\mathbf{x}_{1-k:0}$ are equal to zero. We can rewrite this relation in a matrix form

$$\mathbf{x}_{1:T} + \mathbf{X}_k \mathbf{a}_k = \sigma_k \boldsymbol{\nu}_{1:T}, \quad (7.75)$$

where

$$\mathbf{X}_k \triangleq \begin{pmatrix} 0 & \cdots & \cdots & 0 \\ x_1 & 0 & \cdots & 0 \\ x_2 & x_1 & \cdots & 0 \\ \vdots & & & \\ x_{T-2} & \cdots & \cdots & x_{T-2-k+1} \\ x_{T-1} & \cdots & \cdots & x_{T-1-k+1} \end{pmatrix}, \quad (7.76)$$

$\mathbf{a}_k = (a_1, \dots, a_k)^T$ and σ^2 are unknown parameters and we assume from now that they are random with improper prior distribution

$$p(\mathbf{a}_k, \sigma_k^2) \propto 1/\sigma_k^2. \quad (7.77)$$

Given the observations $\mathbf{x}_{1:T}$, our aim is to estimate the posterior distribution $p(\mathbf{a}_k, \sigma_k^2 | \mathbf{x}_{1:T})$. We recall here that no Monte Carlo method is required to estimate this distribution as it can be obtained in closed-form, as detailed later. However conditional expectations, quantiles or other characteristics cannot be obtained analytically; this can justify the use of Monte Carlo methods.

Several Monte Carlo algorithms to sample from $p(\mathbf{a}_k, \sigma_k^2 | \mathbf{x}_{1:T})$ are presented.

Exact Sampler

In this simple case, direct sampling from this distribution $p(\mathbf{a}_k, \sigma_k^2 | \mathbf{x}_{1:T})$ is possible. Indeed,

$$p(\mathbf{a}_k, \sigma_k^2 | \mathbf{x}_{1:T}) = p(\mathbf{a}_k | \mathbf{x}_{1:T}, \sigma_k^2) p(\sigma_k^2 | \mathbf{x}_{1:T}). \quad (7.78)$$

After a few calculations one obtains

$$\sigma_k^2 | \mathbf{x}_{1:T} \sim \mathcal{IG} \left(\frac{T-k}{2}, \frac{\mathbf{x}_{1:T}^T \mathbf{P}_k \mathbf{x}_{1:T}}{2} \right) \quad (7.79)$$

$$\mathbf{a}_k | (\mathbf{x}_{1:T}, \sigma_k^2) \sim \mathcal{N}(\mathbf{m}_k, \sigma_k^2 \mathbf{M}_k), \quad (7.80)$$

where

$$\mathbf{M}_k^{-1} = \mathbf{X}_k^T \mathbf{X}_k \quad (7.81)$$

$$\mathbf{P}_k = \mathbf{I}_T - \mathbf{X}_k \mathbf{M}_k \mathbf{X}_k^T \quad (7.82)$$

$$\mathbf{m}_k = -\mathbf{M}_k \mathbf{X}_k^T \mathbf{x}_{1:T}. \quad (7.83)$$

\mathbf{M}_k^{-1} is assumed to be invertible.

The following algorithm thus provides I.I.D. samples, distributed according to $p(\mathbf{a}_k, \sigma_k^2 | \mathbf{x}_{1:T})$.

- Iteration $i, i \geq 0$.
 1. • Sample $\sigma_k^{2(i)} \sim p(\sigma_k^2 | \mathbf{x}_{1:T})$.
 - Sample $\mathbf{a}_k^{(i)} \sim p(\mathbf{a}_k | \mathbf{x}_{1:T}, \sigma_k^{2(i)})$.

Sampling from those two distributions is standard [21]. ■

Metropolis Algorithm

We adopt the Metropolis algorithm with random Gaussian increment of covariance matrix $\sigma_{RW}^2 \mathbf{I}_{k+1} > 0$. An MCMC algorithm to obtain (non I.I.D.) samples asymptotically distributed according to $p(\mathbf{a}_k, \sigma_k^2 | \mathbf{x}_{1:T})$ proceeds as follows:

1. Initialization, $i = 0$: set randomly or deterministically $(\mathbf{a}_k^{(0)}, \sigma_k^{2(0)})$.
2. Iteration $i, i \geq 1$.
 - Sample $\mathbf{z} \sim \mathcal{N}(0, \sigma_{RW}^2 \mathbf{I}_{k+1})$ and then evaluate the candidate $(\mathbf{a}_k^T, \sigma_k^2) = \left((\mathbf{a}_k^{(i-1)})^T, \sigma_k^{2(i-1)} \right) + \mathbf{z}^T$.
 - Evaluate the acceptance probability

$$\alpha \left(\mathbf{a}_k^{(i-1)}, \sigma_k^{2(i-1)}; \mathbf{a}_k, \sigma_k^2 \right) = \min \left\{ \frac{p(\mathbf{x}_{1:T} | \mathbf{a}_k, \sigma_k^2) \sigma_k^{2(i-1)}}{p(\mathbf{x}_{1:T} | \mathbf{a}_k^{(i-1)}, \sigma_k^{2(i-1)}) \sigma_k^2}, 1 \right\} \quad (7.84)$$

- Sample $u \sim \mathcal{U}_{[0,1]}$; if $u \leq \alpha(\mathbf{a}_k^{(i-1)}, \sigma_k^{2(i-1)}; \mathbf{a}_k, \sigma_k^2)$ then
 $(\mathbf{a}_k^{(i)}, \sigma_k^{2(i)}) = (\mathbf{a}_k, \sigma_k^2)$ else $(\mathbf{a}_k^{(i)}, \sigma_k^{2(i)}) = (\mathbf{a}_k^{(i-1)}, \sigma_k^{2(i-1)})$.

Note that the choice of σ_{RW}^2 will have a strong influence on the convergence and mixing properties of the Markov chain. A tradeoff between good exploration of the total posterior distribution and a good “average” acceptance probability is required. Some theoretical results on this problem are summarized in [52].

Independent Metropolis-Hastings Algorithm

We adopt as a proposal distribution

$\varphi(\mathbf{a}_k, \sigma_k^2) = \mathcal{N}(\mathbf{a}_k; \mathbf{0}_{k \times 1}, \sigma_{\mathbf{a}_k}^2 \mathbf{I}_k) \mathcal{IG}(\sigma_k^2; \frac{\nu_0}{2}, \frac{\gamma_0}{2})$ ($\nu_0 \gamma_0 > 0$). The MCMC algorithm proceeds as follows.

1. Initialization, $i = 0$: set randomly or deterministically $(\mathbf{a}_k^{(0)}, \sigma_k^{2(0)})$.

2. Iteration $i, i \geq 1$.

- Sample a candidate $\mathbf{a}_k \sim \mathcal{N}(\mathbf{0}_{k \times 1}, \sigma_{\mathbf{a}_k}^2 \mathbf{I}_k)$ and $\sigma_k^2 \sim \mathcal{IG}(\frac{\nu_0}{2}, \frac{\gamma_0}{2})$.
- Evaluate the acceptance probability

$$\alpha(\mathbf{a}_k^{(i-1)}, \sigma_k^{2(i-1)}; \mathbf{a}_k, \sigma_k^2) = \min \left\{ \frac{p(\mathbf{x}_{1:T} | \mathbf{a}_k, \sigma_k^2) \sigma_k^{2(i-1)} \varphi(\mathbf{a}_k^{(i-1)}, \sigma_k^{2(i-1)})}{p(\mathbf{x}_{1:T} | \mathbf{a}_k^{(i-1)}, \sigma_k^{2(i-1)}) \sigma_k^2 \varphi(\mathbf{a}_k, \sigma_k^2)}, 1 \right\} \quad (7.85)$$

- Sample $u \sim \mathcal{U}_{[0,1]}$; if $u \leq \alpha(\mathbf{a}_k^{(i-1)}, \sigma_k^{2(i-1)}; \mathbf{a}_k, \sigma_k^2)$, then
 $(\mathbf{a}_k^{(i)}, \sigma_k^{2(i)}) = (\mathbf{a}_k, \sigma_k^2)$ else $(\mathbf{a}_k^{(i)}, \sigma_k^{2(i)}) = (\mathbf{a}_k^{(i-1)}, \sigma_k^{2(i-1)})$.

Here again the choice of the parameters of $\varphi(\mathbf{a}_k, \sigma_k^2)$ will have a strong influence on the convergence properties of the Markov chain. The closer $\varphi(\mathbf{a}_k, \sigma_k^2)$ is to the target distribution $p(\mathbf{a}_k, \sigma_k^2 | \mathbf{x}_{1:T})$, the more likely the algorithm is to be efficient as the acceptance probability will be close to 1 on the average. Note that the support of the target has to be included in the support of the proposal distribution to ensure that the whole posterior distribution is explored.

MH Sampler One-at-a-Time

In this sampler, we update \mathbf{a}_k using a Metropolis step with a random increment $\sigma_{RW}^2 \mathbf{I}_k > 0$ and σ_k^2 using an independent MH step with proposal distribution $\mathcal{IG}(\sigma_k^2; \frac{\nu_0}{2}, \frac{\gamma_0}{2})$ ($\nu_0 \gamma_0 > 0$). The algorithm proceeds as follows.

1. Initialization, $i = 0$: set randomly or deterministically $(\mathbf{a}_k^{(0)}, \sigma_k^{2(0)})$.

2. Iteration $i, i \geq 1$.

Metropolis Step

- Sample $\mathbf{z} \sim \mathcal{N}(0, \sigma_{RW}^2 \mathbf{I}_k)$ and evaluate the candidate $\mathbf{a}_k = \mathbf{a}_k^{(i-1)} + \mathbf{z}$.
- Evaluate the acceptance probability

$$\alpha(\mathbf{a}_k^{(i-1)}, \sigma_k^{2(i-1)}; \mathbf{a}_k, \sigma_k^{2(i-1)}) = \min \left\{ \frac{p(\mathbf{x}_{1:T} | \mathbf{a}_k, \sigma_k^{2(i-1)})}{p(\mathbf{x}_{1:T} | \mathbf{a}_k^{(i-1)}, \sigma_k^{2(i-1)})}, 1 \right\} \tag{7.86}$$

- Sample $u \sim \mathcal{U}_{[0,1]}$; if $u \leq \alpha(\mathbf{a}_k^{(i-1)}, \sigma_k^{2(i-1)}; \mathbf{a}_k, \sigma_k^{2(i-1)})$, then $\mathbf{a}_k^{(i)} = \mathbf{a}_k$ else $\mathbf{a}_k^{(i)} = \mathbf{a}_k^{(i-1)}$.

Independent MH Step

- Sample a candidate $\sigma_k^2 \sim \mathcal{IG}(\frac{\nu_0}{2}, \frac{\gamma_0}{2})$.
- Evaluate the acceptance probability

$$\alpha(\mathbf{a}_k^{(i)}, \sigma_k^{2(i-1)}; \mathbf{a}_k^{(i)}, \sigma_k^2) = \min \left\{ \frac{p(\mathbf{x}_{1:T} | \mathbf{a}_k^{(i)}, \sigma_k^2) \sigma_k^{2(i-1)} \mathcal{IG}(\sigma_k^{2(i-1)}; \frac{\nu_0}{2}, \frac{\gamma_0}{2})}{p(\mathbf{x}_{1:T} | \mathbf{a}_k^{(i)}, \sigma_k^{2(i-1)}) \sigma_k^2 \mathcal{IG}(\sigma_k^2; \frac{\nu_0}{2}, \frac{\gamma_0}{2})}, 1 \right\}. \tag{7.87}$$

- Sample $u \sim \mathcal{U}_{[0,1]}$, if $u \leq \alpha(\mathbf{a}_k^{(i)}, \sigma_k^{2(i-1)}; \mathbf{a}_k^{(i)}, \sigma_k^2)$ then $\sigma_k^{2(i)} = \sigma_k^2$; else $\sigma_k^{2(i)} = \sigma_k^{2(i-1)}$.

Metropolis/Gibbs Sampler

In this sampler, we first update \mathbf{a}_k using a Gibbs step and σ_k^2 using an independent step with proposal distribution $\mathcal{IG}(\sigma_k^2; \frac{\nu_0}{2}, \frac{\gamma_0}{2})$ ($\nu_0 \gamma_0 > 0$). The algorithm proceeds as follows.

1. Initialization, $i = 0$: set randomly or deterministically $(\mathbf{a}_k^{(0)}, \sigma_k^{2(0)})$.

2. Iteration $i, i \geq 1$.

Gibbs Step

- Sample $\mathbf{a}_k^{(i)} \sim p(\mathbf{a}_k | \mathbf{x}_{1:T}, \sigma_k^{2(i-1)})$.

Independent MH Step

- Sample a candidate $\sigma_k^2 \sim IG\left(\frac{\nu_0}{2}, \frac{\gamma_0}{2}\right)$.
- Evaluate the acceptance probability

$$\alpha\left(\mathbf{a}_k^{(i)}, \sigma_k^{2(i-1)}; \mathbf{a}_k^{(i)}, \sigma_k^2\right) = \min\left\{\frac{p(\mathbf{x}_{1:T}|\mathbf{a}_k^{(i)}, \sigma_k^2)\sigma_k^{2(i-1)}IG\left(\sigma_k^{2(i-1)}; \frac{\nu_0}{2}, \frac{\gamma_0}{2}\right)}{p(\mathbf{x}_{1:T}|\mathbf{a}_k^{(i)}, \sigma_k^{2(i-1)})\sigma_k^2IG\left(\sigma_k^2; \frac{\nu_0}{2}, \frac{\gamma_0}{2}\right)}, 1\right\}. \quad (7.88)$$

- Sample $u \sim \mathcal{U}_{[0,1]}$, if $u \leq \alpha\left(\mathbf{a}_k^{(i)}, \sigma_k^{2(i-1)}; \mathbf{a}_k^{(i)}, \sigma_k^2\right)$ then $\sigma_k^{2(i)} = \sigma_k^2$; else $\sigma_k^{2(i)} = \sigma_k^{2(i-1)}$.

Gibbs Sampler

Finally, we present a Gibbs sampler algorithm to sample from $p(\mathbf{a}_k, \sigma_k^2 | \mathbf{x}_{1:T})$.

1. Initialization, $i = 0$: set randomly or deterministically $(\mathbf{a}_k^{(0)}, \sigma_k^{2(0)})$.
2. Iteration $i, i \geq 1$.
 - Sample $\mathbf{a}_k^{(i)} \sim p(\mathbf{a}_k | \mathbf{x}_{1:T}, \sigma_k^{2(i-1)})$.
 - Sample $\sigma_k^{2(i)} \sim p(\sigma_k^2 | \mathbf{x}_{1:T}, \mathbf{a}_k^{(i)})$.

The expression of $p(\mathbf{a}_k | \mathbf{x}_{1:T}, \sigma_k^{2(i)})$ is given by (7.80) and one obtains $p(\sigma_k^2 | \mathbf{x}_{1:T}, \mathbf{a}_k)$ from Bayes theorem

$$\sigma_k^2 | (\mathbf{x}_{1:T}, \mathbf{a}_k^{(i)}) \sim IG\left(\frac{T}{2}, \frac{(\mathbf{x}_{1:T} + \mathbf{X}_k \mathbf{a}_k^{(i)})^\top (\mathbf{x}_{1:T} + \mathbf{X}_k \mathbf{a}_k^{(i)})}{2}\right). \quad (7.89)$$

We have illustrated in a simple example the different strategies existing to sample from the posterior distribution $p(\mathbf{a}_k, \sigma_k^2 | \mathbf{x}_{1:T})$. The problem of the choice of the strategy is of course easily solved. Indeed it is possible to exactly draw I.I.D. samples from the target distribution with a low computational cost using standard methods. For more complex models, such as those described in Section 7.5, the design of a “good” sampler relies more or less on heuristic considerations and experience.

7.4.2 Bayesian Model Selection of AR Process

Model and Estimation Objectives

We now assume that the model order of the AR process is part of the unknown parameters of the problem, i.e., we assume that the data can be represented by one of the following models \mathcal{M}_k , $k \in \{0, \dots, k_{\max}\}$ ($k_{\max} < T$)

$$\mathcal{M}_k : x_t + \sum_{i=1}^k a_i x_{t-i} = \sigma_k v_t. \tag{7.90}$$

The unknown parameters $(k, \mathbf{a}_k, \sigma_k^2)$ are assumed random with a given prior distribution. We recall here that in a Bayesian framework, the choice of an improper prior for model selection purposes can lead to Lindley’s paradox [9], that is the systematic choice of the simplest model. In this tutorial, we use a proper but vague hierarchical prior such as the one proposed in [48]. The following prior distribution is selected

$$p(k, \mathbf{a}_k, \sigma_k^2, \delta^2) = p(\mathbf{a}_k | k, \sigma_k^2, \delta^2) p(\sigma_k^2) p(\delta^2) p(k). \tag{7.91}$$

For the autoregressive coefficients we chose a data-dependent prior

$$p(\mathbf{a}_k | k, \sigma_k^2, \delta^2) = \left| \frac{\mathbf{X}_k^T \mathbf{X}_k}{2\pi \delta^2 \sigma_k^2} \right|^{1/2} \exp \left[-\frac{\mathbf{a}_k^T \mathbf{X}_k^T \mathbf{X}_k \mathbf{a}_k}{2\delta^2 \sigma_k^2} \right], \tag{7.92}$$

and we assume that

$$\sigma_k^2 \sim \text{IG} \left(\frac{\nu_0}{2}, \frac{\gamma_0}{2} \right), \delta^2 \sim \text{IG} \left(\frac{\nu_{0,\delta^2}}{2}, \frac{\gamma_{0,\delta^2}}{2} \right), k \sim \mathcal{U}_{\{0, \dots, k_{\max}\}}. \tag{7.93}$$

Given the observations $\mathbf{x}_{1:T}$, our aim is to estimate the posterior distribution $p(k, \mathbf{a}_k, \sigma_k^2, \delta^2 | \mathbf{x}_{1:T})$. This distribution does not admit any analytical expression. Bayesian computation is performed by developing a reversible jump MCMC sampler.

Reversible Jump MCMC Algorithm

We propose a reversible jump MCMC sampler that makes the most of the analytical properties of the probabilistic model to estimate the joint posterior distribution $p(k, \mathbf{a}_k, \sigma_k^2, \delta^2 | \mathbf{x}_{1:T})$. For our problem, the following moves have been selected:

1. birth of a new AR coefficient.
2. death of an existing AR coefficient.
3. update of the parameters.

The birth and death moves perform dimension changes from respectively k to $k + 1$ and k to $k - 1$. These moves are defined by heuristic considerations, the only condition to be fulfilled being to maintain the correct invariant distribution. The resulting transition kernel of the simulated Markov chain is then a mixture of the different transition kernels associated with the moves described earlier. This means that at each iteration one of the candidate moves: birth, death or update is randomly chosen. The probabilities for choosing these moves are b_k , d_k and u_k respectively, such that $b_k + d_k + u_k = 1$ for all $0 \leq k \leq k_{\max}$. The move is performed if the algorithm accepts it. For $k = 0$ the death move is impossible, so that $d_0 \triangleq 0$. For $k = k_{\max}$ the birth move is impossible and thus $b_{k_{\max}} \triangleq 0$. Except in these cases, we take the following probabilities $b_k = d_k = u_k$. The algorithm proceeds as follows:

Algorithm for AR Model Order Estimation

1. Initialization, $i = 0$: set $\left(k^{(0)}, \mathbf{a}_{k^{(0)}}^{(0)}, \sigma_{k^{(0)}}^{2(0)}, \delta^{2(0)}\right)$, randomly or deterministically.

2. Iteration i , $i \geq 1$.

- If $(u \sim \mathcal{U}_{[0,1]}) < u_{k^{(i-1)}}$

Update move

- Set $k^{(i)} = k^{(i-1)}$ and sample

$$\left(\mathbf{a}_{k^{(i)}}^{(i)}, \sigma_{k^{(i)}}^{2(i)}\right) \sim p\left(\mathbf{a}_{k^{(i)}}, \sigma_{k^{(i)}}^2 \mid \mathbf{x}_{1:T}, k^{(i)}, \delta^{2(i-1)}\right), \quad (7.94)$$

Else if $u < (u_{k^{(i-1)}} + b_{k^{(i-1)}})$

Birth move

- Evaluate $\alpha_{birth}(k^{(i-1)}, k^{(i-1)} + 1)$, see (7.103), and sample $u_{birth} \sim \mathcal{U}_{[0,1]}$.
- If $u_{birth} < \alpha_{birth}(k^{(i-1)}, k^{(i-1)} + 1)$ set $k^{(i)} = k^{(i-1)} + 1$ and sample

$$\left(\mathbf{a}_{k^{(i)}}^{(i)}, \sigma_{k^{(i)}}^{2(i)}\right) \sim p\left(\mathbf{a}_{k^{(i)}}, \sigma_{k^{(i)}}^2 \mid \mathbf{x}_{1:T}, k^{(i)}, \delta^{2(i-1)}\right), \quad (7.95)$$

Else set $\left(k^{(i)}, \mathbf{a}_{k^{(i)}}^{(i)}, \sigma_{k^{(i)}}^{2(i)}\right) = \left(k^{(i-1)}, \mathbf{a}_{k^{(i-1)}}^{(i-1)}, \sigma_{k^{(i-1)}}^{2(i-1)}\right)$.

Else

Death move

- Evaluate $\alpha_{death}(k^{(i-1)}, k^{(i-1)} - 1)$, see (7.104), and sample $u_{death} \sim \mathcal{U}_{[0,1]}$.

– If $u_{death} < \alpha_{death} (k^{(i-1)}, k^{(i-1)} - 1)$ Set $k^{(i)} = k^{(i-1)} - 1$ and sample

$$\left(\mathbf{a}_{k^{(i)}}^{(i)}, \sigma_{k^{(i)}}^2 \right) \sim p \left(\mathbf{a}_{k^{(i)}}^{(i)}, \sigma_{k^{(i)}}^2 \mid \mathbf{x}_{1:T}, k^{(i)}, \delta^2 \right), \quad (7.96)$$

Else set $\left(k^{(i)}, \mathbf{a}_{k^{(i)}}^{(i)}, \sigma_{k^{(i)}}^2 \right) = \left(k^{(i-1)}, \mathbf{a}_{k^{(i-1)}}^{(i-1)}, \sigma_{k^{(i-1)}}^2 \right)$.

- Sample $\delta^2 \sim p \left(\delta^2 \mid k^{(i)}, \mathbf{a}_{k^{(i)}}^{(i)}, \sigma_{k^{(i)}}^2 \right)$.

The conditional distributions $p \left(\mathbf{a}_k, \sigma_k^2 \mid \mathbf{x}_{1:T}, k, \delta^2 \right)$ and $p \left(\delta^2 \mid \mathbf{x}_{1:T}, k, \mathbf{a}_k, \sigma_k^2 \right)$ satisfy

$$\sigma_k^2 \mid \left(\mathbf{x}_{1:T}, k, \delta^2 \right) \sim \text{IG} \left(\frac{T + \nu_0}{2}, \frac{\gamma_0 + \mathbf{x}_{1:T}^\top \mathbf{P}_k \mathbf{x}_{1:T}}{2} \right) \quad (7.97)$$

$$\mathbf{a}_k \mid \left(\mathbf{x}_{1:T}, k, \sigma_k^2, \delta^2 \right) \sim \mathcal{N} \left(\mathbf{m}_k, \sigma_k^2 \mathbf{M}_k \right) \quad (7.98)$$

$$\delta^2 \mid \left(\mathbf{x}_{1:T}, k, \mathbf{a}_k, \sigma_k^2 \right) \sim \text{IG} \left(\frac{k + \nu_{0,\delta^2}}{2}, \frac{\gamma_{0,\delta^2}}{2} + \frac{\mathbf{a}_k^\top \mathbf{X}_k^\top \mathbf{X}_k \mathbf{a}_k}{2\sigma_k^2} \right), \quad (7.99)$$

where we have defined the following quantities

$$\mathbf{M}_k^{-1} = (1 + \delta^{-2}) \mathbf{X}_k^\top \mathbf{X}_k \quad (7.100)$$

$$\mathbf{m}_k = -\mathbf{M}_k \mathbf{X}_k^\top \mathbf{x}_{1:T} \quad (7.101)$$

$$\mathbf{P}_k = \mathbf{I}_T - \mathbf{X}_k^\top \mathbf{M}_k \mathbf{X}_k. \quad (7.102)$$

The acceptance probabilities of the birth and death moves are equal to

$$\alpha_{birth} (k, k + 1) = \min \{ 1, r_{birth} (k, k + 1) \} \quad (7.103)$$

$$\alpha_{death} (k + 1, k) = \min \{ 1, r_{birth}^{-1} (k, k + 1) \} \quad (7.104)$$

with

$$r_{death}^{-1} (k + 1, k) = r_{birth} (k, k + 1) = \left(\frac{\gamma_0 + \mathbf{x}_{1:T}^\top \mathbf{P}_k \mathbf{x}_{1:T}}{\gamma_0 + \mathbf{x}_{1:T}^\top \mathbf{P}_{k+1} \mathbf{x}_{1:T}} \right)^{\frac{T + \nu_0}{2}} \frac{1}{\sqrt{1 + \delta^2}}.$$

It is worth noticing that the acceptance probabilities of this algorithm for the dimension changes do not depend on the current value of the variance of the noise or on the value of the current parameters.

7.5 Applications

In this section, we show how to successfully apply some Monte Carlo methods to solve complex nonlinear and non Gaussian data analysis problems.

7.5.1 Sequential Importance Sampling for Optimal Filtering

We show in this section how to apply Monte Carlo methods to solve sequential optimal estimation problems for non-linear non-Gaussian problems; see [39], [60] for many potential applications. For the sake of simplicity, we focus on an important application for Radar: bearings-only tracking.

Problem Formulation and Estimation Objectives

Consider the following problem. A target is moving in a two-dimensional plane. The target is located by the parameters $(x_t, y_t)_{t=1, \dots}$ and moves at speed $(\dot{x}_t, \dot{y}_t)_{t \in \mathbb{N}}$. The system fluctuations $(\gamma_t^x, \gamma_t^y)_{t \in \mathbb{N}}$ are independent zero-mean Gaussian white noise,

$$\mathbf{v}_t \sim \mathcal{N}(\mathbf{0}_{2 \times 1}, \text{diag}(\sigma_x^2, \sigma_y^2)). \quad (7.105)$$

Discretization of the classical mechanics equation leads to the following state-space representation of the dynamic system

$$\mathbf{z}_{t+1} = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix} \mathbf{z}_t + \begin{pmatrix} 0.5 & 0 \\ 1 & 0 \\ 0 & 0.5 \\ 0 & 1 \end{pmatrix} \mathbf{v}_{t+1}, \quad (7.106)$$

where $\mathbf{z}_t \triangleq (x_t \quad \dot{x}_t \quad y_t \quad \dot{y}_t)^\top$. The observations are a sequence $(\phi_t)_{t \in \mathbb{N}}$ of bearings

$$\phi_t = \arctan\left(\frac{y_t}{x_t}\right) + n_t, \quad (7.107)$$

where $n_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma_n^2)$ is a white Gaussian noise accounting for measurement errors. The prior distribution for the initial state \mathbf{z}_0 is assumed to be Gaussian with known covariance. The parameters σ_x^2 , σ_y^2 and σ_n^2 are assumed known. Equations (7.106)-(7.107) define a Bayesian model as (7.106) defines a prior distribution and (7.107) defines the likelihood.

Our objective is to estimate *on-line* the configuration \mathbf{z}_t of the object from the observations $\phi_{1:t}$ and more precisely the filtering distribution $p(\mathbf{z}_t | \phi_{1:t})$ and its associated features. This is an optimal filtering problem which does not admit any closed-form analytical solution. A suboptimal solution would consist of using the Extended Kalman filter (EKF) [39]. However, it does not yield, in this case, satisfactory results while simulation-based methods perform very well [35], [27].

In this sequential framework, the iterative MCMC methods cannot be used. We show here how a sequential version of Bayesian importance sampling can be used to solve this problem and more generally any optimal filtering problems; see [35], [24], [27] for additional details on sequential Monte Carlo methods.

Bayesian Computation via Sequential Importance Sampling

Let us consider the joint distribution of the states $p(\mathbf{z}_{0:t}|\phi_{1:t})$. We can approximate this distribution using the Bayesian importance sampling method described in 7.3.5, i.e., we select an arbitrary importance function $q(\mathbf{z}_{0:t}|\phi_{1:t})$ from which it is easy to sample (we emphasize here that the importance function $q(\cdot)$ can depend on the observations $\phi_{1:t}$). Then having obtained $N \gg 1$ samples from $q(\mathbf{z}_{0:t}|\phi_{1:t})$, one straightforwardly obtains an approximation $\hat{p}_N(\mathbf{z}_{0:t}|\phi_{1:t})$ of $p(d\mathbf{z}_{0:t}|\phi_{1:t})$

$$\hat{p}_N(d\mathbf{z}_{0:t}|\phi_{1:t}) = \frac{\sum_{i=1}^N w(\mathbf{z}_{0:t}^{(i)}) \delta_{\mathbf{z}_{0:t}^{(i)}}(d\mathbf{z}_{0:t})}{\sum_{j=1}^N w(\mathbf{z}_{0:t}^{(j)})}, \quad (7.108)$$

where the importance weights satisfy

$$w(\mathbf{z}_{0:t}^{(i)}) \propto \frac{p(\mathbf{z}_{0:t}^{(i)}|\phi_{1:t})}{q(\mathbf{z}_{0:t}^{(i)}|\phi_{1:t})}. \quad (7.109)$$

This method is not a sequential method. However, if one restricts the importance function to have the following form

$$q(\mathbf{z}_{0:t}|\phi_{1:t}) = q(\mathbf{z}_0) \prod_{t=1}^n q(\mathbf{z}_t|\mathbf{z}_{0:t-1}, \phi_{1:t}), \quad (7.110)$$

then we are able to simulate the trajectories $\mathbf{z}_{0:t}^{(i)}$ and to evaluate the associated importance weights $w(\mathbf{z}_{0:t}^{(i)})$ sequentially in time. This specialized version of importance sampling is named sequential importance sampling (SIS).

Algorithms

At time $t-1$, assume that we have N samples with their associated weights $(\mathbf{z}_{0:t-1}^{(i)}, w(\mathbf{z}_{0:t-1}^{(i)}); i = 1, \dots, N)$. Then SIS proceeds as follows at time t to obtain $(\mathbf{z}_{0:t-1}^{(i)}, w(\mathbf{z}_{0:t}^{(i)}); i = 1, \dots, N)$.

SIS for Optimal Filtering

- For $i = 1, \dots, N$, sample $\mathbf{z}_t^{(i)} \sim q(\mathbf{z}_t|\mathbf{z}_{0:t-1}, \phi_{0:t})$ and set $\mathbf{z}_{0:t}^{(i)} \triangleq (\mathbf{z}_{0:t-1}^{(i)}, \mathbf{z}_t^{(i)})$.
- For $i = 1, \dots, N$, evaluate the importance weights up to a normalizing constant

$$w(\mathbf{z}_{0:t}^{(i)}) \propto w(\mathbf{z}_{0:t-1}^{(i)}) \frac{p(\phi_t|\mathbf{z}_t^{(i)}) p(\mathbf{z}_t^{(i)}|\mathbf{z}_{0:t-1}^{(i)})}{q(\mathbf{z}_t^{(i)}|\mathbf{z}_{0:t-1}^{(i)}, \phi_{0:t})}. \quad (7.111)$$

- For $i = 1, \dots, N$, normalize the importance weights

$$\tilde{w}_t^{(i)} = \frac{w(\mathbf{z}_{0:t}^{(i)})}{\sum_{j=1}^N w(\mathbf{z}_{0:t}^{(j)})}. \quad (7.112)$$

This algorithm is very general and it is of interest to note that it can be straightforwardly implemented on a parallel computer. Unfortunately, it can be shown that it is inefficient in the sense that the variance of the importance weights can only increase (stochastically) over time. In practice, after a few time steps, all but one of the normalized importance weights are very close to zero and a large computational effort is devoted to updating trajectories whose contribution to the final estimate is almost zero. The key step to make this algorithm work in practice consists of introducing a resampling step at each time t . This idea was first proposed by Gordon *et al.* [35] and is based on the SIR algorithm. The modified algorithm proceeds as follows at time t .

SIS/Resampling for Optimal Filtering

Sampling Step

- For $i = 1, \dots, N$, sample $\tilde{\mathbf{z}}_t^{(i)} \sim q(\tilde{\mathbf{z}}_t | \mathbf{z}_{0:t-1}^{(i)}, \phi_{0:t})$ and set $\tilde{\mathbf{z}}_{0:t}^{(i)} \triangleq (\mathbf{z}_{0:t-1}^{(i)}, \tilde{\mathbf{z}}_t^{(i)})$.
- For $i = 1, \dots, N$, evaluate the importance weights up to a normalizing constant

$$w_t^{(i)} \propto \frac{p(\phi_t | \mathbf{z}_t^{(i)}) p(\mathbf{z}_t^{(i)} | \mathbf{z}_{t-1}^{(i)})}{q(\mathbf{z}_t^{(i)} | \mathbf{z}_{0:t-1}^{(i)}, \phi_{0:t})}. \quad (7.113)$$

- For $i = 1, \dots, N$, normalize the importance weights

$$\tilde{w}_t^{(i)} = \frac{w_t^{(i)}}{\sum_{j=1}^N w_t^{(j)}}. \quad (7.114)$$

Resampling Step

- For $i = 1, \dots, N$, sample an index $j^{(i)}$ distributed according to the discrete distribution $\Pr\{j^{(i)} = l\} = \tilde{w}_t^{(l)}$ for $l = 1, \dots, N$.
 - For $i = 1, \dots, N$, set $\mathbf{z}_{0:t}^{(i)} = \tilde{\mathbf{z}}_{0:t}^{j^{(i)}}$ and reset the weights $\tilde{w}_t^{(i)} = 1/N$.
-

One obtains at time t the following approximation of $p(\mathbf{z}_{0:t} | \phi_{1:t})$

$$\hat{p}_N(d\mathbf{z}_{0:t} | \phi_{1:t}) = \frac{1}{N} \sum_{i=1}^N \delta_{\mathbf{z}_{0:t}^{(i)}}(d\mathbf{z}_{0:t}). \quad (7.115)$$

Simulations

The parameters of the model are identical to those in [35]. To run the algorithm, it is necessary to chose the importance function $q(\mathbf{z}_t | \mathbf{z}_{t-1}^{(i)}, \phi_{0:t})$. We select here $q(\mathbf{z}_t | \mathbf{z}_{t-1}^{(i)}, \phi_{0:t}) = p(\mathbf{z}_t | \mathbf{z}_{t-1}^{(i)})$, that is the prior distribution. Other choices are presented in [24], [27]. We present an application of the SIS/SIR algorithm for $N = 1000$ samples. The results are displayed in Figure 7.1.




FIGURE 7.1. Left: noisy observations of bearings - Right: true trajectory (solid line) / $\mathbb{E}(\mathbf{z}_t | \mathbf{y}_t)$, estimated trajectory (dotted line)

For this problem, it appears that the EKF algorithm diverges if it is not well initialized, whereas the sequential Monte Carlo method performs remarkably well. Many other developments and applications of sequential Monte Carlo methods can be found in [27].

7.5.2 Blind Deconvolution of Impulsive Process

Numerous phenomena arising in a variety of fields of science are composed of isolated time events occurring at random instants. This is the case in the study of neuronal electrical activity, seismic phenomena, radioactivity etc. Particularly well adapted and frequently used models for such highly discontinuous phenomena are point processes. In many cases the point process cannot be directly observed as it is filtered and corrupted by observation noise [14], [18], [22], [45]limited propagation medium. This medium eventually includes the measurement sensor.

Problem Formulation and Estimation Objectives

The observed signal y_t is modeled as the convolution of the sequence v_t by an AR model $\mathbf{a} = [a_1, \dots, a_k]^T$ observed in white Gaussian noise, i.e.,

$$y_t = x_t + w_t \tag{7.116}$$

$$x_t = \sum_{i=1}^k a_i x_{k-i} + v_t \tag{7.117}$$

where $w_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma_w^2)$, the initial values of the AR process are assumed equal to zero. In this model v_t is an I.I.D. sequence of marginal distribution

$$v_t \stackrel{i.i.d.}{\sim} \lambda \mathcal{N}(0, \sigma_v^2) + (1 - \lambda) \mathcal{N}(0, \alpha^2 \sigma_v^2) \quad (7.118)$$

with $\alpha^2 < 1$ being assumed known, typically $\alpha^2 \ll 1$. The sequence v_t and the parameters $\theta = (\mathbf{a}, \sigma_v^2, \sigma_w^2, \lambda)$ are unknown. It is convenient from an algorithmical point of view to introduce the latent Bernoulli process $r_t \in \{0, 1\}$ so that $\Pr(r_t = 1) = \lambda$ and

$$v_t | (r_t = 1) \sim \mathcal{N}(0, \sigma_v^2), \quad v_t | (r_t = 0) \sim \mathcal{N}(0, \alpha^2 \sigma_v^2) \quad (7.119)$$

We assign a prior distribution to the unknown parameters θ so that

$$p(\theta) = p(\mathbf{a}, \sigma_v^2, \sigma_w^2, \lambda) = p(\mathbf{a} | \sigma_v^2) p(\sigma_v^2) p(\sigma_w^2) p(\lambda). \quad (7.120)$$

For the AR model and excitation noise variance, a normal-inverse gamma prior distribution is selected, i.e.,

$$\mathbf{a} | \sigma_v^2 \sim \mathcal{N}(\mathbf{0}_{k \times 1}, \sigma_v^2 \Sigma_0) \quad \text{and} \quad \sigma_v^2 \sim \mathcal{IG}\left(\frac{\nu_0}{2}, \frac{\gamma_0}{2}\right) \quad (7.121)$$

with Σ_0 a regular matrix and

$$\sigma_w^2 \sim \mathcal{IG}\left(\frac{\alpha_0}{2}, \frac{\beta_0}{2}\right) \quad \text{and} \quad \lambda \sim \mathcal{U}_{[0,1]}. \quad (7.122)$$

The hyperparameters $(\Sigma_0, \nu_0, \gamma_0, \alpha_0, \beta_0)$ can be set so that this distribution is uninformative.

Given the set of observations $\mathbf{y}_{1:T} \triangleq \{y_1, \dots, y_T\}$, our aim is to estimate the joint posterior distribution $p(\theta, \mathbf{r}_{1:T}, \mathbf{x}_{1:T} | \mathbf{y}_{1:T})$. This distribution does not admit an analytical expression. We propose different MCMC samplers to perform Bayesian computation.

Bayesian Computation Using MCMC

We propose two MCMC samplers to compute $p(\theta, \mathbf{r}_{1:T}, \mathbf{x}_{1:T} | \mathbf{y}_{1:T})$. The first sampler is a simple Gibbs sampler. We show that this sampler is unfortunately totally inefficient and can not even converge in important cases. We then propose an alternative Gibbs sampler which solves this problem.

A Gibbs Sampler

A Gibbs sampler to sample from $p(\theta, \mathbf{r}_{1:T}, \mathbf{x}_{1:T} | \mathbf{y}_{1:T})$ proceeds as follows.

Gibbs Sampler for Blind Deconvolution

1. Initialization, $i = 0$: set randomly or deterministically $(\theta^{(0)}, \mathbf{r}_{1:T}^{(0)}, \mathbf{x}_{1:T}^{(0)})$.

2. Iteration i , $i \geq 1$.

- Sample $\theta^{(i)} \sim p\left(\theta \mid \mathbf{y}_{1:T}, \mathbf{r}_{1:T}^{(i-1)}, \mathbf{x}_{1:T}^{(i-1)}\right)$.
- Sample $\mathbf{r}_{1:T}^{(i)} \sim p\left(\mathbf{r}_{1:T} \mid \mathbf{y}_{1:T}, \theta^{(i)}, \mathbf{x}_{1:T}^{(i-1)}\right)$.
- Sample $\mathbf{x}_{1:T}^{(i)} \sim p\left(\mathbf{x}_{1:T} \mid \mathbf{y}_{1:T}, \theta^{(i)}, \mathbf{r}_{1:T}^{(i)}\right)$.

Sampling from $p(\theta \mid \mathbf{y}_{1:T}, \mathbf{r}_{1:T}, \mathbf{x}_{1:T})$ is straightforward as

$$\begin{aligned} & p(\mathbf{a}, \sigma_v^2, \sigma_w^2, \lambda \mid \mathbf{y}_{1:T}, \mathbf{r}_{1:T}, \mathbf{x}_{1:T}) \\ &= p(\mathbf{a}, \sigma_v^2 \mid \mathbf{y}_{1:T}, \mathbf{r}_{1:T}, \mathbf{x}_{1:T}) p(\sigma_w^2 \mid \mathbf{y}_{1:T}, \mathbf{x}_{1:T}) p(\lambda \mid \mathbf{r}_{1:T}) \end{aligned} \quad (7.123)$$

and all these conditional distributions can be evaluated analytically and are easy to sample. Sampling from $p(\mathbf{r}_{1:T} \mid \mathbf{y}_{1:T}, \theta, \mathbf{x}_{1:T}) = p(\mathbf{r}_{1:T} \mid \mathbf{a}, \lambda, \mathbf{x}_{1:T})$ is easily done as $(\mathbf{a}, \mathbf{x}_{1:T})$ allows us to evaluate $\mathbf{v}_{1:T}$ and

$$p(\mathbf{r}_{1:T} \mid \sigma_v^2, \lambda, \mathbf{v}_{1:T}) = \prod_{t=1}^T p(r_t \mid \sigma_v^2, \lambda, v_t), \quad (7.124)$$

where $p(r_t \mid \sigma_v^2, \lambda, v_t)$ is a discrete distribution from which it is easy to sample:

$$\begin{aligned} \Pr(r_t = 1 \mid \sigma_v^2, \lambda, v_t) &= 1 - \Pr(r_t = 0 \mid \sigma_v^2, \lambda, v_t) \\ &= \frac{\lambda \mathcal{N}(v_t; 0, \sigma_v^2)}{\lambda \mathcal{N}(v_t; 0, \sigma_v^2) + (1 - \lambda) \mathcal{N}(v_t; 0, \alpha^2 \sigma_v^2)}. \end{aligned} \quad (7.125)$$

Finally sampling from $p(\mathbf{x}_{1:T} \mid \mathbf{y}_{1:T}, \theta, \mathbf{r}_{1:T}) = p(\mathbf{x}_{1:T} \mid \mathbf{y}_{1:T}, \mathbf{a}, \sigma_v^2, \sigma_w^2, \mathbf{r}_{1:T})$ is a more complex task. This distribution is clearly a Gaussian distribution of mean $\mu(\mathbf{y}_{1:T}, \mathbf{a}, \sigma_v^2, \sigma_w^2, \mathbf{r}_{1:T})$ and covariance $\Sigma(\mathbf{y}_{1:T}, \mathbf{a}, \sigma_v^2, \sigma_w^2, \mathbf{r}_{1:T})$. However, evaluating this $T \times T$ covariance matrix is very computationally intensive. Fortunately, it is not necessary to evaluate it in order to sample $\mathbf{x}_{1:T}$. By using a state-space representation of the data, it is possible to use the efficient forward filtering/backward sampling procedure to sample from this distribution, see [15]. An alternative more computationally efficient method is the simulation smoother [20]. The complexity of these procedures is of $O(T)$.

The resulting Gibbs sampler algorithm is somewhat elegant. Nevertheless, it can be highly inefficient. Indeed, in most practical applications the mixture model for v_t has a component with a very small variance, i.e., $\alpha \ll 1$ [14], [22], [45]. In this case, there is a very strong correlation between r_t and v_t . Consequently the resulting MCMC converges very slowly towards its invariant distribution [22]. When $\alpha = 0$, i.e., v_t is a Bernoulli-Gaussian (BG) process and $p(dv_t \mid r_t = 0) = \delta_0(dv_t)$, then the Markov chain simulated by the Gibbs sampler is not even irreducible !

An Alternative Gibbs Sampler

We present here an alternative Gibbs sampler that circumvents the problems just described. It samples from $p(\theta, \mathbf{r}_{1:T}, \mathbf{x}_{1:T} | \mathbf{y}_{1:T})$. Contrary to what happens in the previous algorithm, the latent Bernoulli process $\mathbf{r}_{1:T}$ is not sampled conditionally on $\mathbf{x}_{1:T}$. This so-called partial sampler proceeds as follows.

A Partial Sampler for Blind Deconvolution

1. Initialization, $i = 0$: set randomly or deterministically $(\theta^{(0)}, \mathbf{r}_{1:T}^{(0)}, \mathbf{x}_{1:T}^{(0)})$.
2. Iteration i , $i \geq 1$.
 - Sample $\theta^{(i)} \sim p(\theta | \mathbf{y}_{1:T}, \mathbf{r}_{1:T}^{(i-1)}, \mathbf{x}_{1:T}^{(i-1)})$.
 - For $t = 1$ to T
 - Sample $r_t^{(i)} \sim p(r_t | \mathbf{y}_{1:T}, \theta^{(i)}, \mathbf{r}_{-t}^{(i)})$ where

$$\mathbf{r}_{-t}^{(i)} \triangleq (r_1^{(i)}, \dots, r_{t-1}^{(i)}, r_{t+1}^{(i-1)}, \dots, r_T^{(i-1)}).$$
 - End For.
 - Sample $\mathbf{x}_{1:T}^{(i)} \sim p(\mathbf{x}_{1:T} | \mathbf{y}_{1:T}, \theta^{(i)}, \mathbf{r}_{1:T}^{(i)})$.

The only modification of this algorithm with respect to the previous one is the sampling step of the latent process $\mathbf{r}_{1:T}$. We sample the indicator variables r_t one-at-a-time, the continuous state $\mathbf{x}_{1:T}$ being integrated out. Sampling from $p(r_t | \mathbf{y}_{1:T}, \theta, \mathbf{r}_{-t})$ can be done in $O(T)$ iterations using a clever backward/forward recursion introduced in [16], see [26] for an extension of this recursion to more general state-space models.

Remark 7.10. *One can check that this partial sampler has $p(\theta, \mathbf{r}_{1:T}, \mathbf{x}_{1:T} | \mathbf{y}_{1:T})$ as invariant distribution using Remark 7.8.*

Remark 7.11. *One could also sample the parameters $\theta \sim p(\theta | \mathbf{y}_{1:T}, \mathbf{r}_{1:T})$ using an MH step, the parameters $\mathbf{x}_{1:T}$ being integrated out. Nevertheless, it is not easy to get a “good” proposal for θ , so we have chosen here to sample them conditional on $\mathbf{x}_{1:T}$.*

Application to Bernoulli-Gauss Processes

We address the problem of blind deconvolution of a Bernoulli-Gaussian sequence, *i.e.* $\alpha = 0$. $T = 500$ data are simulated with the following parameters: $\mathbf{a} = (-1.51, 0.55)^T$, $\sigma_v = 0.3$, $\lambda = 0.05$ and $\sigma_w = 0.25$. This is a good

model of a neutron sensor [22]. The hyperparameters $(\Sigma_0, \nu_0, \gamma_0, \alpha_0, \beta_0) = (100\mathbf{I}_2, 0.01, 0.01, 0.01, 0.01)$.

We are interested in obtaining a point estimate of the filtered point process $\mathbf{v}_{1:T}$. The MMSE estimate is usually of no practical interest in such problems as the posterior distribution $p(\mathbf{v}_{1:T} | \mathbf{y}_{1:T})$ is multimodal [22]. We propose to perform instead a two-step procedure. In the first step, we run the proposed partial sampler algorithm to estimate the posterior distribution $p(\theta | \mathbf{y}_{1:T})$. We deduce the MMAP pointwise estimates of the parameters θ . Then, these parameters now assumed known, we perform a classical detection/estimation procedure [45] to estimate the sequence $\mathbf{r}_{1:T}$ and the dynamic noise $\mathbf{v}_{1:T}$. This second step is easily implemented using a simulated annealing version of the efficient method to sample the latent process $\mathbf{r}_{1:T}$; see [26] for details.

In Figure 7.2 we present the observation and real filtered point process. In Figure 7.3 we present the real and estimated point process.




FIGURE 7.2. Signal x_t and observations y_t (dotted lines).




FIGURE 7.3. Simulated sequence dynamic noise (solid line, circle o) and reconstructed sequence (dotted line, cross x)

7.5.3 Robust Bayesian Spectral Analysis

In this subsection, we address the problem of parameter estimation of sinusoids in non-Gaussian noise. This problem is of great interest in many fields, including seismology, nuclear magnetic resonance and radar. Several algorithms have been proposed in the literature to obtain the maximum likelihood frequency estimator in the case of a white Gaussian observation noise. However, in practice, there are numerous applications where

this noise is non-Gaussian. We develop a Bayesian model and an MCMC sampler to solve this problem. The problem of joint Bayesian detection and estimation of sinusoids in white Gaussian noise as well as the problem of estimation of sinusoids in presence of clipped data are addressed elsewhere [6, 7].

Problem Formulation and Objectives

Let $\mathbf{y}_{0:T-1}$ be an observed vector of T real data samples where

$$\mathbf{y}_t = \sum_{j=1}^k (a_{c_j} \cos[\omega_j t] + a_{s_j} \sin[\omega_j t]) + n_t \quad (7.126)$$

with $\omega_{j_1} \neq \omega_{j_2}$ for $j_1 \neq j_2$. a_{c_j} , a_{s_j} , ω_j are respectively the amplitudes and the radial frequency of the j^{th} sinusoid. In vector-matrix form, we have

$$\mathbf{y}_{0:T-1} = \mathbf{D}(\omega) \mathbf{a} + \mathbf{n}_{0:T-1}, \quad (7.127)$$

where $\mathbf{a} \triangleq (a_{c_1}, a_{s_1}, \dots, a_{c_k}, a_{s_k})^T$, $\omega \triangleq (\omega_1, \dots, \omega_k)^T$ and the $T \times 2k$ matrix $\mathbf{D}(\omega)$ is defined as

$$\begin{aligned} [\mathbf{D}(\omega)]_{i+1, 2j-1} &= \cos[\omega_j t], \quad (t = 0, \dots, T-1, j = 1, \dots, k) \\ [\mathbf{D}(\omega)]_{i+1, 2j} &= \sin[\omega_j t], \quad (t = 0, \dots, T-1, j = 1, \dots, k). \end{aligned} \quad (7.128)$$

The elements of $\mathbf{y}_{0:T-1}$ correspond to the superposition of k ($k \geq 1$) sinusoids corrupted by a non-Gaussian I.I.D. noise n_t . This noise is modeled as a two-component Gaussian mixture

$$n_t \stackrel{i.i.d.}{\sim} \lambda \mathcal{N}(0, \sigma^2) + (1 - \lambda) \mathcal{N}(0, \alpha^2 \sigma^2) \quad (7.129)$$

where $\lambda \in (0, 1)$ and $\alpha^2 \in (0, 1)$. Similarly to the blind deconvolution problem, it is convenient from an algorithmical point of view to introduce the latent Bernoulli process $r_t \in \{0, 1\}$ so that $\Pr(r_t = 1) = \lambda$ and

$$n_t | (r_t = 1) \sim \mathcal{N}(0, \sigma^2), \quad n_t | (r_t = 0) \sim \mathcal{N}(0, \alpha^2 \sigma^2). \quad (7.130)$$

The parameter of the sinusoids (\mathbf{a}, ω) and the parameters of the noise $(\lambda, \sigma^2, \alpha^2)$ are assumed to be unknown and random with the following prior distribution, see [3] for a detailed justification

$$\begin{aligned} p(\mathbf{a}, \omega, \lambda, \sigma^2, \alpha^2, \mathbf{r}_{0:T-1}) = \\ p(\mathbf{a}, \omega | \sigma^2, \alpha^2, \mathbf{r}_{0:T-1}) p(\sigma^2) p(\alpha^2) p(\mathbf{r}_{0:T-1} | \lambda) p(\lambda), \end{aligned} \quad (7.131)$$

where

$$p(\mathbf{a}, \omega | \sigma^2, \mathbf{r}_{0:T-1}) = \frac{1}{[2\pi\sigma^2\boldsymbol{\Sigma}]^{1/2}} \exp\left[-\frac{\mathbf{a}^T \boldsymbol{\Sigma}^{-1} \mathbf{a}}{2\sigma^2}\right] \frac{1}{\pi^k} \mathbb{I}_{(0, \pi)^k}(\omega) \quad (7.132)$$

with Σ defined as

$$\Sigma^{-1} \triangleq \delta^{-2} \mathbf{D}^T(\omega) \Delta^{-1} \mathbf{D}(\omega), \quad \Delta \triangleq \text{diag}(i_0, \dots, i_{T-1}) \quad (7.133)$$

where $i_t = \mathbb{I}_{\{r_t=1\}} + \alpha^2 \mathbb{I}_{\{r_t=0\}}$. The other prior distributions are equal to

$$\sigma^2 \sim \mathcal{IG}\left(\frac{\nu_0}{2}, \frac{\gamma_0}{2}\right), \quad \alpha^2 \sim \mathcal{U}_{(0,1)}, \quad \lambda \sim \mathcal{U}_{(0,1)}. \quad (7.134)$$

Given the observations $\mathbf{y}_{0:T-1}$, our aim is to estimate the posterior distribution $p(\mathbf{a}, \omega, \lambda, \sigma^2, \alpha^2, \mathbf{r}_{0:T-1} | \mathbf{y}_{0:T-1})$. We propose an MCMC sampler to perform Bayesian computation. This MCMC sampler uses the fact that one can evaluate the marginal posterior distribution $p(\omega, \lambda, \alpha^2, \mathbf{r}_{0:T-1} | \mathbf{y}_{0:T-1})$ analytically

$$p(\omega, \lambda, \alpha^2, \mathbf{r}_{0:T-1} | \mathbf{y}_{0:T-1}) = \left[\gamma_0 + \mathbf{y}_{0:T-1}^T \mathbf{P}(\omega, \Sigma) \mathbf{y}_{0:T-1} \right]^{-\frac{\tau+\nu_0}{2}} \frac{1}{\pi^k} \mathbb{I}_{(0,\pi)^k}(\omega) \quad (7.135)$$

with

$$\mathbf{P}(\omega, \Sigma) = \Sigma^{-1} - \Sigma^{-1} \mathbf{D}(\omega) [\mathbf{D}^T(\omega) \Sigma^{-1} \mathbf{D}(\omega)]^{-1} \mathbf{D}^T(\omega) \Sigma^{-1}. \quad (7.136)$$

Bayesian Computation Using MCMC

In this subsection we first present the main steps of the MCMC algorithm, give some details about these different steps and finally present simulation results.

MCMC Algorithm for Spectral Analysis in Impulsive Noise

1. Initialization, $i = 0$: set $\theta^{(0)} = \left\{ \mathbf{a}^{(0)}, \omega^{(0)}, \mathbf{r}_{0:T-1}^{(0)}, \sigma^{2(0)}, \lambda^{(0)}, \alpha^{2(0)} \right\}$ and $i = 1$.
2. Iteration $i, i \geq 1$.
 - For $j = 1, \dots, k$
 - Perform an MH step with $p\left(\omega_{j,k} | \mathbf{y}_{0:T-1}, \mathbf{r}_{0:T-1}^{(i-1)}, \alpha^{2(i-1)}, \omega_{-j}^{(i)}\right)$ as invariant distribution.
 - End_For
 - where $\omega_{-j}^{(i)} \triangleq \left(\omega_1^{(i)}, \dots, \omega_{j-1}^{(i)}, \omega_{j+1}^{(i-1)}, \dots, \omega_k^{(i-1)}\right)$.
 - For $t = 1, \dots, T$
 - Sample $r_t^{(i)} \sim p\left(r_t | \mathbf{y}_{0:T-1}, \alpha^{2(i-1)}, \omega^{(i)}, \mathbf{r}_{-t}^{(i)}, \lambda^{(i-1)}\right)$.
 - End_For
 - where $\mathbf{r}_{-t}^{(i)} \triangleq \left(r_0^{(i)}, \dots, r_{t-1}^{(i)}, r_{t+1}^{(i-1)}, \dots, r_{T-1}^{(i-1)}\right)$.
 - Sample $(\mathbf{a}^{(i)}, \sigma^{2(i)}) \sim p\left(\mathbf{a}, \sigma^2 | \mathbf{y}_{0:T-1}, \mathbf{r}_{0:T-1}^{(i)}, \omega^{(i)}, \alpha^{2(i)}\right)$.

- Sample

$$\left(\alpha^{2(i)}, \lambda^{(i)}\right) \sim p\left(\alpha^2 \mid \mathbf{y}_{0:T-1}, \mathbf{r}_{0:T-1}^{(i)}, \omega^{(i)}, \mathbf{a}^{(i)}, \sigma^{2(i)}\right) p\left(\lambda \mid \mathbf{r}_{0:T-1}^{(i)}\right). \quad (7.137)$$

The details of the different sampling steps are described in [3], [25], however we outline here the specifics of these steps:

- sampling from ω_j relies on a mixture of two MH transition kernels which aim at exploring rapidly the posterior distribution of interest and the modes of this distribution,
- sampling from r_t is standard as it is a discrete probability distribution,
- sampling from \mathbf{a}, σ^2 is standard as it amounts to drawing σ^2 from an inverted gamma distribution and \mathbf{a} from a normal distribution.
- sampling from α^2, λ is also easy as it amounts to draw λ from a beta distribution and α^2 from a truncated inverted gamma distribution; and

Simulations

We have applied the algorithm to synthetic data, whose parameters were $T = 128$, $k = 2$, $\lambda = 0.02$ and $\alpha^2 = 0.01$. The parameters for the two sinusoids are given in Table 7.1 ($E_i \triangleq a_{c_i}^2 + a_{s_i}^2$).

i	E_i	$-\arctan(a_{s_i}/a_{c_i})$	$\omega_i/2\pi$
1	20	0	0.2
2	18	$\pi/4$	0.3

TABLE 7.1. Parameters of the two sinusoids for robust spectral analysis.

We applied our algorithm to the data and ran it for 20,000 iterations. We present in Figure 7.4 the periodogram of the observed data. Note that the second frequency is not visible on the periodogram. In Figure 7.5 we present the time series (truncated to facilitate visualization), the estimates of the probabilities $p(r_t = 1 \mid \mathbf{y}_{0:T-1})$ (for $t = 0, \dots, 127$) and the true sequence of r_t such that $r_t = 1$. In Figure 7.6 we present the estimates of $p(\omega_1/2\pi \mid \mathbf{y}_{0:T-1})$ and $p(\omega_2/2\pi \mid \mathbf{y}_{0:T-1})$. In Figure 7.7 we present an estimate of the posterior distribution of λ , the parameter of the Gaussian mixture, whose true value is 0.02.

7.6 Conclusion

In this chapter, our aim was to provide a simple tutorial review of Monte Carlo methods and their applications to Bayesian data analysis. These




FIGURE 7.4. Periodogram of the data

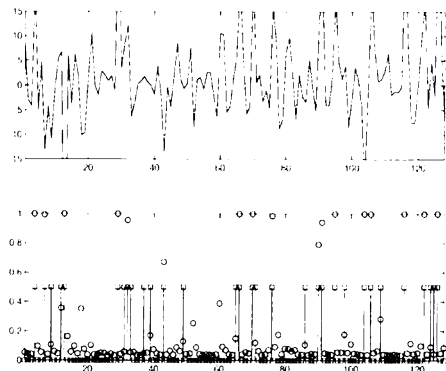


FIGURE 7.5. Top: time series (threshold). Bottom: probability $p(\xi_i = 1 | y)$ (circles) and the original sequence (squares at height 0.5).




FIGURE 7.6. Estimates of the posterior distribution of the two frequencies.




FIGURE 7.7. Estimation of the posterior distribution of λ .

methods provide powerful statistical tools as they allow one to tackle complex non-linear and non-Gaussian real world problems. However, it must be understood that these methods are not free from drawbacks:

- Simulation-based methods are not black boxes and they require some degree of expertise in order to work well in practice.
- They are computationally intensive and cannot be used in many on-line applications of interest.
- Many theoretical questions remain open, in particular the questions related to practical convergence assessments, even if some progress has been made in this direction (see [51] and references therein).

Despite these drawbacks it is the belief of the authors that the revolution that has occurred over the last few years, where MCMC methods have been applied to many problems that could not have been previously addressed is now undergoing a further revolution, sequential Monte Carlo methods.

Indeed, whereas standard MCMC algorithms have been extensively applied in “batch mode” to data problems, only little attention has been paid to real time applications typically met in data analysis or signal processing, and the next few years will show how this can be achieved efficiently by the use of new algorithms and faster computers.

Acknowledgments

C. Andrieu is financially supported by AT&T Laboratories, Cambridge U.K. A. Doucet is supported by EPSRC, U.K.

Notation

- \mathbf{A} : matrix.
- $[\mathbf{A}]_{i,j}$: i^{th} line, j^{th} column of matrix \mathbf{A} .
- \mathbf{A}^T : transpose matrix.
- \mathbf{A}^{-1} : inverse.
- $|\mathbf{A}|$: determinant.
- $\mathbf{z}_{i:j} \triangleq (z_i, \dots, z_j)^T$.
- $\mathbf{z}_{i:j,k:l}$:

$$\mathbf{z}_{i:j,k:l} \triangleq \begin{pmatrix} z_{i,k} & \cdots & z_{i,l} \\ \vdots & & \vdots \\ z_{j,k} & \cdots & z_{j,l} \end{pmatrix}$$

- If $\mathbf{z} \triangleq (z_1, \dots, z_{j-1}, z_j, z_{j+1}, \dots, z_k)^T$ then $\mathbf{z}_{-j} \triangleq (z_1, \dots, z_{j-1}, z_{j+1}, \dots, z_k)^T$.
- $\mathbf{0}_{n \times p}$: null matrix of dimensions $n \times p$.
- \mathbf{I}_n : identity matrix of dimensions $n \times n$.
- $\mathbb{I}_E(\mathbf{z})$: indicator function of the set E (1 if $\mathbf{z} \in E$, 0 else).
- $\lfloor z \rfloor$: highest integer strictly less than z .
- $\mathbf{z} \sim p(\mathbf{z})$: \mathbf{z} is distributed according to $p(\mathbf{z})$.
- $\mathbf{z} | \mathbf{y} \sim p(\mathbf{z})$: \mathbf{z} is conditional upon \mathbf{y} distributed according to $p(\mathbf{z})$.

Name	\mathcal{F}	$f_{\mathcal{F}}(\cdot)$
Inv.-Gamma	$\mathcal{IG}(\alpha, \beta)$	$\frac{\beta^\alpha}{\Gamma(\alpha)} z^{-\alpha-1} \exp(-\beta/z) \mathbb{I}_{[0,+\infty)}(z)$
Gaussian	$\mathcal{N}(\mathbf{m}, \Sigma)$	$ 2\pi\Sigma ^{-1/2} \exp\left(-\frac{1}{2}(\mathbf{z} - \mathbf{m})^T \Sigma^{-1}(\mathbf{z} - \mathbf{m})\right)$
Uniform	\mathcal{U}_A	$\left[\int_A d\mathbf{z}\right]^{-1} \mathbb{I}_A(\mathbf{z})$
Poisson	$\mathcal{P}(\lambda)$	$\exp(-\lambda) \frac{\lambda^z}{z!} \mathbb{I}_{\mathbb{N}}(z)$
Gamma	$\mathcal{G}(\alpha, \beta)$	$\frac{\beta^\alpha}{\Gamma(\alpha)} z^{\alpha-1} \exp(-\beta z) \mathbb{I}_{[0,+\infty)}(z)$

- $\ln(\cdot)$ is the natural logarithm.
- $\log_a(\cdot)$ is the base a logarithm.
- $\delta_{\mathbf{x}}(d\cdot)$ is the Dirac delta measure.

References

- [1] **The MCMC preprint service provides papers in the MCMC field:** <http://www.mcs.surrey.ac.uk/Personal/S.Brooks/MCMC/>
- [2] **The SMC Web server provides some links and papers in the field of Sequential Monte Carlo methods and particle filters:** <http://svr-www.eng.cam.ac.uk/~jfgf/smc.html>.
- [3] C. Andrieu, *Méthodes MCMC pour l'analyse bayésienne de modèles de régression paramétrique non-linéaires. Applications à l'analyse de raies et à la déconvolution impulsionnelle* (in French), Ph.D. Thesis, Université de Cergy-Pontoise, 1998.
- [4] C. Andrieu and A. Doucet, "Simulated Annealing for Bayesian Estimation of Hidden Markov Models", to appear *IEEE Trans. Information Theory*, 2000.
- [5] C. Andrieu and A. Doucet, "Efficient Stochastic MAP Estimation for Harmonic Retrieval", in *Proc. Conf. EUSIPCO*, Sept. 1998.
- [6] C. Andrieu and A. Doucet, "Joint Bayesian Model Selection and Estimation of Noisy Sinusoids via Reversible Jump MCMC", *IEEE Trans. Signal Processing*, vol. 47, no. 10, 2667-2676, 1999.
- [7] C. Andrieu and A. Doucet, "A Bayesian Approach to Harmonic Retrieval with Clipped Data", *Signal Processing*, vol. 74, no. 3, 1999.
- [8] C. Andrieu, L. Breyer and A. Doucet, "Convergence of Simulated Annealing using Foster-Lyapunov Criteria", technical report CUED-F TR 346, 1999.
- [9] J.M. Bernardo and A.F.M. Smith, *Bayesian Theory*, John Wiley & Sons, 1995.
- [10] J.O. Berger and T. Sellke, "Testing a Point Null Hypothesis: the Irreconcilability of P-values and Evidence", *J. Am. Stat. Assoc.*, 82, 112-122, 1987.
- [11] J.O. Berger and L.R. Pericchi, "The Intrinsic Bayes Factor for Model Selection and Prediction", *J. Am. Stat. Assoc.*, 91, 109-122, 1996.
- [12] J. Besag, P.J. Green, D.Higdon and K. Mengersen, "Bayesian Computation and Stochastic Systems", *Stat. Science*, vol. 10, 3-66, 1995.
- [13] S. Brooks, "Markov Chain Monte Carlo Method and its Application", *The Statistician*, vol. 47, 69-100, 1998.
- [14] O. Cappé, A. Doucet, E. Moulines and M. Lavielle, "Simulation-based Methods for Blind Maximum-likelihood Filter Identification", *Signal Processing*, vol. 73, no. 1, 3-25, 1999.
- [15] C.K. Carter and R. Kohn, "On Gibbs Sampling for State Space Models", *Biometrika*, vol. 81, no. 3, 541-553, 1994.
- [16] C. Carter and R. Kohn, "Markov Chain Monte Carlo Methods in Conditionally Gaussian State Space Models", *Biometrika*, vol. 83, 589-601, 1996.
- [17] R. Chen and T. Li, "Blind Restoration of Linearly Degraded Discrete Signals by Gibbs Sampling", *IEEE Trans. Signal Processing*, vol. 43, no. 10, 2410-2413, 1995.
- [18] Q. Cheng, R. Chen and T. Li, "Simultaneous Wavelet Estimation and Deconvolution of Reflection Seismic Signals", *IEEE Trans. Geo. Rem. Sens.*, vol. 34, 377-384, 1996.
- [19] S. Chib and E. Greenberg, "Markov Chain Monte Carlo Simulation Methods in Econometrics", *Econometric Theory*, 12, 409-431, 1996.
- [20] P. DeJong and N. Shephard, "The Simulation Smoother for Time Series Models", *Biometrika*, vol. 82, no. 2, 339-350, 1995.

- [21] L. Devroye, *Non-Uniform Random Variate Generation*, Springer, New-York, 1986.
- [22] A. Doucet and P. Duvaut, "Bayesian Estimation of State Space Models Applied to Deconvolution of Bernoulli-Gaussian Processes", *Signal Processing*, vol. 57, 147-161, 1997.
- [23] A. Doucet, *Algorithmes Monte Carlo pour l'estimation bayésienne de modèles markoviens cachés. Applications aux signaux de rayonnements* (in French), Ph.D. Thesis, Université de Paris-Sud Orsay, 1997.
- [24] A. Doucet, S.J. Godsill and C. Andrieu, "On Sequential Monte Carlo Sampling Methods for Bayesian Filtering", to appear *Statistics and Computing*, 2000. (journal paper version of "On Sequential Simulation-Based Methods for Bayesian Filtering", technical report CUED-F TR 310, Feb. 1998. Available at [http://www-sigproc.eng.cam.ac.uk/~ad2/tr310.ps](http://www.sigproc.eng.cam.ac.uk/~ad2/tr310.ps).)
- [25] A. Doucet and C. Andrieu, "Robust Bayesian Spectral Analysis via MCMC Sampling", *Proc. Conf. EUSIPCO*, Sept. 1998.
- [26] A. Doucet and C. Andrieu, "Iterative Algorithms for Optimal State Estimation of Jump Markov Linear Systems", in *Proc. IEEE ICASSP'99*.
- [27] A. Doucet, J.F.G. de Freitas and N.J. Gordon (editors), *Sequential Monte Carlo Methods in Practice*, Springer-Verlag, 2000.
- [28] A. Doucet, A. Logothetis and V. Krishnamurthy, "Stochastic sampling algorithms for state estimation of jump Markov linear systems", to appear *IEEE Trans. Automatic Control*, 2000.
- [29] A.E. Gelfand and A.F.M. Smith, "Sampling-based Approaches to Calculating Marginal Densities", *J. Am. Statist. Assoc.*, vol. 85, no. 410, 398-409, 1990.
- [30] A. Gelman, G.O. Roberts and W.R. Gilks, "Efficient Metropolis Jumping Rules", in *Bayesian Statistics V*, Clarendon Press, Oxford, 599-608, 1996.
- [31] A. Gelman and D.B. Rubin, "Markov Chain Monte Carlo Methods in Biostatistics", *Stat. Meth. Med. Res.*, 339-355, 1996.
- [32] S. Geman and D. Geman, "Stochastic Relaxation, Gibbs Distribution and the Bayesian Restoration of Images", *IEEE Trans. Patt. Ana. Mac. Int.*, vol. 6, 721-741, 1994.
- [33] J. Geweke, "Bayesian Inference in Econometrics Models using Monte Carlo Integration", *Econometrica*, vol. 57, 1317-1339, 1989.
- [34] W.R. Gilks, S. Richardson and D.J. Spiegelhalter (editors), *Markov Chain Monte Carlo in Practice*, Chapman&Hall, 1996.
- [35] N.J. Gordon, D.J. Salmond and A.F.M. Smith, "Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation", *IEE-Proceedings-F*, vol. 140, no. 2, 107-113, 1993.
- [36] P.J. Green, "Reversible Jump Markov Chain Monte Carlo Computation and Bayesian Model Determination", *Biometrika*, vol. 82, no. 4, 711-732, 1995.
- [37] H. Haario and E. Sacksman, "Simulated annealing in general state space", *Adv. Appl. Prob.*, vol. 23, 866-893, 1991.
- [38] W.K. Hastings, "Monte Carlo Sampling Methods using Markov Chains and their Applications", *Biometrika* 57, 97-109, 1970.
- [39] A.H. Jazwinski, *Stochastic Processes and Filtering Theory*, Academic Press, 1970.
- [40] R.E. Kass and A.E. Raftery, "Bayes Factors", *J. Am. Stat. Assoc.*, 90, 773-796, 1995.

- [41] R.E. Kass and L. Wasserman, "The Selection of Prior Distributions by Formal Rules", *J. Am. Stat. Assoc.*, 91, 1343-1369, 1996.
- [42] J.S. Liu and R. Chen, "Blind Deconvolution via Sequential Imputation", *J. Am. Stat. Assoc.*, vol. 90, no. 430, 567-576, 1995.
- [43] J.S. Liu, "Metropolized Independent Sampling with Comparison to Rejection Sampling and Importance Sampling", *Statistics and Computing*, vol. 6, 113-119, 1996.
- [44] N. Metropolis, N. Rosenblutt, A.W. Rosenblutt, M.N. Teller, A.H. Teller, "Equations of State Calculations by Fast Computing Machines", *Journal of Chemical Physics*, 21, 1087-1092, 1953.
- [45] J.M. Mendel, *Maximum-Likelihood Deconvolution: A Journey into Model-Based Signal Processing*, New York: Springer-Verlag, 1990.
- [46] S.P. Meyn and R.L. Tweedie, *Markov Chains and Stochastic Stability*, Springer-Verlag, 1993.
- [47] J.J.K. O'Ruanaidh and W.J. Fitzgerald, *Numerical Bayesian Methods Applied to Signal Processing*, Springer-Verlag, 1996.
- [48] S. Richardson and P.J. Green, "On Bayesian Analysis of Mixtures with Unknown Number of Components", *J. Roy. Stat. Soc. B*, vol. 59, no. 4, 731-792, 1997.
- [49] B.D. Ripley, *Stochastic Simulation*, Wiley, New York, 1987.
- [50] C.P. Robert, *The Bayesian Choice*, Springer-Verlag, 1996.
- [51] C.P. Robert (editor), *Discretization and MCMC Convergence Assessment* (1998), Lecture Notes 135, Springer-Verlag, New-York.
- [52] C.P. Robert and G. Casella, *Monte Carlo Statistical Methods*, Springer-Verlag, 1999.
- [53] C.P. Robert, A. Doucet and S.J. Godsill, "Maximum A Posteriori Estimation using MCMC", in *Proc. IEEE ICASSP'99*. Available on http://www-sigproc.eng.cam.ac.uk/~ad2/icassp99_same.pdf.
- [54] A.F.M. Smith and A.E. Gelfand, "Bayesian Statistics without Tears: a Sampling-Resampling Perspective", *American Statistician*, vol. 46, no. 2, 84-88, 1992.
- [55] A.F.M. Smith and G.O. Roberts, "Bayesian Computation via the Gibbs sampler and Related Markov Chain Monte Carlo Methods", *J. Roy. Stat. Soc. B*, vol. 55, 3-23, 1993.
- [56] D.J. Spiegelhalter and A.F.M. Smith, "Bayes Factor for Linear and for Log-linear Models with Vague Prior Information", *J. Roy. Stat. Soc. B*, 44, 377-387, 1982.
- [57] M.A. Tanner, *Tools for statistical inference: methods for the exploration of posterior distributions and likelihood functions*, Springer-Verlag, New York, 1993.
- [58] L. Tierney, "Markov Chains for Exploring Posterior Distributions", *The Annals of Statistics*, vol. 22, 1701-1762, 1994.
- [59] P.J. Van Laarhoven and E.H.L. Arts, *Simulated Annealing: Theory and Applications*, Reidel Pub., Amsterdam, 1987.
- [60] M. West and J.F. Harrison, *Bayesian Forecasting and Dynamic Models*, Springer Verlag Series in Statistics, 2nd edition, 1997.

Chapter 8

Constrained Randomization of Time Series for Nonlinearity Tests

Thomas Schreiber¹
Andreas Schmitz

ABSTRACT We discuss the problem of generating time sequences that fulfill given constraints but are random otherwise. This is an important ingredient for generalized nonlinearity tests that use Monte Carlo resampling. We briefly discuss standard methods available for a limited range of problems. Then we put forth a novel scheme in which one can define arbitrary sets of observables and test if these observables give a complete account of the serial correlation structure in the data. The most immediate application is the detection of correlations beyond the two-point autocovariance, even in a non-Gaussian setting. More general constraints, also including multivariate, nonlinear, and nonstationary properties, can be implemented in the form of a cost function to be minimized.

8.1 Introduction

In the statistical evaluation of empirical observations, randomization of data is an important means to obtain probability distributions and confidence intervals in cases where assumptions on the normality of an estimated quantity cannot be made. If the data consist of a time series exhibiting serial correlations, unconditional randomization is not usually desired. Typically, the serial correlations are to some extent explained by the two-point function, while the quantity of interest is meant to probe nonlinear structure. In cases like this, it is desirable to condition the randomization procedure on a given set of observables. For example, in a nonlinearity test, we would like to perform *constrained randomization* that preserves the second-order autocorrelation function in order to isolate and assess the significance of nonlinear correlations. In the dynamical systems community, Monte Carlo resampling is known as the method of *surrogate data*, a term introduced by Theiler et al. [18]. Theiler and Prichard [19] give a readable introduction to

¹Author for correspondence.

the concept of *constrained realizations* of a process, as opposed to *typical realizations* generated by a bootstrap method [2].

The background and motivation for this work lie in the need to justify *by the data* the use of advanced nonlinear time series methods. Many authors have applied methods derived from nonlinear dynamics or even the paradigm of deterministic chaos to signals of supposedly nonlinear origin. However, no matter how certain we may be that a given system constitutes a nonlinear device, take the human brain as an example, nonlinear structure may not be evident in a particular variable recorded with some finite quality. Thus, the supposed nonlinearity in the brain is not in itself a justification for the analysis of electro-encephalographic (EEG) recordings with nonlinear techniques.

Although the proper choice of observable will be essential for the amount of information that can be extracted from a given data set, this is only a side issue here, and it will be discussed only briefly. We will focus on methods for the generation of surrogate data, starting with Fourier-based approaches. After pointing out their merits and limitations, we will proceed to the main topic of this chapter and introduce a general constrained randomization scheme. Its use will be illustrated by a number of examples. Finally, we will discuss the utility of this approach for practical work.

Most of the code to generate surrogate data and measure nonlinearity mentioned in this chapter has been implemented as part of the TISEAN free software package [6]. A review paper containing many technical details on the generation of surrogate data will be published elsewhere [17].

8.2 Probing Nonlinearity

Several quantities have been discussed that can be used to characterize nonlinear time series. For the purpose of nonlinearity testing, we need such quantities that are particularly powerful in discriminating linear dynamics and weakly nonlinear signatures. Traditional measures of nonlinearity are derived from generalizations of the two-point autocovariance function or the power spectrum. One particularly useful third-order quantity is $\sum_{n=\tau+1}^N (s_n - s_{n-\tau})^3$ since it measures the asymmetry of a series under time reversal. (We have dropped all normalization here since the values will only be used in relative comparisons.) When a nonlinearity test is performed with the question in mind if nonlinear deterministic modeling of the signal may be useful, it seems most appropriate to use a test statistic related to a nonlinear deterministic approach [7]. Widely used are test statistics which in some way quantify the nonlinear predictability of the signal. Let $\vec{x}_n = (s_{n-(m-1)\tau}, \dots, s_n)$ be the sequence of time delay embedding vectors obtained from the scalar time series $\{s_n\}$. The nonlinear

prediction error can then be defined (again dropping normalization) by

$$e = \sqrt{\sum [\vec{x}_{n+1} - F(\vec{x}_n)]^2}. \quad (8.1)$$

The prediction by F over one time step can, in the simplest case, be performed by averaging over the future values of all neighboring delay vectors $\vec{x}_{n'}$ closer to \vec{x}_n than ϵ in m embedding dimensions.

In Ref [16], several quantities are compared quantitatively with respect to their power to detect small deviations from linear behavior. Choices that have been made in the literature include concepts like *false nearest neighbors* [8], measures for the continuity of a dynamical system [12], generalized redundancies [10], symbolic descriptions [5], and a fair number of coarse-grained versions of dimensions, entropies, and Lyapunov exponents.

Almost all measures of nonlinearity have in common that their probability distribution on finite data sets is not known analytically. In fact, there are many examples of nonlinearity measures that aren't even approximately normal. It has therefore been advocated since the early days [18] to use robust statistics rather than parametric methods for the actual statistical test. In other words, we discourage the common practice of representing the distribution of the nonlinearity measure by an error bar and deriving the significance from the number of "sigmas" the data lies outside these bounds. Such reasoning implicitly assumes a Gaussian distribution.

Instead, we follow Theiler et al. [18] by using a rank-order test. First, we select a residual probability α of a false rejection, corresponding to a level of significance $(1 - \alpha) \times 100\%$. Then, for a one-sided test (e.g., looking for *small* prediction errors only), we generate $1/\alpha - 1$ surrogate sequences. Thus, including the data itself, we have $1/\alpha$ sets. Therefore, the probability that the data by coincidence has, say, the smallest prediction error is exactly α , as desired. For a two-sided test (e.g., for time asymmetry, which can go both ways), we would generate $2/\alpha - 1$ surrogates, resulting in a probability α that the data give *either* the smallest *or* the largest value.

8.3 Generating Constrained Realizations

Traditional bootstrap methods [2] use explicit model equations that have to be extracted from the data. This *typical realizations* approach can be very powerful for the computation of confidence intervals, provided the model equations can be extracted successfully. As discussed by Theiler and Prichard [19], the alternative approach of *constrained realizations* is more suitable for the hypothesis testing we are interested in here. It avoids the fitting of model equations by directly imposing the desired structures onto the randomized time series. However, the choice of possible null hypotheses is limited by the difficulty of imposing arbitrary structures on otherwise random sequences.

It is essential for the validity of the statistical test that the surrogate series are created properly. If they contain spurious differences to the measured data, these may be detected by the test and interpreted as signatures of nonlinearity. A simple case is the null hypothesis that the data consist of independent draws from a fixed probability distribution. Surrogate time series can be obtained by randomly shuffling the measured data. If we find significantly different serial correlations in the data and the shuffles, we can reject the hypothesis of independence.

8.3.1 *Fourier-Based Methods*

A step toward more interesting null hypotheses is to incorporate the structures reflected by linear two-point autocorrelations. A corresponding null hypothesis is that the data have been generated by some linear stochastic process with Gaussian increments. The statistical test is complicated by the fact that we don't want to test against one particular linear process only (one specific choice of ARMA coefficients), but against a whole class of processes. This is called a *composite* null hypothesis. The unknown coefficients are sometimes referred to as *nuisance parameters*. There are basically three directions we can take in this situation. First, we could try to make the discriminating statistic independent of the nuisance parameters. This approach has not been demonstrated to be viable for any but very simple statistics. Second, we could determine which linear model is most likely realized in the data by a fit for the coefficients, and then test against the hypothesis that the data has been generated by this particular model. Surrogates are simply created by running the fitted model. The main drawback is that we cannot recover the *true* underlying process by any fit procedure.

The null hypothesis of an underlying Gaussian linear stochastic process can also be formulated by stating that all structure to be found in a time series is exhausted by computing first- and second-order quantities, the mean, the variance, and the autocovariance function. (There is a one-to-one correspondence between the coefficients of an ARMA model, the autocovariance function, and the power spectrum.) This means that a randomized sample can be obtained by creating sequences with the same second-order properties as the measured data but are otherwise random. When the linear properties are specified by the squared amplitudes of the Fourier transform (that is, the periodogram estimator of the power spectrum), surrogate time series are readily created by multiplying the (complex) Fourier coefficients of the data by random phases — preserving their moduli — and then transforming back to the time domain. The resulting randomized series are referred to as *phase-randomized* surrogates.

The most obvious deviation from the Gaussian linear process is usually that the data don't follow a Gaussian distribution. There is a simple generalization of the null hypothesis that explains deviations from the normal

distribution by the action of an invertible, static measurement function: $s_n = s(x_n)$ where $\{x_n\}$ is a realization of an ARMA process. We want to regard a time series from such a process as essentially linear since the only nonlinearity is contained in the — in principle invertible — measurement function $s(\cdot)$.

The most commonly used method to create surrogate data sets for this null hypothesis essentially attempts to invert $s(\cdot)$ by rescaling the time series $\{s_n\}$ to conform with a Gaussian distribution. The rescaled version is then phase-randomized (conserving Gaussianity on average), and the result is rescaled to the empirical distribution of $\{s_n\}$. This procedure is referred to as the *amplitude-adjusted Fourier transform* (AAFT) method. In Ref. [15] we argue that for short and strongly correlated sequences this algorithm can yield an incorrect test due to a bias toward a flat spectrum. There, we propose a method which iteratively corrects deviations in spectrum and distribution. In an alternating fashion, the surrogate is filtered toward the correct Fourier amplitudes and rank-ordered to the correct distribution. The accuracy that can be reached depends on the size and structure of the data and is generally more than sufficient for hypothesis testing.

An Example: Southern Oscillation Index

As an illustration, let us perform a statistical test for nonlinearity on a monthly time series of the Southern Oscillation Index (SOI) from 1866 to 1994 (1560 samples). For a reference on analysis of Southern Oscillation data see Graham et al. [3, 4]. Since a discussion of this climatic phenomenon is not relevant to the issue at hand, let us just consider the time series as an isolated data item. Our null hypothesis is that the data are adequately described by its single time probability distribution and its power spectrum. This corresponds to the assumption that an autoregressive moving average (ARMA) process is generating a sequence measured through a static invertible, possibly nonlinear observation function.

For a test at the 99% level of significance ($\alpha = 0.01$), we generate a collection of $1/\alpha - 1 = 99$ surrogate time series, which share the single time sample probability distribution and the periodogram estimator with the data. This is carried out using the iterative method described by Schreiber and Schmitz [15]. Figure 8.1 shows the data with one of the ninety-nine surrogates.

As a discriminating statistic we use a locally constant predictor in embedding space, using three-dimensional delay coordinates at a delay time of one month. Neighborhoods were selected at 0.2 times the rms amplitude of the data. The test is set up in such a way that the null hypothesis may be rejected when the prediction error is smaller for the data than for all of the ninety-nine surrogates. But, as we can see in Figure. 8.2, this is not the case. Predictability is not significantly reduced by destroying




FIGURE 8.1. Monthly values of the Southern Oscillation Index (SOI) from 1866 to 1994 (upper trace) and a surrogate time series exhibiting the same auto-covariance function (lower trace). All linear properties of the fluctuations and oscillations are the same between both tracings. However, any possible nonlinear structure except for a static rescaling of the data is destroyed in the lower tracing by the randomization procedure.

possible nonlinear structure. This negative result can mean several things. The prediction error statistics may just not have any power to detect the kind of nonlinearity present. Alternatively, the underlying process may be linear and the null hypothesis true. It could also be, and this seems the most likely option after all we know about the equations governing climate phenomena, that the process is nonlinear but the single time series at this sampling covers such a poor fraction of the rich dynamics that it must appear linear stochastic to the analysis.

Of course, our test has been carried out disregarding any knowledge of the SOI situation. It is very likely that more informed measures of nonlinearity may be more successful in detecting structure. We would like to point out, however, that if such information is derived from the same data, or literature published on it, a bias is likely to occur. Similarly to the situation of multiple tests on the same sample, the level of significance has to be adjusted properly. Otherwise, if many people try, someone will eventually, and maybe accidentally, find a measure that indicates nonlinear structure.

8.3.2 A General Randomization Scheme

Randomization schemes based on the Fourier amplitudes of the data are appropriate in many cases. However, there remain some flaws, the strongest




FIGURE 8.2. Nonlinear prediction error measured for the SOI data set (see Fig. 8.1) and 99 surrogates. The value for the original data is plotted with a longer impulse. It is evident that the data is not singled out by this property and we are unable to reject the null hypothesis of a linear stochastic stationary process, possibly rescaled by a nonlinear measurement function.

being the severely restricted class of testable null hypotheses. The periodogram estimator of the power spectrum is about the only interesting observable that allows for the solution of the inverse problem of generating random sequences under the condition of its given value.

In the general approach of Ref. [14], constraints (e.g., autocorrelations) on the surrogate data are implemented by a cost function $E(\{s_n\})$ which has a global minimum when the constraints are fulfilled. Any set of constraints that can be written in the form

$$F_i(\{s_n\}) = 0, \quad i = 1, \dots, I \quad (8.2)$$

can be turned into a cost function

$$E(\{s_n\}) = \left(\sum_{i=1}^I |w_i F_i(\{s_n\})|^q \right)^{1/q}. \quad (8.3)$$

The fact that $E(\{s_n\})$ has a global minimum when the constraints are fulfilled is unaffected by the choice of the weights w_i and the order q of the average. The least squares or L^2 average is obtained at $q = 2$, L^1 at $q = 1$, and the maximum distance at $q = \infty$. Geometric averaging is also possible. We have experimented with different choices of q but we haven't found a choice that is uniformly superior to others. It seems plausible to either give uniform weights or to enhance those constraints which are particularly difficult to fulfill. Again, conclusive empirical results are still lacking.

Carrying out the minimization numerically is complicated by the omnipresence of false minima — a typical situation for minimization tasks derived from root-finding problems. One common method to deal with local minima is the method of simulated annealing [20]. Starting with a random permutation of the original time series, the surrogate is modified by

exchanging two points chosen at random. The modification will be accepted if it yields a lower value for the cost function or with a probability $p = \exp(-\Delta E/T)$. The “system temperature” T will be lowered slowly to let the system settle down to a minimum. With an appropriate cooling scheme, the annealing procedure can reach any desired accuracy.

Of course, the data itself constitute a formal solution of the minimization procedure. However, except for very short sequences or extremely over-specified constraints, it is very unlikely to ever arrive at this absolute minimum. If it does become a problem, closeness to the original data and trivial transformations of it can be explicitly penalized by extra terms in the cost function.

The particular constraint that the autocovariances of the surrogate $C'(\tau)$ should be the same as those of the data $C(\tau)$ can be realized by specifying the discrepancy as a cost function, for example,

$$E = \sum_{\tau=0}^{N-1} |C'(\tau) - C(\tau)|. \quad (8.4)$$

$E(\{\tilde{s}_n\})$ is minimized among all permutations $\{\tilde{s}_n\}$ of the original time series $\{s_n\}$. Even for the null hypothesis underlying the Fourier-based methods (AAFT and iterative), the cost function approach is superior in two respects. First, any desired accuracy can be specified and reached at the expense of computer time. Second, the definition of linear correlations is not restricted to the periodogram spectral estimator. The most dangerous caveat using Fourier transforms here is that they silently assume that the data is one full cycle of a periodic function. If that were the case, the null hypothesis would be inadequate right away.

Constrained randomization using combinatorial minimization is a very flexible method since, in principle, it can realize arbitrary constraints of the form $F(\{s_n\}) = 0$. It can be quite useful to be able to incorporate into the surrogates any feature of the data that is understood already or that is considered uninteresting. The price for the accuracy and generality of the method is its high computational cost. Rather than dwelling more on theoretical aspects, let us give a few instructive examples of the use of annealed surrogates. The superiority of the method in terms of higher accuracy for the case of a standard null hypothesis has been demonstrated in Ref. [14], where multivariate and nonstationary examples are also presented.

The SOI data discussed earlier are rather well behaved with little end-to-end mismatch and small deviation from Gaussianity. Therefore, the iterative Fourier-based scheme is completely satisfactory, and we don't have to redo the analysis with the more time-consuming annealing method.

Example: Unevenly Sampled Data

Let us show how the constrained randomization method can be used to test for nonlinearity in time series taken at time intervals of different length.

Unevenly sampled data are quite common; examples include drill core data, astronomical observations, and stock price notations. Most observables and algorithms cannot easily be generalized to this case, which is particularly true for nonlinear time series methods. (See [11] for material on irregularly sampled time series.) Interpolating the data to equally spaced sampling times is not recommendable for a test for nonlinearity since one could not *a posteriori* distinguish between genuine structure and nonlinearity introduced spuriously by the interpolation process.

For data that are evenly sampled except for a moderate number of gaps, surrogate sequences can be simply produced by assuming the value zero during the gaps and minimizing the standard cost function Eq.(8.4) while excluding the gaps from the permutations tried. The error made in estimating correlations would then be identical for the data and surrogates and could not affect the validity of the test. For data sampled at incommensurate times, such a strategy can no longer be adopted. We then need different means to specify the linear correlation structure.

Consider a time series sampled at times $\{t_n\}$ that need not be equally spaced. The power spectrum can then be estimated by the Lomb periodogram $P(\omega)$, as discussed, for example, in Ref. [13]. Here we give the final formula:

$$P(\omega) = \frac{1}{2\sigma^2} \left\{ \frac{[\sum_n (y_n - \bar{y}) \sin \omega(t_n - \tau)]^2}{\sum_n \sin^2 \omega(t_n - \tau)} + \frac{[\sum_n (y_n - \bar{y}) \cos \omega(t_n - \tau)]^2}{\sum_n \cos^2 \omega(t_n - \tau)} \right\} \quad (8.5)$$

where τ is defined by $\tan(2\omega\tau) = \sum_n \sin 2\omega t_n / \sum_n \cos 2\omega t_n$ and \bar{y}, σ^2 are the mean and variance of the data, respectively. The result can be derived by fitting a least squares model $y_n = a \cos \omega t_n + b \sin \omega t_n$ to the data for each given frequency ω . Therefore, Lomb periodograms are often referred to as *least squares periodograms*.

For time series sampled at constant time intervals, the Lomb periodogram yields the standard squared Fourier transformation. Except for this particular case, it does not have any inverse transformation, which makes it impossible to use the standard surrogate data algorithms mentioned in Section. 8.3.1. Therefore, we use the Lomb periodogram of the data as a constraint for the surrogates. It can be expressed as a cost function, for example, by: $E = \sum_{k=1}^{N m_f} |P'(k\omega_0) - P(k\omega_0)|$. We use P at N_f equally spaced frequencies $k\omega_0$; other choices are possible. Consider a series [1] of the time-integrated intensity of light observed from a variable star; see Figure. 8.3. It consists of seventeen parts with different numbers of points, the time range of which may overlap or show gaps of up to 10,000 seconds. Between gaps, the (down-sampled) data are evenly sampled with $\Delta = 120$ seconds, the total number of points is $N = 2260$. Since the length of the gaps is not a multiple of the sampling rate, we cannot use a simple estimator of autocorrelation. We use the Lomb periodogram (8.5) instead. The linear null hypothesis was not rejected by the time reversibility statistic. One




FIGURE 8.3. The down-sampled data set **E** with one corresponding surrogate. Gaps of different sizes prevents reasonable interpolation.

surrogate is shown in Figure. 8.3.

Imposing a given Lomb periodogram is very time-consuming because at each annealing step, the $O(N)$ spectral estimator has to be computed at $O(N_f)$ frequencies with $N_f \propto N$. We are currently developing a scheme that uses binned autocorrelations instead. By limiting the range of lags involved, this is expected to lead to a dramatic reduction in computation time. Let us finally note that the very common situation that the data are given by event times, or interevent intervals, is different from the case studied here in that the sampling instances $\{t_n\}$ are not independent of the measured values. In fact, between these instances, the value of $x(t)$ is undefined. Work in progress is concerned with methods for generating surrogates in this case.

Example: Higher-Order Correlations

As an example of a more exotic cost function, let us show the randomization of 500 iterates of the Hénon map, Figure. 8.4 (a)). Panel (b) shows a Fourier-based (iterated) surrogate data set having the same spectrum and distribution. Starting from a random permutation (Panel (c)), the cost function

$$\begin{aligned}
 C &= \langle x_{n-1}x_n \rangle + \langle x_{n-2}x_n \rangle \\
 &+ \langle x_{n-1}^2x_n \rangle + \langle x_{n-1}x_n^2 \rangle + \langle x_{n-2}^2x_n \rangle + \langle x_{n-2}x_{n-1}x_n \rangle \\
 &+ \langle x_{n-1}^2x_n^2 \rangle + \langle x_{n-1}x_n^3 \rangle + \langle x_{n-1}^3x_n \rangle
 \end{aligned} \tag{8.6}$$

is minimized. It involves all the higher-order autocorrelations which would be needed for a least squares fit with the ansatz $x_n = c - ax_{n-1}^2 + bx_{n-2}$

and in this sense fully specifies the quadratic structure of the data. In other words, we are testing the hypothesis that the Hénon data can be faithfully represented by the preceding ansatz, but without ever fitting it explicitly. Of course, this is a highly contrived example that serves the sole purpose of demonstrating the flexibility of the approach.

Figure. 8.4 shows the effect of minimizing the cost function Eq.(8.6) by simulated annealing. The random shuffle (Panel (c)) yields $C = 2400$, panels (d)-(f) correspond to $C = 150, 15, \text{ and } 0.002$, respectively. While the accuracy is increased by six orders of magnitude, the full fractal structure is not yet recovered. Small- scale structure is not represented very efficiently by the higher-order autocorrelations. This example may serve as an illustration of why higher order correlations and poly-spectral methods are not very popular with researchers with a deterministic chaos background. While they formally give a full account of the nonlinearity in a time series, they are not particularly adequate to capture the characteristic structure emerging in the limiting case of pure determinism. Conversely, one can argue that this limiting case itself is not very relevant for field research where noise sources abound.

8.4 Remaining Caveats

Strictly speaking, the concept of constrained realizations requires the constraints to be fulfilled *exactly*, a practical impossibility. Most of the research efforts reported in this chapter have their origin in the attempt to increase the accuracy with which the constraints are implemented, that is, to minimize the bias resulting from any remaining discrepancy. Since most measures of nonlinearity are also sensitive to linear correlations, a side effect of the reduced bias is a reduced variance of such estimators. Paradoxically, the enhanced accuracy may result in false rejections of the null hypothesis on the ground of tiny differences in some nonlinear characteristics. This important point was recently put forth by Kugiumtzis [9]. He suggests testing the validity of the surrogate sample by performing a test using a linear statistic for normalization. Currently, this seems to be the only way around the problem, and we thus recommend following his suggestion. It may be possible, however, to generate unbiased ensembles of surrogates by specifying a cost function that explicitly minimizes the bias, involving the whole collection of surrogates at the same time. In any event, this will be a very cumbersome procedure, in terms of implementation and execution speed.

The most severe caveat, in our point of view, remains the danger of spurious interpretation of a test result. Often in the literature, rejection of a rather simple null hypothesis is taken as an indication for the validity of a particular alternative. For example, in the case of the rescaled linear

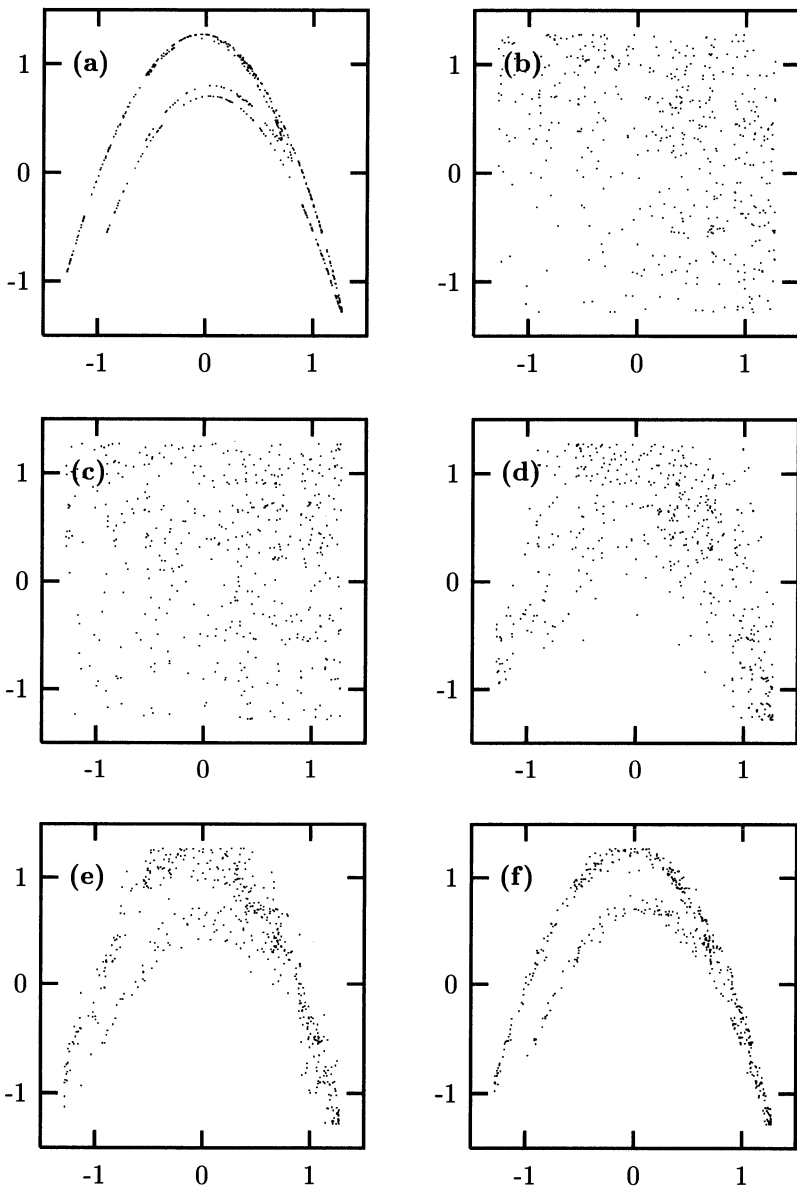


FIGURE 8.4. Randomization of 500 points generated by the Hénon map. (a) Original data; (b) Same autocorrelations and distribution; (c)-(f) Different stages of annealing with a cost function C involving three and four-point correlations. (c) A random shuffle, $C = 2400$; (d) $C = 150$; (e) $C = 15$; (f) $C = 0.002$. See text.

stationary process, any other static nonlinearity, for example, a nonlinear moving average, can lead to a rejection as well as nonstationarity or inadequate sampling. There is no immediate justification to conclude that nonlinear *dynamics* is acting in the system.

8.5 Discussion

We have set up a statistical hypothesis test of nonlinearity in a general sense. How interesting its outcome is depends on the specific null hypothesis chosen. The most meaningful test can be performed if the null hypothesis is plausible enough that we are prepared to believe it in the lack of evidence against it. In general, we will specify a set of observables we believe to be complete to describe the structure found in the data. The surrogates will then share these properties with the data and any significant discrepancy between data and surrogates can guide to a more complete understanding.

Recent efforts on the generalization of randomization schemes try to broaden the repertoire of null hypotheses we can test against. The hope is that we can eventually choose one that is general enough to be acceptable if we fail to reject it with the methods we have. Still, we cannot prove that there isn't any structure in the data beyond what is covered by the null hypothesis. From a practical point of view, however, there is not much of a difference between structure that isn't there and structure that is undetectable with our observational means.

We thank Dimitris Kugiumtzis and Holger Kantz for useful discussions. This work was supported by the SFB 237 of the Deutsche Forschungsgemeinschaft.

References

- [1] J. Ch. Clemens. Whole earth telescope observation of the white dwarf star PG 1159-035 (data set E). In A. S. Weigend and N. A. Gershenfeld, editors, *Time Series Prediction: Forecasting the Future and Understanding the Past*, pages 139–150. Addison Wesley, Reading, MA, 1993.
- [2] B. Efron. The jackknife, the bootstrap, and other resampling plans. *SIAM, Philadelphia*, 1982.
- [3] N. E. Graham, J. Michaelsen, and T. P. Barnett. An investigation of the El Nino-Southern Oscillation cycle with statistical models - 1. Predictor field characteristics. *J. Geophys. Res.*, 92:1425–1427, 1987.
- [4] N. E. Graham, J. Michaelsen, and T. P. Barnett. An investigation of the El Nino-Southern Oscillation cycle with statistical models - 2. Model results. *J. Geophys. Res.*, 92:1427–1428, 1987.
- [5] B.-L. Hao. *Experimental Study and Characterization of Chaos*. World Scientific, Singapore, 1990.

- [6] R. Hegger, H. Kantz, and T. Schreiber. Practical implementation of nonlinear time series methods: The TISEAN package. *Chaos*, 9:413, 1999.
- [7] H. Kantz and T. Schreiber. *Nonlinear Time Series Analysis*. Cambridge Univ. Press, Cambridge, UK, 1997.
- [8] M. B. Kennel and S. Isabelle. Method to distinguish chaos from colored noise and to determine embedding parameters. *Phys. Rev. A*, 46:3111–3118, 1992.
- [9] D. Kugiumtzis. Test your surrogate data before you test for nonlinearity. *Phys. Rev. E*, 60(3):2808–2816, 1999.
- [10] M. Palus. Nonlinearity in normal human EEG: Cycles, temporal asymmetry, non-stationarity and randomness, not chaos. *Biol. Cybern.*, 75(5):389–396, 1996.
- [11] E. Parzen, editor. *Time Series Analysis of Irregularly Observed Data*, volume 25 of *Lecture Notes in Statistics*. Springer-Verlag, Berlin, Heidelberg, New York, 1983.
- [12] L. M. Pecora, T. L. Carroll, and J. F. Heagy. Statistics for continuity and differentiability: An application to attractor reconstruction from time series. In C. D. Cutler and D. T. Kaplan, editors, *Nonlinear Dynamics and Time Series*, volume 11 of *Fields Inst. Communications*, pages 49–62. American Math. Soc., Providence, Rhode Island, 1997.
- [13] W. H. Press, S. T. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes*. Cambridge University Press, Cambridge, England, 1992.
- [14] T. Schreiber. Constrained randomization of time series data. *Phys. Rev. Lett.*, 90(10):2105–2108, 1998.
- [15] T. Schreiber and A. Schmitz. Improved surrogate data for nonlinearity tests. *Phys. Rev. Lett.*, 77(4):635–638, 1996.
- [16] T. Schreiber and A. Schmitz. Discrimination power of measures for nonlinearity in a time series. *Phys. Rev. E*, 55(5):5443–5447, 1997.
- [17] T. Schreiber and A. Schmitz. Surrogate time series. *Physica D (accepted for publication)*, 2000.
- [18] J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, and J. D. Farmer. Testing for nonlinearity in time series: The method of surrogate data. *Physica D*, 58:77–94, 1992.
- [19] J. Theiler and D. Prichard. Constrained-realization monte-carlo method for hypothesis testing. *Physica D*, 94:221, 1996.
- [20] R. V. V. Vidal, editor. *Applied Simulated Annealing*, volume 396 of *Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, Berlin, Heidelberg, New York, 1993.

Chapter 9

Removing the Noise from Chaos Plus Noise

Steven P. Lalley

ABSTRACT The problem of extracting a “signal” x_n generated by a dynamical system from a time series $y_n = x_n + e_n$, where e_n is an observational error, is considered. It is shown that consistent signal extraction is impossible when the errors are distributed according to a density with unbounded support, and the underlying dynamical system admits homoclinic pairs. It is also shown that consistent signal extraction is possible when the errors are uniformly bounded by a suitable constant and the underlying dynamical system has the “weak orbit separation property”. Simple algorithms for signal recovery are described in the latter case.

9.1 Introduction

Is it possible to consistently recover a “signal” $\{x_n\}_{n \in \mathbb{Z}}$ generated by a chaotic dynamical system from a time series of the form

$$y_n = x_n + e_n \tag{9.1}$$

where e_n is observational noise? This is the *noise removal*, or *signal separation* problem, and it has been discussed in a number of papers, including [4, 5, 8, 2, 1]. Various sophisticated methods for noise removal have been proposed, nearly all requiring a degree of smoothness in the underlying dynamical system, and some requiring rather detailed a priori knowledge of the dynamics. The issue of convergence seems not to have been broached until now. The purpose of this paper is to state some general results concerning the theoretical possibility of *consistent filtering* and to propose some fairly simple general-purpose filters for use in high-signal/noise ratio problems.

In many circumstances, *scalar* measurements will be made on a dynamical system at equally spaced times to produce the series x_n , which is then observed with error. We shall assume here, however, that x_n is the actual *state vector* of the system at time n . This assumption is probably harmless, in view of the embedding theorem [9]. Moreover, we shall assume throughout that the *noise* e_n consists of I.I.D. mean zero random vectors that are independent of the state vectors x_n . Although we shall make only weak

assumptions about the dynamics, we shall limit our attention to dynamical systems with compact invariant sets. Compactness will be essential in Theorems 9.2 and 9.3.

Definition 9.1. A dynamical system is a homeomorphism $F : \Lambda \rightarrow \Lambda$ of a compact subset Λ of a Euclidean space \mathbb{R}^d . For any point $x \in \Lambda$, the orbit of x is the doubly infinite sequence $\{x_n = F^n(x)\}_{n \in \mathbb{Z}}$, where F^n denotes the n -fold composition of F .

Definition 9.2. A filter \hat{x} is a collection of functions

$$\hat{x}_n(y_0, y_1, y_2, \dots, y_m) = \hat{x}_n^{(m)}(y_0, y_1, y_2, \dots, y_m).$$

A filter \hat{x} is weakly consistent if for every orbit $\{x_n\}_{n \in \mathbb{Z}}$ and every $\varepsilon > 0$,

$$\frac{1}{m} \sum_{n=1}^m |\hat{x}_n - x_n| \xrightarrow{P} 0. \tag{9.2}$$

If μ is an F -invariant probability measure on Λ , then the filter \hat{x} is μ -weakly consistent if (9.2) holds for μ -almost every orbit $\{x_n\}$.

Note that if a filter is weakly consistent, then it is μ -weakly consistent for any F -invariant probability measure μ . Informally, a filter is weakly consistent if, for large m , most of the fitted values \hat{x}_n are close to the corresponding state vectors x_n . Other notions of consistency are undoubtedly worthy of consideration. In certain situations one might regard the requirement in (9.2) that only “most” points on the orbit be well approximated as too weak. For a stronger notion of consistency, see [6], theorem 2.

9.2 Homoclinic Pairs

A common and important dynamical feature of many chaotic systems is the occurrence (or even abundance) of *homoclinic pairs*. Two distinct points x, x' are said to be *homoclinic* if their orbits $\{x_n\}_{n \in \mathbb{Z}}$ and $\{x'_n\}_{n \in \mathbb{Z}}$ satisfy

$$\lim_{n \rightarrow \pm\infty} |x_n - x'_n| = 0. \tag{9.3}$$

In smooth systems, it is commonly (but not always) the case that if the convergence (9.3) occurs then it is exponentially fast. Say that two homoclinic points x, x' are *strongly homoclinic* if their orbits satisfy

$$\sum_{n=-\infty}^{\infty} |x_n - x'_n| < \infty. \tag{9.4}$$

In uniformly hyperbolic systems, most points are members of strongly homoclinic pairs; if the stable and unstable manifolds through x intersect at

x' , then (x, x') is a strongly homoclinic pair. In systems admitting “symbolic dynamics” (that is, systems conjugate [or nearly conjugate] to subshifts of finite type), *all* points will be members of homoclinic pairs. This class includes all mixing *Axiom A* diffeomorphisms; see [6].

The occurrence of strongly homoclinic pairs is a fundamental obstruction to the existence of consistent filters. For any error density ϕ on \mathbb{R}^d , say that ϕ is in the class Φ if it is strictly positive, has mean zero, and satisfies

$$\limsup_{y \rightarrow 0} \frac{1}{|y|} \int \left| \log \frac{\phi(x+y)}{\phi(x)} \right| \phi(x) dx < \infty. \tag{9.5}$$

Note that all mean zero Gaussian densities are of class Φ .

Proposition 9.1. *Let x, x' be a strongly homoclinic pair, with F -orbits $\{x_n\}_{n \in \mathbb{Z}}$ and $\{x'_n\}_{n \in \mathbb{Z}}$, respectively. Assume that the noise density is of class Φ . Then there is no sequence of (measurable) functions $\xi_n(y_{-n}, y_{-n+1}, \dots, y_n)$ such that*

$$\begin{aligned} \xi_n(x_{-n} + e_{-n}, x_{-n+1} + e_{-n+1}, \dots, x_n + e_n) &\xrightarrow{P} x \text{ and} \\ \xi_n(x'_{-n} + e_{-n}, x'_{-n+1} + e_{-n+1}, \dots, x'_n + e_n) &\xrightarrow{P} x' \end{aligned} \tag{9.6}$$

Proof Sketch. Define probability measures Q, Q' on the sequence space $(\mathbb{R}^d)^{\mathbb{Z}}$ to be the distributions of the doubly infinite sequence $\{y_n\}_{n \in \mathbb{Z}}$ when y_n is defined by

$$y_n = x_n + e_n \quad (Q), \tag{9.7}$$

$$y_n = x'_n + e_n \quad (Q'), \tag{9.8}$$

with the random vectors e_n I.I.D. from a density ϕ in class Φ . Then the measures Q, Q' are mutually absolutely continuous, because (9.4) and the assumption that $\phi \in \Phi$ guarantees the almost-sure convergence of the infinite product

$$\frac{dQ}{dQ'} = \prod_{n=-\infty}^{\infty} \frac{\phi(y_n - x_n)}{\phi(y_n - x'_n)} \tag{9.9}$$

to a strictly positive limit. But if Q and Q' are mutually a.c., then there can be no sequence of functions $\xi_m(y_{-m}, \dots, y_m)$ such that as $m \rightarrow \infty$,

$$\xi_m(y_{-m}, \dots, y_m) \xrightarrow{Q} x \quad \text{and} \quad \xi_m(y_{-m}, \dots, y_m) \xrightarrow{Q'} x'. \tag{9.10}$$

□

Although Proposition 9.1 does not by itself preclude the existence of weakly consistent filters, it indicates that consistent orbit identification is

impossible when there are strongly homoclinic pairs. Moreover, if homoclinic pairs are sufficiently common, as in the case of an Axiom A diffeomorphism, then there may not be weakly consistent filters.

Theorem 9.1. *Let μ be an ergodic, F -invariant probability measure μ on Λ such that on some probability space there are Λ -valued random variables X, X' satisfying the following (a) the marginal distributions of X and X' are both μ ; (b) X and X' are either equal or strongly homoclinic, with probability one; and (c) with positive probability, X and X' are strongly homoclinic. If the noise density is of class Φ , then there is no μ -weakly consistent filter.*

Proof: We may assume that the error random variables $\{e_n\}_{n \in \mathbb{Z}}$ are defined on the same probability space as X and X' . Set

$$\begin{aligned} X_n &= F^n(X) & Y_n &= X_n + e_n, \\ X'_n &= F^n(X') & Y'_n &= X'_n + e_n. \end{aligned}$$

Each of the sequences $\{X_n\}, \{X'_n\}, \{Y_n\}, \{Y'_n\}$ is stationary, and the random variables X_n, X'_n have distribution μ ; moreover, the random variables X_n and X'_n are bounded (since Λ is compact).

Suppose there were a μ -weakly consistent filter $\{\hat{x}_n^{(m)}\}$. Without loss of generality, we may assume that the functions $\hat{x}_n^{(m)}$ are bounded, since truncation does not affect the validity of (9.2). Hence,

$$\begin{aligned} \frac{1}{m} \sum_{n=1}^m |\hat{x}_n^{(m)}(Y_1, Y_2, \dots, Y_m) - X_n| &\xrightarrow{L^2} 0, & (9.11) \\ \frac{1}{m} \sum_{n=1}^m |\hat{x}_n^{(m)}(Y'_1, Y'_2, \dots, Y'_m) - X'_n| &\xrightarrow{L^2} 0. \end{aligned}$$

This implies, by the triangle inequality for the L^2 -norm and the stationarity of the sequences $\{X_n\}, \{Y_n\}$, that, for every n , the random variable X_n is in the closed L^2 -span of the random variables $\{Y_n\}_{n \in \mathbb{Z}}$. Therefore, there is a sequence of measurable functions $\xi^{(m)} = \xi^{(m)}(y_{-m}, y_{-m+1}, \dots, y_m)$ such that

$$X = L^2 - \lim \xi^{(m)}(Y_{-m}, Y_{-m+1}, \dots, Y_m) \quad \text{and} \quad (9.12)$$

$$X' = L^2 - \lim \xi^{(m)}(Y'_{-m}, Y'_{-m+1}, \dots, Y'_m), \quad (9.13)$$

the second convergence follows from the first because $(X', \{Y'_n\})$ has the same joint distribution as $(X, \{Y_n\})$. But this contradicts Proposition 9.1, since with positive probability the random variables X, X' form a homoclinic pair. \square

9.3 Sensitive Dependence on Initial Conditions

Definition 9.3. *The dynamical system $F : \Lambda \rightarrow \Lambda$ has sensitive dependence on initial conditions if there exists a constant $\Delta > 0$, called a separation threshold, such that for any two distinct points $x, x' \in \Lambda$, there exists $n \in \mathbb{Z}$ such that*

$$|F^n(x) - F^n(x')| > \Delta. \quad (9.14)$$

Dynamical systems with sensitive dependence on initial conditions often have homoclinic pairs; for example, topologically mixing Axiom A diffeomorphisms has both sensitive dependence and homoclinic pairs. For systems with sensitive dependence on initial conditions, consistent noise removal is possible if the noise level is sufficiently low. Consistent filters are easily described and implemented.

Smoothing Algorithm D: The algorithm takes as input a finite sequence $\{y_n\}_{0 \leq n \leq m}$ and produces as output a sequence $\{\hat{x}_n\}_{0 \leq n \leq m}$ of the same length that will approximate the unobservable signal $\{x_n\}_{0 \leq n \leq m}$. Let κ_m be an increasing sequence of integers such that

$$\lim_{m \rightarrow \infty} \kappa_m = \infty \quad \text{and} \quad \lim_{m \rightarrow \infty} \frac{\kappa_m}{\log m} = 0; \quad (9.15)$$

e.g., $\kappa_m = \log m / \log \log m$. For each integer $1 \leq n \leq m$, define A_n to be the set of indices $\nu \in \{0, 1, \dots, m\}$ such that

$$\max_{|j| \leq \kappa_m} |y_{\nu+j} - y_n| < 3\delta, \quad (9.16)$$

with the convention that $|y_j - y_i| = \infty$ if either i or j is not in the range $[0, m]$. Observe that $n \in A_n$, so A_n is nonempty; and for $n \leq \kappa_m$ or $n \geq m - \kappa_m$, the set A_n is the singleton $\{n\}$. In rough terms, A_n consists of the indices of those points in the time series whose orbits “shadow” the orbit of x_n for κ_m time units. Now define

$$\hat{x}_n = \frac{1}{|A_n|} \sum_{\nu \in A_n} y_\nu. \quad (9.17)$$

Theorem 9.2. *Suppose that the dynamical system $f : \Lambda \rightarrow \Lambda$ has sensitive dependence on initial conditions, with separation threshold Δ . If the errors e_n have mean zero and are uniformly bounded in absolute value by δ , where $\delta < \Delta/5$, then Smoothing Algorithm D is weakly consistent.*

This is a generalization of Theorem 1 of [6], which applies only to smooth, uniformly hyperbolic systems, where sensitive dependence on initial conditions can be “quantified”. Theorem 9.2 requires no smoothness of the underlying dynamical system at all. Furthermore, the hypotheses may be

relaxed in several ways: (1) It is not necessary that the errors e_n be identically distributed. If, for example, the distribution of e_n is allowed to depend on the state vector x_n , then weak consistency of Smoothing Algorithm D will still hold, *provided* that the errors are conditionally independent, given the orbit $\{x_n\}_{n \in \mathbb{Z}}$, that they are uniformly bounded by δ , and that $E(e_n | \{x_n\}_{n \in \mathbb{Z}}) = 0$. (2) It is not necessary even that the errors be mutually independent. If $\{e_n\}$ is a mean zero, stationary sequence satisfying suitable mixing requirements, then the conclusion of Theorem 9.2 remains valid.

Explanation of Theorem 9.2. A complete proof is given in [7]; here we shall give only a brief indication of the argument. Observe that the average (9.17) may be rewritten as

$$\hat{x}_n = \frac{1}{|A_n|} \sum_{\nu \in A_n} x_\nu + \frac{1}{|A_n|} \sum_{\nu \in A_n} e_\nu. \tag{9.18}$$

Thus, to establish weak consistency, it suffices to show that for most of the indices $n \in [1, m]$, (a) the cardinality of A_n is large, and (b) if $\nu \in A_n$ then $|x_n - x_\nu|$ is small. Property (b) will guarantee that the first average in (9.18) is close to x_n , while property (a), together with the law of large numbers, will imply (with some work!) that with high probability the second average is near zero.

Property (b) follows easily from sensitive dependence on initial conditions. This implies that if x, x' are any two points whose orbits remain within distance 5δ for all times $-\kappa < n < \kappa$, then $|x - x'|$ must be small, provided κ is large. Since the errors e_n are of magnitude less than δ , if $\nu \in A_n$ then, by the triangle inequality, the orbits of x_n and x_ν must remain within distance 5δ for all $j \in [-\kappa_m, \kappa_m]$. Thus, $\nu \in A_n$ implies that $|x_n - x_\nu|$ is small.

Property (a) follows from the assumption that $\kappa_m = o(\log m)$. Let H be a finite subset of Λ that is δ -dense in Λ , and denote by H^* the set of all H -valued sequences of length $2\kappa_m + 1$. For every F -orbit segment of length $2\kappa_m + 1$, there is at least one sequence in H^* that δ -shadows it. Since $\kappa_m = o(\log m)$, the cardinality of H^* is $o(m)$. Thus, by the pigeonhole principle, for most indices $n \in [1, m]$ there will be *many* indices ν such that

$$\max_{|j| \leq \kappa_m} |x_{n+j} - x_{\nu+j}| < 2\delta. \tag{9.19}$$

All such indices ν must be included in A_n . □

Note that this is not a complete proof, because the sets A_n are random, not fixed, and so the use of the law of large numbers is problematic.

9.4 Example: The Henon Mapping

Smoothing Algorithm D is easily implemented, and simple variations of the algorithm can be made to run in $O(m \log m)$ steps. In practical terms this means that, for simple low-dimensional systems, with $m = 10^5$, the procedure can be run in “real time” (e.g., 10 to 20 seconds on a 200-MHz Power Macintosh). This implies that experimentation with the parameters κ_m and δ may be done in real time. In simple examples, choosing δ to be one-fifth to one-tenth the apparent diameter of the attractor has been effective; for $m \approx 10^5$, choosing κ_m so that most bins A_n have 20 to 50 points has provided the best results.

Figures 9.1 to 9.2 show the results of using the filter for a noise-corrupted orbit of length 10^5 generated by the Henon mapping. The figures show (a) 10^5 points on the orbit of a randomly chosen point in the basin of attraction of the attractor Λ ; (b) the orbit corrupted by noise; and (c) the reconstructed orbit. The author is indebted to Jason Stover for coding the algorithm. Similar figures for a noise-corrupted orbit of Smale’s solenoid mapping may be found on the author’s Web page at

<http://galton.uchicago.edu/~lalley>




FIGURE 9.1. A partial orbit (10^5 points) of a randomly chosen point in the basin of attraction of the attractor of the Henon map.




FIGURE 9.2. The orbit of Fig.9.1 contaminated by noise.




FIGURE 9.3. The reconstructed orbit.

9.5 The Weak Orbit Separation Property

Although sensitive dependence on initial conditions is sometimes taken to be a necessary condition for chaos (see, e.g., [3], section 1.8), there are important systems for which sensitive dependence does not hold, which nonetheless share many of the dynamical features of chaotic systems. Noteworthy among these are the time-1 mappings induced by smooth hyperbolic flows. If $\phi_t : \Lambda \rightarrow \Lambda$ is a smooth flow, and if $F = \phi_1$, then sensitive dependence cannot hold, for an obvious and trivial reason: If two points x, x' are on the same flow line (that is, if $x' = \phi_s(x)$ for some $s \neq 0$), then their orbits (under F) remain on the same flow line, at (roughly) the same distance, forever. However, if the flow ϕ_t is hyperbolic (see [10] for the definition), then the orbits of all neighboring points *not on a common flow line* will eventually separate. Such systems satisfy a *weak orbit separation property*, defined as follows.

Definition 9.4. For any pair of points $x, x' \in \Lambda$ and any $\Delta > 0$, define

$$\tau_+^\Delta(x, x') = \min\{n \geq 0 : |F^n(x) - F^n(x')| > \Delta\}, \tag{9.20}$$

$$\tau_-^\Delta(x, x') = \max\{n \leq 0 : |F^n(x) - F^n(x')| > \Delta\}, \tag{9.21}$$

with the convention that $\tau_+^\Delta = \infty$ and/or $\tau_-^\Delta = -\infty$ if there are no such integers n . The dynamical system $F : \Lambda \rightarrow \Lambda$ has the weak orbit separation property if there exist constants $\Delta > 0$ (a separation threshold) and $\alpha > 0$ such that for any two distinct points $x, x' \in \Lambda$, the inequality

$$|F^n(x) - F^n(x')| > \alpha|x - x'| \tag{9.22}$$

holds for all integers n satisfying

$$0 \leq n \leq \tau_+^\Delta(x, x') \quad \text{if} \quad \tau_+^\Delta(x, x') < \infty; \tag{9.23}$$

$$0 \geq n \geq \tau_-^\Delta(x, x') \quad \text{if} \quad \tau_-^\Delta(x, x') > -\infty; \text{ and} \tag{9.24}$$

$$-\infty < n < \infty \quad \text{if} \quad \tau_+^\Delta(x, x') = -\tau_-^\Delta(x, x') = \infty. \tag{9.25}$$

Perhaps the simplest nontrivial dynamical systems satisfying the weak orbit separation property are the rotations R_α of the unit circle. It is trivial to verify that the weak orbit separation property holds, because

$$|R_\alpha^n x - R_\alpha^n y| = |x - y| \quad \forall n \in \mathbb{Z} \text{ and } \forall x, y. \tag{9.26}$$

The weak orbit separation property holds not only for highly rigid, non-chaotic systems such as rotations, but also for highly chaotic systems, such as topologically mixing, Axiom A diffeomorphisms restricted to their non-wandering sets. This is not difficult to check. There are other examples that arise naturally, for example, when an Axiom A system is weakly coupled

with an almost-periodic system. In particular, the product

$$S \times R : X \times \mathbb{T}^d \longrightarrow X \times \mathbb{T}^d$$

of an Axiom A diffeomorphism $S : X \rightarrow X$ restricted to its nonwandering set X with a system $R : \mathbb{T}^d \rightarrow \mathbb{T}^d$ that is bi-Lipschitz conjugate to an ergodic rotation of the d -torus \mathbb{T}^d satisfies the weak orbit separation property. Finally, the most important dynamical systems satisfying the weak orbit separation property are the time-1 mappings of smooth flows with compact, hyperbolic invariant sets Λ .

Consistent noise removal is possible for dynamical systems satisfying the weak orbit separation property, provided the noise level is sufficiently low. A consistent filter is easily described, although it is not as easily implemented as Smoothing Algorithm D.

Smoothing Algorithm W. The filter is defined by averaging, as in Smoothing Algorithm D, but the selection of indices over which to average is now done differently. Let κ_m be a sequence of integers satisfying the conditions (9.15), and, for each $1 \leq n \leq m$, let A_n again be the set of indices $\nu \in \{0, 1, \dots, m\}$ for which inequality (9.16) is satisfied. Define B_n to be the subset of A_n consisting of those $\lceil |A_n|/\log |A_n| \rceil$ indices ν for which the residual sums of squares

$$SS(\nu, n; m) = \sum_{|j| \leq \kappa_m} |y_{n+j} - y_{\nu+j}|^2 \tag{9.27}$$

are the smallest. Now define

$$\hat{x}_n = \frac{1}{|B_n|} \sum_{\nu \in B_n} y_\nu. \tag{9.28}$$

Theorem 9.3. *Suppose that the dynamical system $F : \Lambda \rightarrow \Lambda$ satisfies the weak orbit separation property, with separation threshold Δ . If the errors e_n have mean zero and are uniformly bounded by $\delta < \Delta/5$, then Smoothing Algorithm W is weakly consistent.*

Explanation of Theorem 9.3. A complete proof is given in [7]; what follows is the skeleton of the argument. For dynamical systems that satisfy the weak orbit separation property, orbits of nearby points need not diverge to a fixed distance Δ , and so neighboring points cannot be identified in the same simple manner as in the case of dynamical systems with sensitive dependence on initial conditions. In particular, it is no longer necessarily the case that $\nu \in A_n$ (that is, $|y_{n+j} - y_{\nu+j}| < 3\delta$ for all $|j| < \kappa_m$) will guarantee that $|x_n - x_\nu|$ is small, even for large m . However, the weak orbit separation property *does* imply that if two orbits fail to diverge in time κ_m

(either forward in time or backward in time), and if κ_m is sufficiently large, then

$$|x_{n+j} - x_{\nu+j}| > \alpha|x_n - x_\nu| \quad (9.29)$$

for either all $0 \leq j \leq \kappa_m$, all $-\kappa_m \leq j \leq 0$, or all $-\kappa_m \leq j \leq \kappa_m$.

Now consider $SS(n, \nu; m)$; since the random vectors e_n have mean zero, with high probability,

$$\begin{aligned} SS(n, \nu; m) &= \sum_{|j| \leq \kappa_m} |y_{n+j} - y_{\nu+j}|^2 \\ &\approx \sum_{|j| \leq \kappa_m} |x_{n+j} - x_{\nu+j}|^2 + \sum_{|j| \leq \kappa_m} |e_{n+j} - e_{\nu+j}|^2 \\ &\approx \sum_{|j| \leq \kappa_m} |x_{n+j} - x_{\nu+j}|^2 + 4\kappa_m E|e_0|^2. \end{aligned} \quad (9.30)$$

Thus, by (9.29), $SS(\nu, n; m)$ is, with high probability, considerably smaller for those indices ν such that $|x_n - x_\nu|$ is small. Selection of indices according to the values of $SS(\nu, n; m)$, will, therefore, tend to identify those ν such that x_ν is near x_n . Averaging over these indices will, with high probability, yield an estimate close to x_n , by an argument like that used in the proof of Theorem 9.2. \square

9.6 Concluding Remarks

(1) Much of the published work on the signal separation problem (and, indeed, most work on statistical inference for chaotic dynamical systems) makes no distinction between discrete-time systems and continuous-time systems. However, the preceding results suggest that there may, in fact, be a significant difference, at least for the signal separation problem. This is certainly the case for hyperbolic systems; discrete-time hyperbolic systems have the sensitive dependence property, but continuous-time systems do not — they satisfy only the weak orbit separation property.

(2) The case of *hyperbolic flows* deserves further attention. It may be shown that certain large classes of hyperbolic flows — including (a) mixing geodesic flows on compact, negatively curved manifolds and (b) ergodic suspensions of hyperbolic toral automorphisms — admit homoclinic pairs. However, it may also be shown that for such flows homoclinic pairs are *rare*, in the sense that the set of points x that belong to such pairs has SRB-measure 0. Thus, it may be possible to construct weakly consistent filters, or filters which, although not weakly consistent in the sense of Definition 9.2, nevertheless satisfy the consistency relation (9.2) for *almost every* orbit.

(3) Practical aspects of the signal separation problem have not been systematically studied. Various authors have investigated the efficacy of various filtering schemes for one or two low-dimensional systems, but no comparative studies have been made of the relative merits of these schemes. Perhaps somewhere an enterprising graduate student will find this a worthwhile project.

Acknowledgment

Thanks to Andrew Nobel for pointing out a serious error in the original statement of Theorem 1.

References

- [1] H. Abarbanel. *Analysis of observed chaotic data*. Springer-Verlag, 1996.
- [2] M. Davies. Noise reduction schemes for chaotic time series. *Phys. D*, 79:174–192, 1992.
- [3] R. Devaney. *An Introduction to Chaotic Dynamical Systems*. Benjamin and Cummings, 1986.
- [4] E. Kostelich and T. Schreiber. Noise reduction in chaotic time-series data: a survey of common methods. *Phys. Rev. E*, 48:1752–1763, 1993.
- [5] E. Kostelich and J. Yorke. The simplest dynamical system consistent with the data. *Physica D*, 41:183–196, 1990.
- [6] S. Lalley. Beneath the noise, chaos. *Annals of Statistics*, 27, 1999.
- [7] S. Lalley. More noise, more chaos. 1999.
- [8] T. Sauer. A noise reduction method for signals from nonlinear systems. *Physica D*, 58:193–201, 1992.
- [9] T. Sauer, J. Yorke, and M. Casdagli. Embedology. *J. Statistical Phys.*, 65:579–616, 1991.
- [10] S. Smale. Differentiable dynamical systems. *Bull. Amer. Math. Soc.*, 73:747–817, 1967.

Chapter 10

Embedding Theorems, Scaling Structures, and Determinism in Time Series

Colleen D. Cutler

ABSTRACT In this chapter we discuss specific definitions of deterministic and stochastic for stationary time series. Our main purpose in doing so is to create a convenient rigorous framework in which to examine the interplay between state-space reconstruction (embedding theorems), scaling or fractal structures (the Grassberger-Procaccia algorithm), and the predictability properties of time series. Thus the definitions in and of themselves are not as important as the clarity and precision they provide within the context. In spite of the various pitfalls and limitations involved, choosing and adhering to a specific appropriate definition of determinism provides a firm foundation for proving theorems and constructing examples in those areas of chaos theory and time series concerned with reconstruction of the underlying source (or generating mechanism) of a time series. In this chapter we provide some examples where our approach enables us to show that such reconstruction cannot be done.

10.1 Introduction

The past few decades have seen considerable attention paid to the phenomenon of erratic, apparently stochastic, time series arising from the sampling of low-dimensional deterministic systems—the phenomenon popularly known as *chaos*. To some extent chaos is surprising (perhaps intuitively we do not expect randomness, or even pseudo randomness, to arise from a deterministic rule with a few degrees of freedom) and this intuition is backed up by the knowledge that the long-run behavior of standard nonlinear systems cannot give rise to chaos. Fixed points and limit cycles, even quasi-periodic motions, cannot produce erratic random-looking time series. On the other hand, most of us are comfortable with the use of pseudo random number generators, the majority of which attempt to generate a sequence of pseudo independent random variables based on a few deterministic rules. In fact it is not too difficult to coax a deterministic map into producing a truly stochastic sequence by choosing the correct functional. Consider the

binary shift map on the unit interval

$$\eta(y) = \begin{cases} 2y & \text{for } 0 \leq y \leq 1/2 \\ 2(y - 1/2) & \text{for } 1/2 < y \leq 1 \end{cases} \quad (10.1)$$

and define the functional $h : [0, 1] \rightarrow \mathbb{R}$ by $h(y) = 0$ if $0 \leq y \leq 1/2$ and $h(y) = 1$ if $1/2 < y \leq 1$. Noting that the uniform distribution (that is, Lebesgue measure restricted to the unit interval) is invariant and ergodic for η , choose an initial condition Y randomly from the uniform distribution. Then the sampled time series $X_n = h(\eta^{n-1}(Y))$, $n = 1, 2, \dots$ obtained by iterating Y under η will consist of an independent and identically-distributed (*I.I.D.*) sequence of 0's and 1's. (It is not possible to actually carry out this iteration on the computer because we cannot store the initial condition to infinite precision, and finite decimals are attracted to the fixed point at 0 under the action of η .)

Of course η is a bounded nonlinear map. Moreover, it possesses the other necessary ingredient for chaos — it has a positive Lyapunov exponent and thus is sensitive to initial conditions (nearby points separate quickly under the action of η). So this is one part of the recipe for producing a stochastic time series. But the other part is the functional h itself, which in this case destroys the structural link between the deterministic source η and the resulting stochastic output X_n , $n = 1, 2, \dots$. It is not possible to work backward from the X_n series to η ; the source of the X_n output could just as easily be a sequence of random coin tosses. Hence we say that reconstruction of η is not possible from the sampled time series.

At this point we need to back up and clarify our terminology. The sequence of 0s and 1s produced by η and h is in fact *I.I.D.* and hence stochastic by anyone's definition of the term. But this is not the case with the majority of time series resulting from sampling chaotic systems like η . These time series may appear erratic and may possess many stochastic properties (like decaying autocorrelations and a broadband spectrum) but often sufficient structural link exists between the deterministic source and the observed time series that the source *can* be reconstructed from the time series. Demonstrating the existence of this structural link (in the form of time-delay embeddings) is the purpose behind both the theorem of Takens (1981) and the theorems of Sauer et al. (1991). We will argue later that a time series from which a finite-dimensional deterministic mechanism can be reconstructed (in a manner to be made precise later) should properly be viewed as deterministic, regardless of any apparent stochastic properties it might also possess.

Hence we believe it useful to have precise definitions of deterministic and stochastic in order to help us make distinctions between certain types of time series. It will also aid us in distinguishing, at least in some cases, between those sampled time series which permit system reconstruction and those which do not.

Our definitions of deterministic and stochastic (given here in Section 10.2 and earlier in Cutler (1997)) are motivated in part by the desire to equate a deterministic time series with a finite-dimensional dynamical system evolving on an appropriate state space (in other words, we take the concept of finite-dimensional dynamical system as a starting point for our notion of deterministic). They are further motivated by the common intuition that “deterministic” translates to “perfectly predictable”. However, our definitions also fit seamlessly into the framework of time-delay embeddings and system reconstruction (discussed briefly earlier), and so they are ideally suited for that environment.

Throughout the following we will be concerned with time series X_n , $n = 1, 2, \dots$ taking values in a closed subset $K \subseteq \mathbb{R}^u$ for some $1 \leq u < \infty$. We will let P denote the distribution of the time series on the Borel sets \mathcal{B}^∞ of the infinite product space $K^\infty = \times_{m=1}^\infty K$ (equipped with the usual product topology) and let P_n denote the distribution of (X_1, \dots, X_n) on the Borel sets \mathcal{B}^n of $K^n = \times_{m=1}^n K$. The collection P_n , $n = 1, 2, \dots$ comprises the finite-dimensional or *joint distributions* of the time series. The time series is said to be *stationary* if, for each $n \geq 1$, the distribution of $(X_{k+1}, X_{k+2}, \dots, X_{k+n})$ coincides with P_n for each $k \geq 1$. Our attention will be confined to stationary time series.

The remainder of the chapter is organized as follows: Section 10.2 presents definitions of deterministic and stochastic for stationary time series, provides added rationale for these definitions beyond those presented in Cutler (1997), and describes the link with finite-dimensional and infinite-dimensional dynamical systems. Section 10.3 makes the connection to time-delay embeddings and state-space reconstruction, and gives a small survey of the competing Takens (1981) and Sauer et al. (1991) methods. In Section 10.4 we examine the relationship between scaling properties (fractal dimensions), the Grassberger-Procaccia algorithm, and determinism in time series. Finally, Section 10.5 consists of some examples where we show that system reconstruction is not possible by proving that the sampled time series are stochastic. These examples also show that coordinate projections (the most natural of functionals) can be “bad” functionals in some cases.

Discussions on the relationships between chaos, determinism, and time series analysis can be found in the papers by Packard et al. (1980), Casdagli (1989, 1992), Kaplan (1992), Kennel et al. (1992), Abarbanel (1993), Cutler (1997), and in the collections of papers edited by Smith and Tong (1992) and by Cutler and Kaplan (1997).

10.2 Dynamical Systems and Determinism

As noted in the Introduction, we take as our starting point for a fundamental notion of determinism the finite-dimensional dynamical system

(f.d.d.s.) (defined later for discrete time). Part of the motivation for this comes from the physics literature, where such systems often appear to be informally equated with determinism. We are also motivated by the question of predictability — an f.d.d.s. can be said to be completely determined (and therefore predictable) after observation of only a finite number of coordinates (equivalently, after acquisition of only a finite amount of information). As such, a human observer can hope to acquire sufficient information within a finite length of time to completely determine and predict an f.d.d.s. Of course there is a distinction between theoretical and “real-world” observations; the latter are always contaminated by errors, and so actual predictability of the system will depend on continuity and smoothness properties of the f.d.d.s. We will also address this issue. We begin with general measure-preserving systems.

Definition 10.1. (general dynamical system) *Let \mathcal{Y} be a metric space, and let $\varphi : \mathcal{Y} \rightarrow \mathcal{Y}$ be a Borel-measurable mapping. For each $y \in \mathcal{Y}$ the sequence of iterates $y, \varphi(y), \varphi^2(y), \dots$ is called the orbit with initial condition y . The pair (\mathcal{Y}, φ) constitutes the dynamical system. If Q is a distribution on the Borel sets \mathcal{B} of \mathcal{Y} satisfying $Q = Q\varphi^{-1}$ (that is, Q is an invariant measure for φ) then the quadruple $(\mathcal{Y}, \varphi, \mathcal{B}, Q)$ is called a measure-preserving dynamical system (m.p.d.s.).*

The inclusion of a measure Q introduces a kind of randomness into the dynamical system via the initial condition — the initial condition becomes a random variable Y with distribution Q . However, this randomness is only observable under repeated sampling of the initial condition; it cannot be seen from a single orbit of the system. In the future we will always assume the presence of a measure (generally invariant) which governs the distribution of the initial condition.

Note also that if (\mathcal{Y}, φ) is a dynamical system and the initial condition Y is selected randomly according to some distribution Q then iterating Y under φ produces a time series $Y_n = \varphi^{n-1}(Y)$, $n = 1, 2, \dots$ with $\varphi^0(Y) = Y$, called *the time series corresponding to φ and Q* . The joint distributions of this time series will be denoted by Q_n , $n = 1, 2, \dots$ where Q_n is the distribution of (Y_1, \dots, Y_n) . It follows that $(\mathcal{Y}, \varphi, \mathcal{B}, Q)$ is an m.p.d.s. (that is, Q is invariant) if and only if the corresponding time series Y_n , $n = 1, 2, \dots$ is stationary with $Q_1 = Q$. Note also that the conditional distribution of (Y_1, \dots, Y_n) given $Y = y$ is a Dirac point mass with all probability concentrated at the point $(y, \varphi(y), \dots, \varphi^{n-1}(y))$. See Lasota and Mackey (1994) for a discussion of the properties of an iterated m.p.d.s.

There are two special cases of measure-preserving dynamical system in which we will be particularly interested, the one where $\mathcal{Y} = K$ and K is a closed subset of \mathbb{R}^u for $1 \leq u < \infty$, and the one where $\mathcal{Y} = K \subseteq \mathbb{R}^\infty$.

Definition 10.2. (finite-dimensional dynamical system) We say that a *m.p.d.s.* $(K, \varphi, \mathcal{B}, Q)$ is a *finite-dimensional dynamical system (f.d.d.s.)* if K is a closed subset of \mathbb{R}^u where $1 \leq u < \infty$.

Definition 10.3. (infinite-dimensional dynamical system) Here we restrict ourselves to the case where $K \subseteq \mathbb{R}^\infty$ and call this an *infinite-dimensional dynamical system (i.d.d.s.)*. However, we note that it will sometimes be possible to meaningfully identify an i.d.d.s. with a system evolving on a finite-dimensional domain, in which case the process may be regarded as an f.d.d.s.

Any reasonable definition of *deterministic* for a stationary time series should, at the very least, always result in the conclusion that the stationary time series $Y_n = \varphi^{n-1}(Y)$, $n = 1, 2, \dots$ corresponding to a measure-preserving f.d.d.s. is deterministic. The approach we take below does just that — the preceding Y_n series becomes a special case of a deterministic time series. Some aspects of this approach and the definitions below were given earlier in Cutler (1997).

Definition 10.4. (predictive dimension) Let X_n , $n = 1, 2, \dots$ be a strictly stationary time series taking values in a closed subset $K \subseteq \mathbb{R}^u$, $1 \leq u < \infty$. We define the *predictive dimension p* of the time series to be

$$p = \min\{n \geq 1 \mid \text{there exists } T : K^n \rightarrow K \text{ such that } X_{n+1} = T(X_1, \dots, X_n) \text{ w. p. 1}\}.$$

If no such function T exists for any $n \geq 1$, we set $p = \infty$. In the case that $p < \infty$ the function T corresponding to $n = p$ is called the *predictor function* of the time series. If a predictor function T exists then $X_{p+1} = T(X_1, \dots, X_p)$ w. p. 1. and clearly T is P_p -a.s. unique. Moreover, as a consequence of stationarity we also have

$$X_{m+p+1} = T(X_{m+1}, X_{m+2}, \dots, X_{m+p}) \text{ w. p. 1} \quad (10.2)$$

for all $m \geq 0$.

The notion of predictive dimension leads to precise definitions of the terms *deterministic* and *stochastic* for stationary time series, as well as to a precise notion of *predictability*:

Definition 10.5. We say that a stationary time series X_n , $n = 1, 2, \dots$ taking values in \mathbb{R}^u , $1 \leq u < \infty$, is *deterministic* if $p < \infty$ and *stochastic* if $p = \infty$, where p is the predictive dimension. Moreover, in the deterministic case, we say that the time series is also *predictable* if the predictor function T is P_p -a.s. continuous.

Thus a *deterministic time series* is one which can predict itself perfectly based on a fixed finite number of past (error-free) observations. There is no

assumption made or inferred about the nature of the underlying generating mechanism (or source) of the time series — the time series is evaluated only in reference to itself. It will be a point of later interest to consider whether a time series and its underlying source need share common deterministic/stochastic properties. (The example in the Introduction indicates that they need not.) Moreover, a *predictable time series* is a deterministic one in which the predictor function is continuous. This allows for some degree of prediction even in the case of observational error. Thus “predictable” might also be called “real-world deterministic”. Note also that a *stochastic time series* is one which *cannot* predict itself based on a fixed finite number of past observations; this constitutes a very large class of processes. For some comments on the meaning of stochastic, see points 5), 6), 7), and 8) later in this section.

We note that under Definition 10.5, the time series $Y_n = \varphi^{n-1}(Y)$, $n = 1, 2, \dots$ corresponding to φ is always deterministic with $p = 1$ and $T = \varphi$. Conversely, if X_n , $n = 1, 2, \dots$, is a time series with $p = 1$ and predictor function T , then the time series must correspond to the f.d.d.s. $\varphi = T$. Thus it follows that a deterministic time series with $p \geq 2$ cannot correspond to a measure-preserving f.d.d.s. However, it is possible to identify such a time series with a coordinate projection of a measure-preserving f.d.d.s. evolving on the finite product space K^p ; we call this system the *canonical f.d.d.s.* associated with the time series.

Theorem 10.1. (canonical f.d.d.s.) *Let X_n , $n = 1, 2, \dots$ be a stationary time series taking values in $K \subseteq \mathbb{R}^u$, $1 \leq u < \infty$, and having joint distributions P_n , $n = 1, 2, \dots$, finite predictive dimension p , and predictor function T . Then the mapping $\varphi_T : K^p \rightarrow K^p$ defined by*

$$\begin{aligned} \varphi_T(x) &= T(x) && \text{if } p = 1 \\ \varphi_T(x_1, \dots, x_p) &= (x_2, \dots, x_p, T(x_1, \dots, x_p)) && \text{if } p \geq 2 \end{aligned} \tag{10.3}$$

defines a measure-preserving f.d.d.s. on K^p with invariant distribution $Q = P_p$. Taking the initial condition to be $Y = (X_1, \dots, X_p)$ and defining the coordinate projection $\pi : (\mathbb{R}^u)^p \rightarrow \mathbb{R}^u$ by $\pi(x_1, \dots, x_p) = x_1$, the original time series is then recovered as the projection $X_n = \pi(\varphi_T^{n-1}(Y))$ for $n = 1, 2, \dots$

If $p = 1$, the canonical f.d.d.s. coincides with the measure-preserving f.d.d.s. generating the time series. If $p \geq 2$, Theorem 10.1 exhibits a dynamic correspondence between the points (x_1, \dots, x_p) of K^p (and their evolution) and the realizations of the time series (and their evolution). In order to see the significance of this, we first note that *every* stationary time series can be represented as a coordinate projection of a dynamical system if we make the state space large enough — specifically, if we take the state space to be the infinite product space $K^\infty = \times_{m=1}^\infty K$. This construction leads to an i.d.d.s.

Definition 10.6. (left-shift dynamical system) We let L be the left-shift operator $L : K^\infty \rightarrow K^\infty$ given by

$$L(x_1, x_2, x_3, \dots) = (x_2, x_3, x_4, \dots). \tag{10.4}$$

Stationarity of the time series is equivalent to P being invariant under L , so $(K^\infty, L, \mathcal{B}^\infty, P)$ is a measure-preserving dynamical system. The time series X_1, X_2, \dots can be represented as $X_n = \pi(L^{n-1}(Z))$ where $Z = (X_1, X_2, \dots)$ is the initial condition (selected according to P) and π is the projection onto the first coordinate; that is, $\pi(x_1, x_2, x_3, \dots) = x_1$. We note that the left-shift operator defines an infinite-dimensional dynamical system as discussed in Definition 10.3.

It follows from the preceding discussion and Theorem 10.1 that in the case of finite predictive dimension p we obtain two representations of our time series, one as a coordinate projection of an f.d.d.s. and the other as a coordinate projection of an i.d.d.s. However, it is straightforward to establish that these two representations are dynamically equivalent in the following sense.

Definition 10.7. (dynamical equivalence) *M.p.d.s.* $(\mathcal{Y}, \varphi, \mathcal{B}_1, Q)$ and $(\mathcal{X}, \psi, \mathcal{B}_2, V)$ are called *measure conjugate* or *dynamically equivalent* if there exists a mapping $\Psi : \mathcal{Y} \rightarrow \mathcal{X}$ satisfying the following:

1. Ψ is a measurable mapping and $V = Q\Psi^{-1}$.
2. There exist sets $F_0 \in \mathcal{B}_1$ and $G_0 \in \mathcal{B}_2$ such that $Q(F_0) = V(G_0) = 1$ and Ψ is 1-1 and onto between F_0 and G_0 .
3. $\psi \circ \Psi = \Psi \circ \varphi$ (*conjugacy property*).

Ψ is called a *measure conjugacy* between the two m.p.d.s.

Billingsley (1978, p. 53) uses the term *isomorphism* for what we have called *measure conjugacy*. We chose a more explicit descriptive term because it is often necessary to impose further structure (e.g., topological and/or differential) on the dynamical systems and conjugacies in order to preserve desired properties. In the absence of such additional structure, measure conjugacy is a rather weak relation which may identify many (apparently dissimilar) systems.

In the case of our two representations for finite predictive dimension p , it is easy to establish a measure conjugacy between them:

$$(K^\infty, L, \mathcal{B}^\infty, P) \xleftrightarrow{\Psi} (K^p, \varphi_T, \mathcal{B}^p, P_p). \tag{10.5}$$

The mapping $\Psi : K^\infty \rightarrow K^p$ defined by $\Psi(x_1, x_2, \dots) = (x_1, \dots, x_p)$ has the property that $P_p = P\Psi^{-1}$ and also that $\Psi \circ L(x_1, x_2, \dots) = \varphi_T \circ$

$\Psi(x_1, x_2, \dots)$ w. p. 1 using (10.2). It also follows from (10.2) that Ψ is one-to-one on a set of P -measure one. Consequently (10.5) defines a measure conjugacy and the two dynamical systems may be regarded as dynamically equivalent. Additionally, we see that Ψ is a bi-Lipschitz mapping provided T itself is Lipschitz continuous. This preserves metric properties between the two spaces.

We are now in a position to discuss several points:

1) First note that the preceding paragraphs illustrate the fact that a time series may arise as a functional of more than one dynamical system (in fact there will generally be several possible sources for a given time series). The two different sources for a deterministic time series given in Theorem 10.1 and Definition 10.6 were shown to be dynamically equivalent, but we may ask whether this is always, or almost always, the case. Moreover, dynamical equivalence can be a very weak relation which does not necessarily imply equivalence of other properties of interest which may be determined by topological and metric structures on the spaces. Scaling quantities such as fractal dimensions, for example, may not be preserved between the systems unless the conjugacy is bi-Lipschitz.

2) Definition 10.6 shows that every stationary time series can be identified with a measure-preserving dynamical system (an i.d.d.s.); hence such an identification in and of itself is not useful in arriving at a definition of determinism. The appropriate distinguishing criterion appears to be *finiteness*. Can it be identified with an f.d.d.s.?

3) The preceding question brings us back to an important point noted prior to Definition 10.1 — the problem of observational error and continuity of the predictor function T . We have chosen to use the word “predictable” in the case where T is *a.s.* continuous, since for any desired level of prediction accuracy we can find a level of observational accuracy which will provide that wished-for prediction accuracy. From the point of view of prediction in nonlinear chaotic systems, there seems to be no advantage in imposing a great deal more in the way of smoothness. The study of chaos has shown us that arbitrarily smooth nonlinear systems can produce orbits which diverge very rapidly, so that minor errors in the initial conditions are quickly magnified and predictability is lost (to the observer) within a few iterations. However we note that the embedding theorems of the next section, and the scaling results of the section following it, require some additional smoothness (at least Lipschitz continuous) on the predictor function T .

We might also comment on the related problem of *estimating* T from data in the more typical situation where T is unknown. Such estimation methods generally rely on some smoothness properties of T as well as on ergodicity of the time series.

4) We retain the notion of “deterministic” along with “predictable” because at this point it seems good, in principle, to have a definition which is separate from questions of smoothness and problems of observational or system error. (It remains to be seen if one can in fact separate these two notions.) Moreover, sufficiently interesting phenomena can be found even in the idealized case of exact error-free observations.

5) Note that our definition of deterministic was carefully restricted to time series taking values in \mathbb{R}^u for finite u . Any dynamical system obviously has $p = 1$, and we wanted to exclude the case where the domain was \mathbb{R}^∞ , on the assumption that this required information on an infinite number of coordinates (and, in any case, would render *all* stationary time series deterministic by identification with the left-shift i.d.d.s.). However, some i.d.d.s. might properly be considered deterministic (that is, as f.d.d.s.) by an identification such as given in (10.5).

6) We also note that the introduction of observational or other types of noise (for example, system noise) creates added sources of randomness beyond that contained in the initial condition. The dynamics are no longer the same. In the case of system noise, the result may even be a time series whose qualitative behavior is very different from that of the noise-free version. We consider these to be “higher-level” problems where the relevant question becomes one of quantifying the *degree* of noise or stochastic component present rather than a question of deterministic vs. stochastic — they are clearly stochastic by Definition 10.5. Another approach, deviating somewhat from Definition 10.5, would be to consider several different levels or grades of “deterministic” that allow for some kinds of noise. For results on stochastic dynamical systems of the form

$$X_n = T(X_{n-p}, \dots, X_{n-1}) + \epsilon_n \quad (10.6)$$

(that is, deterministic systems plus system noise) see Cheng and Tong (1992, 1995). Also see the last three chapters of Lasota and Mackey (1994) as well as Stark et al. (1996).

7) It has been pointed out that, in practice, a deterministic time series with a very large predictive dimension p (in other words, a very high-dimensional system) behaves for all intents and purposes like a stochastic process. Of course this is due to limitations on observational accuracy and, especially, sample size. Stochastic models may be more useful for predicting such systems from limited data, but this does not affect the theoretical validity of our definitions.

8) Finally, we have been discussing Definition 10.5 almost exclusively from the point of view of what “deterministic” should mean. But it is also worth considering what we want “stochastic” to mean. Models such as (10.6) are generally considered stochastic by everyone — frequently the

noise sequence ϵ_n , $n = 1, 2, \dots$ is *I.I.D.* or, at the very least, involves a great deal of independence. The sequence of 0's and 1's generated by the binary shift in (10.1) is *I.I.D.* and also certainly regarded as stochastic by all. However Definition 10.5 allows for stochastic time series which involve far less randomness than we might normally consider; some of these time series might bridge the gap between what we regard as truly deterministic and what we regard as truly stochastic. The examples in Section 10.5 should be viewed at least once from this perspective.

10.3 Time-Delay Embeddings and Determinism

In the previous section we considered an arbitrary stationary time series X_n , $n = 1, 2, \dots$. In this section we want to focus on those time series resulting from sampling an f.d.d.s. as discussed following (10.1). The basic situation is this: we are interested in the evolution of a measure-preserving f.d.d.s. φ (as given in Definition 10.2) and, ideally, would like to record the corresponding time series $Y_n = \varphi^{n-1}(Y)$, $n = 1, 2, \dots$. In practice, however, we are often unable to view the complete system evolving on its state space $K \subseteq \mathbb{R}^v$ but rather observe a *functional* or *projection* $h : K \rightarrow \mathbb{R}^u$ of the system (typically $u < v$, often $u = 1$). This produces an *observed* or *sampled* time series X_n , $n = 1, 2, \dots$ taking values in \mathbb{R}^u , where $X_n = h(Y_n) = h(\varphi^{n-1}(Y))$ and Y is the random initial condition of φ chosen according to an invariant distribution Q . The assumption that φ is measure-preserving ensures that X_n , $n = 1, 2, \dots$ is stationary.

The sampled time series is generally less predictable (or, if you like, more stochastic) than the underlying f.d.d.s. because h almost always condenses information by collapsing several pre-images to a common value. If we observe $Y_1 = y$ in the f.d.d.s. then Y_2 is uniquely determined by $Y_2 = \varphi(y)$. However, if we observe $X_1 = x$ in the sampled time series, we only obtain the information that $Y_1 \in \{y \mid h(y) = x\}$ and hence that $X_2 \in \{h(\varphi(y)) \mid h(y) = x\}$. Thus the conditional distribution of X_2 given $X_1 = x$ is not a Dirac point mass and the predictive dimension of the sampled time series exceeds 1. Another way of saying this is that X_1 does not uniquely determine the initial condition Y_1 and so the future evolution of the system is still uncertain to the observer. Abarbanel et al. (1993) point out that this feature often contributes to the erratic chaotic appearance of the sampled time series. The idea behind the method of *time-delay embeddings* is that while observation of a single value $X_1 = x$ will not pin down the underlying initial condition Y_1 , observation of a sufficiently long string, say $X_1 = x_1, \dots, X_d = x_d$, may do so. (This is essentially equivalent to obtaining a unique solution to a system of d equations.) Moreover, the sequence of delay vectors $(X_k, X_{k+1}, \dots, X_{k+d-1})$, $k = 1, 2, \dots$ may then sketch out in $(\mathbb{R}^u)^d$ a “copy” of the evolution of Y_1, Y_2, \dots making it possible to study

the hidden underlying system via the sampled time series. Abarbanel et al. (1993) refer to this process as *unfolding the attractor*; it is also known as *reconstructing the state space*. We describe this formally in the following:

Definition 10.8. *The delay coordinate mapping $\Phi_d : K \subseteq \mathbb{R}^v \rightarrow (\mathbb{R}^u)^d$ with embedding dimension d is defined by*

$$\Phi_d(y) = (h(y), h(\varphi(y)), \dots, h(\varphi^{d-1}(y))). \quad (10.7)$$

For $d = 1$ we have $\Phi_1(Y) = h(Y) = X$ which has distribution P_1 . For $d \geq 2$ we have $\Phi_d(Y) = (X_1, \dots, X_d)$ which has distribution P_d . Note also that $\Phi_d(\varphi^{k-1}(Y)) = \Phi_d(Y_k) = (X_k, X_{k+1}, \dots, X_{k+d-1})$ so this is the method of generating the delay vectors discussed in the preceding paragraph.

We now make the link with the material in Section 10.2.

Theorem 10.2. *Suppose for some $d \geq 1$ the delay coordinate mapping Φ_d is one-to-one on a set of Q -measure one. Then the sampled time series X_n , $n = 1, 2, \dots$ is deterministic with predictive dimension $p \leq d$ and predictor function $T = \Phi_p \circ \varphi \circ \Phi_p^{-1}$. Furthermore, Φ_p is a measure conjugacy between the two measure-preserving f.d.d.s. $(K, \varphi, \mathcal{B}, Q)$ and $(K_0, \varphi_T, \mathcal{B}^p, P_p)$; that is, the canonical f.d.d.s. of the sampled time series is dynamically equivalent to the original underlying f.d.d.s.*

Proof: Since P_p is the distribution of (X_1, \dots, X_p) , it follows from Definition 10.8 that $P_p = Q\Phi_p^{-1}$. Now, assuming that Φ_d is an embedding for some d , it follows that the predictive dimension p of the sampled time series will be the smallest value of d for which Φ_d is one-to-one w. p. 1; hence Φ_p will be one-to-one on a set of Q -measure one. It is then easy to see that $T = \Phi_p \circ \varphi \circ \Phi_p^{-1}$ and $\Phi_p \circ \varphi = \varphi_T \circ \Phi_p$. Thus Φ_p is a measure conjugacy. \square

It is worth stating explicitly that state-space reconstruction via delay coordinates leads directly to creation of the canonical f.d.d.s. $(K_0, \varphi_T, \mathcal{B}^p, P_p)$; it is this dynamical system which is being sketched out by the delay vectors. Theorem 10.2 shows that this system is, under certain circumstances, dynamically equivalent to the underlying f.d.d.s. of interest, so we can learn about the latter from the former. However, the delay coordinate space $K_0 \subseteq (\mathbb{R}^u)^d$ may not be the natural coordinate system for φ and interpretation of the canonical f.d.d.s. in terms of a physical model can be difficult. Very often when the approach taken is to estimate T directly from time series data we are engaging in “black box” modeling; the predictor T is estimated as closely as possible by some statistical criterion and method but there is no physics involved. The alternative (one which makes physicists happier) is to use knowledge of the physical system, when possible, to develop a meaningful dynamical model, a set of difference equations or differential equations, back on the original state space K .

It is natural now to wonder when the delay coordinate mapping Φ_d will be one-to-one for some $d \geq 1$. In fact we often require additional structural properties on Φ_d beyond one-to-one; for example, we might require that Φ_d be bi-Lipschitz in order that metric properties, such as fractal dimensions, are preserved. This leads to the following definition:

Definition 10.9. (embedding) *A mapping $\Phi : K \subseteq \mathbb{R}^v \rightarrow K_0 \subseteq \mathbb{R}^w$ is called a (diffeomorphic) embedding if Φ is one-to-one and bi-Lipschitz, and both Φ and its inverse Φ^{-1} are diffeomorphisms.*

The bi-Lipschitz property ensures that metric properties, such as scaling properties, are preserved between K and K_0 . The diffeomorphism property preserves local differential structure between the spaces; see Sauer et al. (1991) and Cutler (1997).

Takens (1981) and Sauer et al. (1991) provide theorems which basically state that the delay coordinate mapping will, in most cases, yield an embedding for sufficiently large d . That is, given a smooth f.d.d.s. φ under certain mild technical conditions, “almost all” smooth functionals h result in the delay coordinate mapping Φ_d being an embedding for sufficiently large d . Here “smooth” means twice-differentiable, sometimes only continuously differentiable. We will not supply other technical conditions of the theorems here. The most interesting difference between the Takens approach and the Sauer et al. approach lies in the interpretation of “almost all”. Takens (1981) interprets this topologically, and the set of h for which an embedding is obtained (the set of *good* h) is *generic*; that is, it is an open dense subset of the space of all smooth h . Hunt et al. (1992) and Sauer et al. (1991) point out that open dense sets can often be very small in terms of measure, and so they prefer a measure-theoretic approach. They show that the collection of good h forms a *prevalent* subset of the set of all smooth h ; that is, the complementary set consisting of all *bad* h is a *shy* set. A shy set is the equivalent of a set of zero Lebesgue measure in an infinite-dimensional space (in this case the space of all smooth h). We close this section with the following points:

- 1) One difficulty with the theorems of Takens (1981) and Sauer et al. (1991) is that there is no criterion for determining whether a specific functional h is good or bad. Moreover, being good or bad is not a property of the functional h in isolation; it is a property of h and φ together. A functional may be good for one f.d.d.s. but bad for another. We should add that, genericity and prevalence properties notwithstanding, it is not difficult to construct pairs (φ, h) where no embedding results.

- 2) A functional can be bad in more than one way. It may be that an embedding is achieved between the canonical f.d.d.s. and only one component (a subsystem) of the underlying f.d.d.s. In this case the sampled time series

has finite predictive dimension but there is no measure conjugacy with the desired original f.d.d.s. Some examples of this are provided in Section 10.5. It may also happen that the functional h produces a sampled time series which is actually stochastic and hence no embedding is possible with φ or any other f.d.d.s. In fact if one can prove that a time series is stochastic according to Definition 10.5 then it always follows that no embedding can exist with any f.d.d.s. Examples of this are also given in Section 10.5.

3) We also note that Stark et al. (1996) develop embedding theorems, based on Takens' genericity approach, for certain stochastic time series arising from both forced and noisy systems. This is a new and important direction which needs to be integrated with ours.

4) Finally we note again that a time series can arise mathematically as a functional of many different systems. If the time series has finite predictive dimension then an embedding will exist between the canonical f.d.d.s. and some of these systems but not others. It would be of interest to keep the sampled time series fixed and examine the collection of all theoretically possible or potential generating mechanisms. This is the "inverse problem" to the Takens and Sauer et al. embedding theorems.

10.4 Scaling Structures and Determinism

One of the first techniques for searching for determinism in observed time series was the *Grassberger-Procaccia (GP) algorithm*; see Grassberger and Procaccia (1983) and the discussion of the algorithm and estimation methods in Smith (1992). This technique was predicated on the assumption, or rule-of-thumb, that in a deterministic time series the fractal dimensions associated with the joint distributions P_n would level off at fixed finite values once $n \geq d$, where d is the correct embedding dimension, whereas in a stochastic time series the fractal dimensions would grow without bound as $n \rightarrow \infty$. Experimentalists would estimate the fractal dimension D_n of the joint distribution P_n then plot \hat{D}_n vs. n . If the plotted values approached a finite asymptote as $n \rightarrow \infty$, this was taken as evidence of determinism or low-dimensional dynamics. If, in addition, the limiting asymptote was fractional in value, this was taken as evidence of chaos. We will not discuss this latter aspect of the algorithm here.

The definitions of deterministic and stochastic were not quite clear in the GP algorithm. Moreover, it was not certain what distinction, if any, was made between determinism in the sampled time series and determinism in the underlying generating mechanism, or how these related to the scaling properties of the joint distributions of the sampled time series. Using our terminology and the results obtained so far, it is possible to obtain a complete picture of the theoretical GP algorithm (by theoretical we mean that

we use the true values D_n rather than estimates).

There are many possible choices for the fractal dimension D_n associated with P_n , and these different choices generally produce different values. The GP algorithm was initially developed in terms of the *correlation dimension* ν_n , where

$$\nu_n = \lim_{r \rightarrow 0} \frac{\log C_n(r)}{\log r} \quad (10.8)$$

and

$$C_n(r) = P_n \times P_n(\{(x, y) \mid \|x - y\| \leq r\}). \quad (10.9)$$

$C_n(r)$ is the (spatial) *correlation integral* of P_n . In fact $C_n(r)$ is the probability that two random independent points from P_n are within distance r of each other, and ν_n describes how this probability scales as a power of r as $r \rightarrow 0$. The paper by Osborne and Provenzale (1989) purported to show a counterexample to the GP rule-of-thumb by exhibiting a stochastic process in which the sequence of correlation dimensions ν_n converged to a finite value as $n \rightarrow \infty$. Theiler (1991) and Cutler (1994), taking different approaches, showed that this counterexample was false. In particular, Cutler (1994) argued that this counterexample can be viewed as an artifact of the simulation method used by the authors. Wolff (1990) gives some results on the behavior of sample versions of the correlation integral in the time series setting.

Another popular choice for fractal dimension is *information dimension* σ_n . There are several ways of defining σ_n . The following is one approach if P_n is supported on a compact subset $K_n \subseteq (\mathbb{R}^u)^n$. We first define the local scaling exponent or *pointwise dimension* at $x \in K_n$ by

$$\sigma_n(x) = \lim_{r \rightarrow 0} \frac{\log P_n(B(x, r))}{\log r} \quad (10.10)$$

where $B(x, r)$ is the ball of radius r centered at x , and then set

$$\sigma_n = \int \sigma_n(x) P_n(dx). \quad (10.11)$$

Thus information dimension is the average pointwise dimension of P_n . In many cases of interest the pointwise dimension will be constant almost surely, and the information dimension will coincide with this constant. We always have the inequality $\sigma_n \geq \nu_n$ and this inequality can be strict. We have argued elsewhere (Cutler (1998)) that information dimension is generally a more informative quantity than correlation dimension. However, it is often useful to look at both quantities, and yet other fractal dimensions simultaneously.

The following is our basic result concerning the GP algorithm. While the rule-of-thumb is not strictly correct, in many ways it is very close to being correct.

Theorem 10.3. *Let P_n , $n = 1, 2, \dots$ be the joint distributions of a stationary time series X_n , $n = 1, 2, \dots$ taking values in a compact set $K_0 \subseteq \mathbb{R}^u$. Let D_n be the fractal dimension of P_n , where $D_n = \nu_n$ or $D_n = \sigma_n$. Then the following are true:*

1. *The sequence D_n is increasing in n ; that is, $D_n \leq D_{n+1}$ for every n .*
2. *If the time series arises as a functional of an f.d.d.s., where both the functional and the f.d.d.s. are at least Lipschitz continuous, then the sequence D_n converges to a finite value D_0 . (This convergence occurs whether the sampled time series is deterministic or stochastic; it may be stochastic.)*
3. *If $\lim_{n \rightarrow \infty} D_n = \infty$ and the only possible predictor functions are Lipschitz continuous, then the time series is stochastic. Moreover, the time series cannot arise as a Lipschitz functional of a Lipschitz f.d.d.s.*
4. *It is possible for a stochastic time series to yield $\lim_{n \rightarrow \infty} D_n = D_0 < \infty$. Some of these time series can be represented as a Lipschitz functional of a Lipschitz f.d.d.s., others cannot.*

Remark 10.1. *Proofs of the first two items are given in Cutler (1997). The third item is simply the contrapositive of the second. These first three items taken together comprise what is right about the rule-of-thumb in the GP algorithm. It is the fourth item which deviates from the rule-of-thumb and constitutes a most interesting grey area. It includes stochastic time series which are Lipschitz functionals of smooth f.d.d.s. and might be regarded as almost deterministic (both examples in Section 10.5 are of this type). But it also includes other time series which involve so much independence that they must be regarded as truly stochastic. Examples of this type, such as randomly iterated function systems, are given in Cutler (1998).*

10.5 Examples

The following are two examples of real-valued time series arising as smooth functionals of chaotic f.d.d.s. evolving on the unit square. In each case the plotted time series appears random and erratic, and we will show that both are in fact stochastic by our definition. However, it follows from Theorem 10.3 that the GP algorithm would yield a finite fractal dimension D_0 for each of these time series. Also in each of these cases we find that any function of either coordinate projection $\pi_1(x, y) = x$ or $\pi_2(x, y) = y$ will fail to result in an embedding with the f.d.d.s. This does not contradict any of the embedding theorems — the set of all smooth functions of the coordinate mappings is a shy subset in the space of all smooth functionals $h : [0, 1]^2 \rightarrow \mathbb{R}$. However, coordinate projections are common choices for

functionals, and we see that if we keep the functional fixed (choose a coordinate projection) then obtaining a good or bad result depends on the coordinate system we choose for our f.d.d.s.

Example 10.1. Define $\tilde{\psi} : [0, 1]^2 \rightarrow [0, 1]^2$ by $\tilde{\psi}(x, y) = (\psi(x), \psi(y))$ where ψ is the tent map

$$\psi(x) = \begin{cases} 2x & \text{for } 0 \leq x \leq 1/2 \\ 2(1-x) & \text{for } 1/2 \leq x \leq 1 \end{cases} . \quad (10.12)$$

Lebesgue measure, restricted to the unit interval, is an ergodic invariant measure for ψ , and it follows that two-dimensional Lebesgue measure, confined to the unit square, is ergodic and invariant for $\tilde{\psi}$. Note that this latter f.d.d.s. is uncoupled as the x and y coordinates iterate independently of one another.

Suppose now we define $f : [0, 1] \rightarrow [0, 1]$ by

$$f(x) = \begin{cases} 1 & \text{for } 0 \leq x \leq 1/2 \\ 2(1-x) & \text{for } 1/2 \leq x \leq 1 \end{cases} \quad (10.13)$$

and then define the functional $h : [0, 1]^2 \rightarrow [0, 1]$ by $h(x, y) = f(x)y$. We will denote the deterministic time series resulting from iterating the i.d.d.s. $\tilde{\psi}$ (with a random initial condition) by $(X_1, Y_1), (X_2, Y_2), \dots$. Applying h produces the stationary time series Z_1, Z_2, \dots where $Z_n = h(X_n, Y_n)$. A plot of 300 points of the Z series is given in Figure 10.1. We note that the functional h is Lipschitz continuous but not differentiable because of the behavior of f at the point $x = 1/2$. We could modify f , even make it infinitely differentiable, without affecting the relevant qualitative behavior of the Z series. However, to do so would complicate our analysis and we prefer to leave f as it is.

Time-delay embeddings, applied to the Z series, will fail to reconstruct the underlying system $\tilde{\psi}$. In other words, h is a bad functional. We can prove this by showing that the Z series is stochastic. Let $n \geq 2$ and suppose the initial condition X_1 of the x -coordinate satisfies $0 < X_1 \leq 2^{-n}$. (This event occurs with probability 2^{-n} .) Then X_1, X_2, \dots, X_n all lie in the interval $(0, 1/2]$ and so $f(X_k) = 1$ for $k = 1, \dots, n$. Thus $Z_k = Y_k$ for $k = 1, \dots, n$ (and no additional information is obtained about the x coordinate during this run of observed Y 's). We can detect the occurrence of such a run in the Z series by noting that it corresponds to $Z_k = \psi(Z_{k-1})$ for $k = 2, \dots, n$. Now if $f(X_k) = 1$ for $k = 1, \dots, n$ (equivalently, if $Z_k = \psi(Z_{k-1})$ for $k = 2, \dots, n$) then, using Lebesgue measure, there is a 50% probability that $0 < X_n \leq 1/4$ and a 50% probability that $1/4 < X_n \leq 1/2$. In the first case we get $f(X_{n+1}) = 1$ and $Z_{n+1} = \psi(Z_n)$ (that is, the run continues) and in the second case the outcome is $Z_{n+1} = 2(1-2X_n)\psi(Z_n)$. The variable $2(1-2X_n)$ is uniformly distributed over $[0, 1]$ and its value cannot be predicted by an observer based on the knowledge $Z_k = \psi(Z_{k-1})$ for $k = 2, \dots, n$. This shows that, for each $n \geq 2$, there are particular choices of (z_1, \dots, z_n) (which constitute a set of positive probability) for which the conditional distributions




FIGURE 10.1. Plot of 300 points of the Z series.

$P(Z_{n+1} \in \cdot / Z_n = z_n, \dots, Z_1 = z_1)$ are not Dirac point masses, and hence an embedding of dimension n cannot exist with the underlying f.d.d.s. $\tilde{\psi}$. Moreover, due to ergodicity of the system $\tilde{\psi}$, arbitrarily long runs of $f(X_{k+j}) = 1$, $j = 1, \dots, m$ will occur infinitely often in the output sequence. Figure 10.2 shows a lag plot of Z_{n+1} vs. Z_n ; note that some of the underlying deterministic structure can be detected in the plot.




FIGURE 10.2. Lag Plot of Z_{n+1} vs. Z_n .

Let us add that had we chosen either of the coordinate projections $\pi_1(x, y) = x$ or $\pi_2(x, y) = y$ as our functionals, then we would also have been unable to

reconstruct our f.d.d.s. but for a different reason. Each functional time series would have been deterministic but of predictive dimension $p = 1$, picking up only the one-dimensional system ψ evolving in the corresponding coordinate. This is a consequence of the coordinates of $\tilde{\psi}$ being uncoupled.

Example 10.2. In this example we construct an f.d.d.s. where the coordinates are coupled but nonetheless one of the coordinate projections produces a stochastic time series and thus fails to reconstruct the f.d.d.s. Moreover, the f.d.d.s. is a type of *threshold model* commonly considered in nonlinear modeling; see Tong (1990). We define $\varphi : [0, 1]^2 \rightarrow [0, 1]^2$ by $\varphi(x, y) = (S(x, y), L(y))$, where L is the standard logistic map $L(y) = 4y(1 - y)$ and S is defined by

$$S(x, y) = \begin{cases} L(x) & \text{if } y \leq 1/2 \\ (1 - c(y))L(x) + c(y)x & \text{if } y > 1/2 \end{cases} \quad (10.14)$$

and $c(y) = 2(y - 1/2)$. The basic idea is to switch back and forth between two regimes: $L(x)$ for $y \leq 1/2$ and x for $y > 1/2$. The factor $c(y)$ is introduced to create a smooth linear transition between the two regimes. It is not necessary to select the function x for the second regime; it could be replaced by some other function $V(x)$. However, choosing x produces some interesting behavior; it shares the fixed point $x = 3/4$ with the logistic map but of course is not repelling—the result is that the x -coordinate series $X_n = \pi_1(\varphi^{n-1}(X_1, Y_1))$ tends to gravitate around $x = 3/4$ for stretches of time before being driven away by a change in regime and the repelling behavior of the logistic map. Figure 10.3 shows a plot of 700 points of the X series. The horizontal line corresponds to the fixed point $x = 3/4$.

Although we do not know the stationary ergodic behavior of the overall system φ , it is well known that the logistic map has the arcsine (or beta($\frac{1}{2}, \frac{1}{2}$)) distribution as ergodic invariant measure. If the initial condition Y_1 of the y -coordinate is close to 0 (and hence several subsequent iterates of $L(Y_1)$ fall below $1/2$), then the 1st coordinate of the system spends several iterations in the first regime; that is, $S(x, y) = L(x)$ or $X_n = L(X_{n-1})$. When the Y coordinate finally exits to the second regime $Y > 1/2$, the value of Y (and hence of $S(X, Y)$) cannot be predicted based on the observed X series up to that point; this is very similar in spirit to Example 10.1. Consequently the X -coordinate series is stochastic, and an embedding with φ cannot be constructed from it. Figure 10.4 shows a lag plot of X_{n+1} vs. X_n ; the “angel wings” illustrate the two boundaries $L(x) = 4x(1 - x)$ and $y = x$ as well as the stochastic region in between. The dark crossover point corresponds to the fixed point $x = 3/4$.

We also note that φ cannot be reconstructed from the y -coordinate projection, since lack of coupling in that direction (x has no influence on y) results in predictive dimension $p = 1$ and reconstruction only of $L(y)$.




FIGURE 10.3. Plot of 700 points of x coordinate of the threshold model.




FIGURE 10.4. Lag plot of threshold model showing “angel wings”.

Acknowledgments

I am indebted to David Brillinger for some gentle prodding on the concept of predictability, to the referee for forcing me to clarify some points, and

to Luis Gorostiza and Gail Ivanoff for their infinite patience.

A version of this paper was presented at the International Conference on Stochastic Models in June 1998 in honour of Professor Donald Dawson. Don has been my teacher, my Ph.D. supervisor, and a co-author. He is responsible for introducing me to the notions of fractal dimension and scaling in measure theory and to many other areas of mathematics and probability. But in addition to his mathematical generosity, Don creates around him a unique atmosphere of integrity and enthusiasm for research, learning, and academic life which will not be duplicated. It has been a privilege to be his student and it is a pleasure to thank him.

References

- [1] Abarbanel, H.D.I., Brown, R., Sidorowich, J.J., and Tsimring, L.S. (1993). The analysis of observed chaotic data in physical systems. *Rev. Modern Phys.* **65** 1331–1392.
- [2] Billingsley, P. (1978). *Ergodic Theory and Information*. Kriger, New York.
- [3] Casdagli, M.C. (1989). Nonlinear prediction of chaotic time series. *Physica D* **35** 335–356.
- [4] Casdagli, M.C. (1992). Chaos and deterministic *versus* stochastic and non-linear modelling. *J.R. Statist. Soc. B* **54** 303–328.
- [5] Cheng, B. and Tong, H. (1992). On consistent nonparametric order determination and chaos. *J. R. Statist. Soc. B* **54** 427–449.
- [6] Cheng, B. and Tong, H. (1995). Orthogonal projection, embedding dimension, and sample size in chaotic time series from a statistical perspective. In: *Chaos and Forecasting* (Ed.: H. Tong) World Scientific, Singapore, pp. 1–29.
- [7] Cutler, C.D. (1994). A theory of correlation dimension for stationary time series. *Philos. Trans. Roy. Soc. London A* **348** 343–355. Expanded version with proofs in *Chaos and Forecasting* (1995) (Ed.: H. Tong) World Scientific, Singapore, pp. 31–56.
- [8] Cutler, C.D. (1997). A general approach to predictive and fractal scaling dimensions in discrete-index time series. In: *Nonlinear Dynamics and Time Series* (Eds.: C.D. Cutler and D.T. Kaplan) *Fields Institute Communications* **11** American Mathematical Society, Providence RI, pp. 29–48.
- [9] Cutler, C.D. (1998). Computing pointwise fractal dimension by conditioning in multivariate distributions and time series. To appear in *Bernoulli*.
- [10] Cutler, C.D. and Kaplan, D.T. (Eds.) (1997). *Nonlinear Dynamics and Time Series*. Fields Institute Communications **11** American Mathematical Society, Providence RI.
- [11] Grassberger, P. and Procaccia, I. (1983). Characterization of strange attractors. *Phys. Rev. Lett.* **50** 346–349.
- [12] Hunt, B.R., Sauer, T., and Yorke, J.A. (1992). Prevalence: A translation-invariant “almost every” on infinite-dimensional spaces. *Bull. Amer. Math. Soc.* **27** 217–238.
- [13] Kaplan, D.T. and Glass, L. (1992). Direct test for determinism in a time series. *Phys. Rev. Lett.* **68** 427–430.

- [14] Kennel, M.B., Brown, R., and Abarbanel, H.D.I. (1992). Determining embedding dimension for phase space reconstruction using a geometrical construction. *Phys. Rev. A* **45** 3403–3411.
- [15] Lasota, A. and Mackey, M.C. (1994). *Chaos, Fractals, and Noise*. (2nd Ed.) Springer-Verlag, New York.
- [16] Osborne, A.R. and Provenzale, A. (1989). Finite correlation dimension for stochastic systems with power-law spectra. *Physica D* **35** 357–381.
- [17] Packard, N.H., Crutchfield, J.P., Farmer, J.D., and Shaw, R.S. (1980). Geometry from a time series. *Phys. Rev. Lett.* **45** 712–716.
- [18] Sauer, T., Yorke, J.A., and Casdagli, M. (1991). Embedology. *J. Statist. Phys.* **65** 579–616.
- [19] Smith, R.L. (1992). Estimating dimension in noisy chaotic time series. *J. R. Statist. Soc. B* **54** 329–351.
- [20] Smith, R.L. and Tong, H. (Eds.) (1992). *Proceedings of the Royal Statistical Society Meeting on Chaos*. In: *J. R. Statist. Soc. B* **54**.
- [21] Stark, J., Broomhead, D.S., Davies, M.E., and Huke, J. (1996). Takens embedding theorems for forced and stochastic systems. Submitted to the *Proceedings of the 2nd World Congress of Nonlinear Analysis*, Athens, Greece.
- [22] Takens, F. (1981). Detecting strange attractors in turbulence. In *Lecture Notes in Mathematics*. **898** Springer-Verlag, New York, pp. 366–381.
- [23] Tong, H. (1990). *Non-linear Time Series: A Dynamical System Approach*. Oxford University Press, New York.
- [24] Theiler, J. (1991). Some comments on the correlation dimension of $1/f^\alpha$ noise. *Phys. Lett. A* **155** 480–493.
- [25] Wolff, R.C.L. (1990). A note on the behaviour of the correlation integral in the presence of a time series. *Biometrika* **77** 689–697.

Chapter 11

Consistent Estimation of a Dynamical Map

Andrew Nobel

ABSTRACT Estimation of a nonlinear map F governing the evolution of an observed dynamical system is considered in two specific models. In the first model, F is successively applied to a fixed initial vector in the absence of noise, so that the the observed states of the system constitute a trajectory of F . In the second, dynamical noise model, the system is perturbed by independent noise between each application of F . Estimates of F are proposed for each model, and are shown to be consistent under general conditions. No assumptions are made regarding mixing rates of the observations. Both continuous and general measurable maps F are considered.

11.1 Introduction

The advent of modern computing and the recent interest in chaos have focused increasing attention on deterministic systems that exhibit random behavior. While there is no universally accepted definition of chaos, phenomena termed “chaotic” have frequently been studied in the context of dynamical systems. Dynamical systems are mathematical models in which a physical system of interest is described by a family of differential equations. The solution of the equations describes the time evolution of the dynamical system, starting from any initial condition.

In many situations, the data arising from measurements of a physical system are obtained at discrete, equally spaced instants of time. In some cases direct measurement of the state of the system is possible. More commonly, one makes periodic measurements of some scalar function ϕ of the state of the system. It is known from Takens’s embedding theorem [35, 3, 34] that, for generic functions ϕ and suitable integers d , one may study the dynamics of the system by means of time-delay vectors that consist of d successive scalar measurements. Thus, when it is in a steady state and no noise is present, observations of the physical system can be modeled by iteration of a fixed, nonlinear map $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$.

The statistical analysis of measurements from dynamical systems is complicated by the fact that such measurements can exhibit very long range dependence. Even when a system is perturbed by independent noise, its

measurements may fail to satisfy standard mixing assumptions, and existing statistical estimation theory may not apply to the analysis of these measurements. Nevertheless, it is often reasonable to assume that measurements of a dynamical system are stationary and ergodic, or when this assumption fails, that the system converges asymptotically to an ergodic steady state.

11.1.1 Overview

The subject of this chapter is estimation of a nonlinear map F governing an observed dynamical system. Two dynamical models are considered. In the first, F is successively applied to a fixed initial vector in the absence of noise, so that the measured states of the system constitute a trajectory of F . In the second, often referred to as a dynamical noise model, the system is perturbed by independent noise between each application of F . Histogram estimates of F are proposed and analyzed for each model. The primary goal of the chapter is to describe estimates of F that are consistent under general conditions. Throughout, no assumptions are made regarding the mixing rates of the observations. Both continuous and more general measurable maps F are considered. Although in many cases the proposed estimates can be implemented on a computer, we have not attempted to assess their empirical performance.

The problem of estimating an iterated map governing a dynamical system has previously been considered by a number of authors, working in several fields. Most often F is estimated with the ultimate goal of prediction, estimating Lyapunov exponents, or estimating the dimension of an attractor. Representative work and additional references can be found in the papers of Farmer and Sidorowich [14], Casdagli [8, 9], Kostelich and Yorke [21], Nychka *et al.* [31], Lu and Smith [24], and the book by Tong [36]. See also the surveys by Eckmann and Ruelle [13], Jensen [20], and Isham [19]. In most of this work iterates of the map are perturbed by observational or dynamical noise. Bosq and Guégan [7] studied kernel estimates of uniformly mixing transformations in the absence of noise. Lalley [23] describes a general means of reconstructing the orbit of a smooth diffeomorphism F , acting on a hyperbolic attractor, when the iterates of F are corrupted by observation noise.

In these references it is commonly assumed that the map under study is continuous or differentiable. Hofbauer and Keller [18], Mayer [25], and Denker and Keller [12] established central limit theorems for U-statistics and smooth functionals of noiseless dynamical systems that are generated by piecewise-monotone maps. More irregular transformations are of interest to ergodic theorists and may arise, for example, when one considers the Poincaré return map of a smooth flow to a low-dimensional set $A \subseteq \mathbb{R}^d$.

11.1.2 Outline

The two noise models studied in the chapter are defined in the next section. A description of the multivariate histograms used to define the estimates is given in Section 11.3. In Section 11.4 the problem of estimating a μ -preserving ergodic map $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$ from one of its trajectories is considered. It is shown in Proposition 11.1 that consistent estimation of F is possible if μ is comparable to a known reference measure with compact support.

In Section 11.5 the problem of estimating a map F in the presence of dynamical noise is considered. Theorem 11.1 gives a general consistency result for continuous maps. Density estimates of Györfi and Masry [16] are briefly discussed, and consistent estimates of the noise variance are described. It is shown in Theorem 11.3 that one may estimate a measurable map F in the dynamical noise model provided only that the observations are bounded and ergodic and that the distribution of the noise has a density.

11.2 Two Models for Dynamical Data

Let $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a fixed nonlinear map governing the one-step evolution of a dynamical system under study. Let \mathcal{B} denote the Borel subsets of \mathbb{R}^d . Here and in what follows it is assumed that F is measurable, so that $F^{-1}A \in \mathcal{B}$ for every $A \in \mathcal{B}$.

Model I: No Noise

In the simplest model of a dynamical system, the evolution of the system is determined by repeated application of F , in the absence of observational or dynamical noise. Starting from an initial vector $x \in \mathbb{R}^d$ successive observations of the system are described by the trajectory

$$x, Fx, F^2x, \dots \in \mathbb{R}^d. \quad (11.1)$$

Here F^i denotes the i -fold composition of F with itself. In this case the system evolves in a purely deterministic fashion. Thus, from (complete) knowledge of F and any single observation one can, in principle, reconstruct every subsequent measurement. In analyzing the model (11.1) it is typically assumed that the map F is measure preserving and ergodic.

Definition 11.1. A Borel-measurable map $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is said to preserve a probability measure μ on $(\mathbb{R}^d, \mathcal{B})$ if

$$\mu(F^{-1}A) = \mu(A) \quad \text{for every } A \in \mathcal{B}, \quad (11.2)$$

and is said to be ergodic if $F^{-1}A = A$ implies $\mu(A) = 0$ or 1. Equivalently,

F is ergodic if and only if

$$\frac{1}{n} \sum_{i=1}^n \mu(A \cap F^{-i}B) \rightarrow \mu(A)\mu(B) \text{ for every } A, B \in \mathcal{B}. \tag{11.3}$$

The books of Petersen [32] and Walters [37] give comprehensive introductions to ergodic theory.

Model II: Dynamical Noise

When the evolution of a system is mediated by independent noise, its state may be represented by a simple nonlinear autoregression. The resulting dynamical noise model has applications in the analysis of chaotic data.

Definition 11.2 A system is said to obey a dynamical noise model if its state evolves according to the recursion

$$X_{i+1} = F(X_i) + Z_{i+1} \quad i \geq 0 \tag{11.4}$$

where $X_0 \in \mathbb{R}^d$ is the (random) initial state of the system, $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a fixed map, and $Z_1, Z_2, \dots \in \mathbb{R}^d$ are I.I.D., independent of X_0 , and such that $EZ_i = 0$. The map F describes the deterministic component of the dynamics, while $\{Z_i\}$ are random perturbations whose influence on subsequent measurements is mediated by the action of F .

Definition 11.3 A stationary process $X_0, X_1, \dots \in \mathbb{R}^d$ is said to be ergodic if for every $l \geq 1$, and every pair of Borel sets $A, B \subseteq \mathbb{R}^d$,

$$\frac{1}{n} \sum_{i=1}^n \mathbb{P}\{X_0^{l-1} \in A, X_i^{i+l-1} \in B\} \rightarrow \mathbb{P}\{X_0^{l-1} \in A\}\mathbb{P}\{X_0^{l-1} \in B\} \tag{11.5}$$

as n tends to infinity, where $X_i^j = (X_i, \dots, X_j)$ for $i \leq j$.

For example, if $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$ preserves a measure μ and is ergodic, and if X_0 is a random vector with distribution μ , then (11.2) and (11.3) imply that the process X_0, FX_0, F^2X_0, \dots is stationary and ergodic. By contrast, due to the presence of the perturbations Z_i , ergodicity of the dynamical noise process (11.4) does not require that F preserve the distribution of the measurements X_i , or that F be ergodic. (To take a trivial example, let $F : \mathbb{R} \rightarrow \mathbb{R}$ be identically zero, and let $\{Z_i\}$ be any real valued I.I.D. sequence taking more than one value.)

The dynamical noise process (11.4) is a discrete time Markov chain. Results such as those in Nummelin [30], and Meyn and Tweedie [26], provide general conditions under which such chains are ergodic, or converge to an ergodic steady state. Dynamical noise models are also a special case of random dynamical systems. Many of the theoretical properties of such systems, e.g., topological dynamics, the existence of invariant measures, and

Lyapunov exponents, have been studied. See Kifer [22] or Arnold [4] for more details.

11.3 Multivariate Histograms

Let $X_0, X_1, \dots \in \mathbb{R}^d$ be a stationary ergodic sequence of random vectors. The one-step autoregression of $\{X_i\}$ is the (vector-valued) conditional expectation

$$G(x) = E(X_1 | X_0 = x),$$

which minimizes $E\|X_1 - G_0(X_0)\|^2$ over all measurable functions $G_0 : \mathbb{R}^d \rightarrow \mathbb{R}^d$. It may readily be verified that in both the dynamical noise model (11.4) and the deterministic model with $X_i = F^i X_0$, the map F of interest is equal to the one-step autoregression function G of the available measurements. Thus the problem of estimating F is equivalent to the problem of estimating an autoregression function. Proposition 11.1 and Theorem 11.3 are based on recent results of Nobel and Adams [28] concerning regression estimation from ergodic processes. Related work can be found in the papers of Delecroix [10], Yakowitz [38, 39], Delecroix and Rosa [11], Yakowitz *et al.* [40] and Morvai *et al.* [29]. For reviews of regression estimation from dependent processes, see Györfi, Härdle, Sarda and Vieu [17], Rosenblatt [33], Beran [5], and Bosq [6].

It has recently been shown by Adams [1], Nobel [27], and Yakowitz and Heyde [41] that no regression estimation scheme can be consistent for every stationary ergodic process, so restrictions must be placed on the observations in order to establish the consistency of a regression scheme under study. Most work to date places assumptions on the dependence (mixing) structure of the observations, on the regularity of the unknown (auto)regression function, or on both. For the models studied here it is sufficient to place assumptions on the one-dimensional distribution of the observations.

Each of the estimates proposed herein is based on a multivariate histogram, which is obtained by dividing the measurements X_i into disjoint cells and then averaging the corresponding values X_{i+1} within each cell. The family of histograms employed here will be defined later.

Partitions. Fix a nested sequence π_0, π_1, \dots of finite partitions of \mathbb{R}^d such that $\pi_0 = \{\mathbb{R}^d\}$, and such that for each vector $x \in \mathbb{R}^d$,

$$\lim_{k \rightarrow \infty} \text{diam}(\pi_k[x]) = 0. \quad (11.6)$$

Here $\pi_k[x]$ is the unique cell of π_k containing x , and we denote by $\text{diam}(A) = \sup_{u, v \in A} \|u - v\|$ the maximum Euclidean distance between any two points in A . Condition (11.6) allows the partitions π_k to have unbounded cells, provided the sequence of cells containing each fixed vector x eventually

shrinks down to x . The partitions π_l may be obtained, for example, by dividing $[-l, l]^d$ into cubes of side-length 2^{-l} , and letting the complement of $[-l, l]^d$ comprise a single cell.

Candidate Histograms and Empirical Error

Let measurements X_0, \dots, X_{n-1} be given. For each $k \geq 1$ define a multivariate histogram based on the partition π_k ,

$$G_{k,n}(u) = \frac{\sum_{i=0}^{n-2} X_{i+1} I\{X_i \in \pi_k[u]\}}{\sum_{i=0}^{n-2} I\{X_i \in \pi_k[u]\}}. \tag{11.7}$$

If a cell $\pi_k[u]$ contains no vector X_i , then set $G_{k,n}(u) = 0$. Furthermore, let

$$\Delta_{k,n}^2 = \frac{1}{n-1} \sum_{i=0}^{n-2} \|G_{k,n}(X_i) - X_{i+1}\|^2 \tag{11.8}$$

denote the empirical error of the histogram $G_{k,n}$ on X_0, \dots, X_{n-1} . An estimate of F is chosen from among the candidates $\{G_{k,n} : k \geq 1\}$ by selecting a suitable partition index k . Selection of k is based on the available measurements X_0, \dots, X_{n-1} , and on information about the distribution of X_0 .

11.4 Ergodic Systems without Noise

Consider the problem of estimating a μ -preserving ergodic map $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$ in the noiseless model, in which the available measurements form a trajectory

$$x, Fx, F^2x, \dots \tag{11.9}$$

of F starting from an initial value $x \in \mathbb{R}^d$. In this case, the measurements are purely deterministic and will usually fail to satisfy standard mixing assumptions: the rate at which the sample average $n^{-1} \sum_{i=0}^{n-1} h(F^i x)$ of a bounded function $h : \mathbb{R}^d \rightarrow \mathbb{R}$ converges to $\int h d\mu$ is typically unknown, and may be arbitrarily slow.

If the measurements are grouped into pairs $(x, Fx), (Fx, F^2x), \dots$, then each pair is a point on the graph of F . When F is continuous and $d = 1$, connecting neighboring points with straight lines will give pointwise consistent estimates of F on the support of μ . Similar piecewise linear estimates may be used in higher dimensions. Of interest here is the case when F may be highly irregular, so that some sort of local averaging is necessary to obtain good estimates.

11.4.1 Estimation of Irregular Maps

It will be shown that consistent estimation of an ergodic μ -preserving map F is possible, provided only that μ is comparable to a known reference

measure. Recall that the support of a distribution μ on \mathbb{R}^d is the smallest closed set $\Lambda \subseteq \mathbb{R}^d$ such that $\mu(\Lambda) = 1$.

Definition 11.4. Let μ_0 be a reference measure on $(\mathbb{R}^d, \mathcal{B})$ with compact support, and let $M \geq 1$ be a constant. Let $\mathcal{D}(\mu_0, M)$ be the family of all probability measures μ on $(\mathbb{R}^d, \mathcal{B})$ such that $\mu \equiv \mu_0$ and

$$\int \frac{d\mu_0}{d\mu} d\mu_0 \leq M^2. \tag{11.10}$$

The condition $\mu \equiv \mu_0$ means that $\mu(A) = 0$ if and only if $\mu_0(A) = 0$. This implies that the densities $d\mu/d\mu_0$ and $d\mu_0/d\mu$ exist, and that μ and μ_0 share the same support set. The constant M controls the deviation of μ from μ_0 . As M increases the measures in $\mathcal{D}(\mu_0, M)$ can differ more substantially from μ_0 .

Example 11.1. Let Λ be a compact subset of \mathbb{R}^d with Lebesgue measure $\lambda(\Lambda) > 0$, and let $\mu_0(A) = \lambda(A \cap \Lambda)/\lambda(\Lambda)$ be normalized Lebesgue measure on Λ . The family $\mathcal{D}(\mu_0, M)$ contains every probability measure μ having a density $f = d\mu/d\lambda$ such that $\{x : f(x) > 0\} = \Lambda$ and $\int_{\Lambda} (1/f) dx \leq M^2 \lambda(\Lambda)$. The latter condition says that f is suitably bounded away from zero on Λ , and is satisfied, for example, if $f(x) \geq M^{-2}$ for $x \in \Lambda$. In general the reference measure need not be absolutely continuous. For example, μ_0 might be Hausdorff measure on a known compact manifold of dimension less than d , or on some other bounded low-dimensional subset of \mathbb{R}^d .

Description of the Estimate. Suppose that a measure μ_0 with compact support and a constant $M \geq 1$ have been fixed. Let $\epsilon_1, \epsilon_2, \dots$ be any sequence of positive numbers tending monotonically to zero. Given measurements $x, Fx, \dots, F^{n-1}x$ with initial vector x , define the estimate

$$\hat{F}_n^{(1)}(u) = \frac{\sum_{i=0}^{n-1} F^{i+1}x \cdot I\{F^i x \in \pi_{k(n)}[u]\}}{\sum_{i=0}^{n-1} I\{F^i x \in \pi_{k(n)}[u]\}}, \tag{11.11}$$

where $k(n)$ is the largest integer $k \geq 1$ such that

$$\int \|G_{i,n} - G_{j,n}\| d\mu_0 \leq 2M\Delta_{i,n} + 2(1+M)\epsilon_i \text{ for } 1 \leq i \leq j \leq k. \tag{11.12}$$

Here $G_{i,n}$ and $\Delta_{i,n}$ are defined as in Equations (11.7) and (11.8) above, with $X_i = F^i x$. Each of the quantities appearing in (11.12) is either specified in advance of the measurements or may be evaluated once the measurements are obtained. Thus $\hat{F}_n^{(1)}$ is well defined. A key feature of $\hat{F}_n^{(1)}$ and other estimates in the chapter, is the adaptive, data-dependent choice of the partition index k_n . A version of the following result appears in Nobel and Adams [28].

Proposition 11.1. *For every $\mu \in \mathcal{D}(\mu_0, M)$ and every μ -preserving ergodic transformation $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$, the estimates $\hat{F}_n^{(1)}$ defined in (11.11)-(11.12) are such that*

$$\int \|\hat{F}_n^{(1)} - F\|^2 d\mu \rightarrow 0 \text{ as } n \rightarrow \infty$$

for μ -almost every initial vector $x \in \mathbb{R}^d$.

Overview of Proof. The proof of Proposition 11.1 proceeds roughly as follows. It is first shown, using the ergodic theorem and some elementary approximation arguments, that the partition index $k(n)$ chosen according to (11.12) will tend to infinity with n , if $\mu \in \mathcal{D}(\mu_0, M)$. Then inequality (11.12) is used to show that the estimates $\{\hat{F}_n^{(1)} : n \geq 1\}$ form, with probability one, a Cauchy sequence in $L_1(\mu_0)$. As $L_1(\mu_0)$ is complete, the estimates must converge in $L_1(\mu_0)$ (and hence also in $L_1(\mu)$) to some integrable limit F^* . Finally, it is shown that $F^* = F$ with probability one. Note that, in contrast with standard histogram estimates, the partition index $k(n)$ may not increase monotonically with n , nor need it grow at any pre-specified rate. It is precisely this flexibility that enables the estimates to perform effectively under very weak assumptions.

11.4.2 Remarks and a Corollary

The estimates $\hat{F}_n^{(1)}$ of Proposition 11.1 are multivariate histograms. In some cases kernel type estimates with similar properties can be defined, but we have not pursued this line of inquiry. Note that both μ_0 and M must be known before the estimate $\hat{F}_n^{(1)}$ can be computed.

Bosq and Guégan [7] consider the estimation of uniform mixing continuous maps F using kernel density estimators. The assumptions here concern the measure preserved by F ; no conditions are placed on the regularity of F or on its mixing properties.

Under additional conditions on the derivative $d\mu/d\mu_0$, one can strengthen the mode of convergence in Proposition 11.1 and obtain estimates consistent for individual sequences and non-ergodic transformations. See [2] for more details. When consistent estimates of the invariant density $f = d\mu/d\mu_0$ are available, the assumption that $\mu \in \mathcal{D}(\mu_0, M)$ can be dropped. In the special case when $d = 1$ and μ_0 is equal to Lebesgue measure, the density estimates of Nobel, Morvai, and Kulkarni [29] can be used to obtain consistent estimates of every ergodic transformation $F : \mathbb{R} \rightarrow \mathbb{R}$ which preserves a measure μ having compact support and a density $f = d\mu/dx$ whose variation is less than a known constant.

The family of transformations F for which the estimates $\hat{F}_n^{(1)}$ are consistent may be quite large. To illustrate this, let $d = 1$ and let μ_0 be Lebesgue measure on $[0, 1]$. Setting $M = 1$, the following corollary of Proposition 11.1 is immediate.

Corollary 11.1. *For every ergodic Lebesgue measure preserving map $F : [0, 1] \rightarrow [0, 1]$ the estimates defined in (11.11)-(11.12) are such that $\int_0^1 |\hat{F}_n^{(1)} - F|^2 du \rightarrow 0$ as $n \rightarrow \infty$ for almost every initial value $x \in [0, 1]$.*

Among the ergodic Lebesgue measure preserving transformations of $[0, 1]$ there is an uncountable sub-family \mathcal{S} with the property that no two transformations in \mathcal{S} are isomorphic. The histograms $\hat{F}_n^{(1)}$ give consistent estimates of each transformation in \mathcal{S} from almost every one of its trajectories.

11.5 Ergodic Systems with Dynamical Noise

Recall that a system is said to obey a dynamical noise model if its state evolves according to the recursion

$$X_{i+1} = F(X_i) + Z_{i+1} \tag{11.13}$$

where $X_0 \in \mathbb{R}^d$ is the initial state of the system, $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a fixed map, and

$$Z_1, Z_2, \dots \in \mathbb{R}^d \text{ are I.I.D., independent of } X_0, \text{ and } EZ_i = 0. \tag{11.14}$$

Estimation of F from ergodic measurements obeying the dynamical noise model (11.13)–(11.14) is considered later, first for continuous maps, and then for arbitrary maps with absolutely continuous noise.

11.5.1 Continuous Maps

When F is assumed to be continuous, smoothing the measurements with a suitable histogram and then employing a nearest-neighbor-type rule will give uniformly consistent estimates of F .

Description of the estimate: Let π_1, π_2, \dots be nested, finite partitions of \mathbb{R}^d satisfying (11.6). Let \bar{l}_n be the largest $l \geq 0$ such that $\sum_{j=0}^l |\pi_j| \leq n$, where $|\pi_l|$ denotes the number of cells in π_l . Given measurements X_0, \dots, X_{n-1} , define

$$\tilde{F}_n(u) = \frac{\sum_{i=0}^{n-1} X_{i+1} I\{X_i \in \pi_{l(n)}[u]\}}{\sum_{i=0}^{n-1} I\{X_i \in \pi_{l(n)}[u]\}},$$

where $l(n)$ is the largest $l \leq \bar{l}_n$ such that for each cell $A \in \pi_l$ either

$$\sum_{i=0}^{n-1} I\{X_i \in A\} \geq \frac{n}{\log n} \quad \text{or} \quad \sum_{i=0}^{n-1} I\{X_i \in A\} = 0.$$

Define the estimate

$$\hat{F}_n^{(2)}(u) = \tilde{F}_n(X_{\gamma(u)}) \quad \text{where} \quad \gamma(u) = \underset{0 \leq j \leq n-1}{\operatorname{argmin}} \|u - X_j\|.$$

Thus $\hat{F}_n^{(2)}$ is a nearest-neighbor estimate based on the histogram \tilde{F}_n .

Theorem 11.1. *Let $\{X_i : i \geq 0\}$ be a stationary ergodic sequence obeying the dynamical noise model (11.13)–(11.14), and let μ be the distribution of X_i . If F is continuous and Z_i is bounded, then for every compact set Λ contained in the support of μ ,*

$$\sup_{u \in \Lambda} \|\hat{F}_n^{(2)}(u) - F(u)\| \rightarrow 0$$

with probability one as $n \rightarrow \infty$. In particular, $\int \|\hat{F}_n^{(2)} - F\|^2 d\mu \rightarrow 0$ with probability one as $n \rightarrow \infty$ if F is bounded.

Remark 11.1. The assumption that the noise Z_i is bounded may be replaced by suitable tail conditions. When the distribution of Z_i has a continuous density, a similar result can be deduced from work of Delecroix [10] (see Györfi *et al.* [17], Theorem 3.5.1).

11.5.2 Absolutely Continuous Noise

When the map F in (11.13) is not continuous, the estimates $\hat{F}_n^{(2)}$ may fail to be consistent. However, if X_i has a density f , then it is possible to estimate irregular maps F when both f and the noise variance $E Z^2$ are estimable. As is shown below, estimating these quantities is possible under very general conditions.

Density Estimation

Suppose that $\{X_i : i \geq 0\}$ is a stationary ergodic sequence obeying the dynamical noise model (11.13)–(11.14). Let μ be the distribution of the observed X_i and let ν be the distribution of the perturbations Z_i . Equation (11.13) implies that

$$\mu = (\mu \circ F^{-1}) * \nu \tag{11.15}$$

where $*$ denotes convolution. It follows from (11.15) that if the noise distribution ν has a density with respect to Lebesgue measure λ on \mathbb{R}^d , then the same is true of the distribution μ of X_i . In other words, if $d\nu/d\lambda$ exists, so does $f = d\mu/d\lambda$.

Györfi [15] and later Györfi and Masry [16] proposed and studied recursive kernel density estimates that can be used to estimate $f = d\mu/d\lambda$. Their estimates are of the form

$$\hat{f}_n(x) = \frac{1}{n} \sum_{i=0}^{n-1} \frac{1}{h_i^d} K\left(\frac{x - X_i}{h_i}\right), \tag{11.16}$$

where $K : \mathbb{R}^d \rightarrow \mathbb{R}$ is a non-negative kernel such that $\int K dx = 1$, and such that for each $x \in \mathbb{R}^d$ the function $K(cx)$ is non-increasing in $0 < c < \infty$.

Suppose the bandwidth sequence is of the form $h_i = i^{-b}$ for some fixed $b \in (0, 1/d)$. It is shown in Theorem 3.2 of Györfi and Masry [16] that the estimates \hat{f}_n are L_1 -consistent for every ergodic process $\{X_i : -\infty < i < \infty\}$ such that the conditional distribution of X_1 given X_0, X_{-1}, \dots is absolutely continuous with probability one. The next theorem follows directly from this result.

Theorem 11.2. *Let $\{X_i : i \geq 0\}$ be a stationary ergodic sequence obeying the dynamical noise model (11.13)–(11.14), and let μ be the distribution of X_i . If the distribution ν of Z_i has a density, then $f = d\mu/d\lambda$ exists and $\int |\hat{f}_n - f| dx \rightarrow 0$ with probability one.*

Estimation of Noise Variance

Although the problems appear to be unrelated, it turns out that consistent estimates of the noise variance $E\|Z\|^2$ can be used to obtain consistent estimates of the map F governing a dynamical noise model. Indeed, estimating the noise variance may be of independent interest. Let π_1, π_2, \dots be nested, finite partitions of \mathbb{R}^d satisfying (11.6). Let $m(n)$ be the largest integer m such that $|\pi_m| \leq \log n$, and define

$$\hat{\Gamma}_n = \frac{1}{n-1} \sum_{i=0}^{n-2} \|G_{m(n),n}(X_i) - X_{i+1}\|^2 \tag{11.17}$$

where $G_{k,n}$ is the multivariate histogram defined in equation (11.7).

Proposition 11.2. *For every bounded stationary ergodic sequence $\{X_i : i \geq 0\}$ obeying the dynamical noise model (11.13)–(11.14) the estimates $\hat{\Gamma}_n \rightarrow E\|Z\|^2 = \text{Var}(Z)$ with probability one.*

Remark 11.2. If measurements X_0, X_1, \dots are generated by a dynamical noise model then the variance ratio

$$\kappa = \text{Var}(Z) / \text{Var}(X) = \text{Var}(Z) / (\text{Var}(F(X)) + \text{Var}(Z))$$

measures the size of the noise in relation to the overall variability of the dynamics. It indicates (in a rough sense) the extent to which the measurements are deterministic or random. If $\kappa = 0$, then each Z_i is identically zero and the measurements constitute a trajectory of F . On the other hand, if $\kappa = 1$ then $F(X_i)$ is constant, and X_0, X_1, \dots are I.I.D. with mean $F(X_0)$. It is interesting then that one can estimate κ directly from the measurements. In particular, if S_{n-1}^2 is the empirical variance of X_0, \dots, X_{n-1} , then $\hat{\Gamma}_n / S_{n-1}^2$ is a consistent estimate of κ .

Estimation of General Maps

Using the estimates \hat{f}_n and $\hat{\Gamma}_n$ defined earlier to select a suitable partition, one may obtain histogram estimates consistent for general measurable maps F .

Description of the Estimate. Let $\epsilon_1, \epsilon_2, \dots$ be any sequence of positive numbers tending monotonically to zero. Given measurements X_0, \dots, X_{n-1} , let the density estimate \hat{f}_n and variance estimate $\hat{\Gamma}_n$ be defined as in (11.16) and (11.17), respectively. Define the multivariate histogram

$$\hat{F}_n^{(3)}(u) = \frac{\sum_{i=0}^{n-1} X_{i+1} I\{X_i \in \pi_{r(n)}[u]\}}{\sum_{i=0}^{n-1} I\{X_i \in \pi_{r(n)}[u]\}},$$

where $r(n)$ is the largest integer $r \geq 1$ such that

$$\int \|G_{i,n} - G_{j,n}\| \hat{f}_n dx \leq 2(\Delta_{i,n}^2 - \hat{\Gamma}_n)_+^{1/2} + 5\epsilon_i \text{ for } 1 \leq i \leq j \leq r. \tag{11.18}$$

Here $G_{i,n}$ and $\Delta_{i,n}$ are defined as in equations (11.7) and (11.8), and the integral is taken with respect to d -dimensional Lebesgue measure. Each of the quantities appearing in (11.18) may be evaluated once the measurements are obtained, so $\hat{F}_n^{(3)}$ is well defined. By combining Theorem 11.2 and Proposition 11.2 with a modification of Theorem 1 in Nobel and Adams [28], one obtains the following result.

Theorem 11.3. *Let $\{X_i : i \geq 0\}$ be a stationary ergodic sequence obeying the dynamical noise model (11.13)–(11.14), where F is assumed only to be measurable. If X_i is bounded and the distribution of Z_i has a density, then*

$$\int \|\hat{F}_n^{(3)} - F\|^2 d\mu \rightarrow 0$$

with probability one, where μ is the distribution of X_i .

Remark: The proof of Theorem 11.3 follows the same sequence of steps as the proof of Proposition 11.1. Under the conditions of the theorem the partition index $r(n)$ used to define $\hat{F}_n^{(3)}$ will tend to infinity with the sample size n . However the index $r(n)$ may not increase monotonically with n , nor will it necessarily grow at a pre-specified rate.

Acknowledgments

The author wishes to thank D. Guégan and an anonymous referee for their helpful comments and suggestions. This work was supported in part by NSF Grant DMS-9501926 and Grant DMS-9971964.

References

- [1] Adams, T.M. and Nobel, A.B. (1999). Finitary reconstruction of a measure preserving transformation. Preprint.
- [2] Adams, T.M. (1999). Families of ergodic processes without consistent density or regression estimates. Preprint.
- [3] Aeyels, D. (1981). Generic observability of differential systems. *SIAM J. Control and Optimization*, 19:595–603.
- [4] Arnold, L. (1995). Random dynamical systems. In *Dynamical Systems*, R.Johnson, editor, Springer Lecture Notes in Mathematics 1609.
- [5] Beran, J. (1994). *Statistics for Long-Memory Processes*. Chapman and Hall, New York.
- [6] Bosq, D. (1998). *Nonparametric Statistics for Stochastic Processes: Estimation and Prediction*, 2nd edition. Springer Lecture Notes in Statistics, v.110.
- [7] Bosq, D. and Guégan, D. (1995). Nonparametric estimation of the chaotic function and the invariant measure of a dynamical system. *Stat. and Prob. Let.*, 25:201–212.
- [8] Casdagli, M. (1989). Nonlinear prediction of chaotic time series. *Physica D*, 35:335–356.
- [9] Casdagli, M. (1992). Chaos and deterministic *versus* stochastic nonlinear modeling. *J. R. Stat. Soc. B*, 54:303–328.
- [10] Delecroix, M. (1987). *Sur l'estimation et la prévision non-paramétrique des processus ergodiques*. Ph.D. Thesis, University of Lille Flandres Artois, Lille, France.
- [11] Delecroix, M. and Rosa, A.C. (1996). Nonparametric estimation of a regression function and its derivatives under an ergodic hypothesis. *Nonparam. Statistics*, 6:367–382.
- [12] Denker, M., and Keller, G. (1986). Rigorous statistical procedures for data from dynamical systems. *J. Stat. Phys.*, 44:67–93.
- [13] Eckmann, J.-P., and Ruelle, D. (1985). Ergodic theory of chaos and strange attractors. *Rev. Mod. Phys.*, 57(3):617–656.
- [14] Farmer, J.D. and Sidorowich, J.J. (1987). Predicting chaotic time series. *Phys. Rev. Let.*, 59:845–848.
- [15] Györfi, L. (1981). Strongly consistent density estimate from ergodic sample. *J. Multivariate Analysis*, 11:81–84.
- [16] Györfi, L. and Masry, E. (1990). The L_1 and L_2 strong consistency of recursive kernel density estimation from dependent samples. *IEEE Trans. Inform. Theory*, 36:531–539.
- [17] Györfi, L., Härdle, W., Sarda, P., and Vieu, P. (1989). *Nonparametric Curve Estimation from Time Series*. Springer-Verlag, Berlin.
- [18] Hofbauer, F., and Keller, G. (1982). Ergodic properties of invariant measures for piecewise monotonic functions. *Mathematische Zeitschrift*, 180:119–140.
- [19] Isham, V. (1993). Statistical aspects of chaos: a review. In *Networks and Chaos – Statistical and Probabilistic Aspects*, Barndorff-Nielsen, O.E., Jensen, J.L., and Kendall, W.S. eds, Chapman and Hall, London.
- [20] Jensen, J.L. (1993). Chaotic dynamical systems with a view towards statistics: a review. In *Networks and Chaos – Statistical and Probabilistic Aspects*, Barndorff-Nielsen, O.E., Jensen, J.L., and Kendall, W.S. eds, Chapman and Hall, London.

- [21] Kostelich, E.J. and Yorke, J.A. (1990). Noise reduction: finding the simplest dynamical system consistent with the data. *Physica D*, 41:183-196.
- [22] Kifer, Y. (1986). *Ergodic Theory of Random Transformations*. Birkhäuser, Boston.
- [23] Lalley, S.P. (1999). Beneath the noise, chaos. *Annals of Statistics*, 27:461-479.
- [24] Lu, Z.-Q., and Smith, R. L. (1997). Estimating local Lyapunov exponents. *Fields Institute Communications*, 11:135-151.
- [25] Mayer, D.H. (1984). Approach to equilibrium for locally expanding maps in R^k . *Communications in Mathematical Physics*, 95:1-15.
- [26] Meyn, S.P. and Tweedie, R.L. (1993). *Markov Chains and Stochastic Stability* Springer Verlag, London.
- [27] Nobel, A.B. (1999). Limits to classification and regression estimation from ergodic processes. *Annals of Statistics* 27:262-273.
- [28] Nobel, A.B., and Adams, T.M. (1999). On regression estimation from ergodic samples with additive noise. Technical Report #2363, Department of Statistics, University of North Carolina, Chapel Hill. Submitted for publication.
- [29] Nobel, A.B., Morvai, G., and Kulkarni, S. (1998). Density estimation from an individual numerical sequence. *IEEE Trans. Info. Theory* 44:537-541.
- [30] Nummelin, E. (1984). *General Irreducible Markov Chains and Non-Negative Operators*. Cambridge Univ. Press.
- [31] Nychka, D., Ellner, S., Gallant, A.R., and McCaffrey, D. (1992). Finding chaos in noisy systems. *J. R. Stat. Soc. B*, 54:399-426.
- [32] Petersen, K. (1989). *Ergodic Theory*. Cambridge Univ. Press.
- [33] Rosenblatt, M. (1991). *Stochastic Curve Estimation*. NSF-CBMS Regional Conference Series in Probability and Statistics, Inst. Math. Stat., Hayward, CA..
- [34] Sauer, T., Yorke, J.A., and Casdagli, M. (1991). Embedology. *J. Stat. Phys.*, 65:579-616.
- [35] Takens, F. (1980). Detecting strange attractors in turbulence In *Dynamical Systems and Turbulence, Warwick 1980*, D.A. Rand and L.-S. Young, editors. Springer Lecture Notes in Mathematics 898.
- [36] Tong, H. (1990). *Nonlinear Time Series: a Dynamical System Approach*. Oxford University Press.
- [37] Walters, P. (1981). *An Introduction to Ergodic Theory*. Springer, New York.
- [38] Yakowitz, S. (1989). Nonparametric density and regression estimation for Markov sequences without mixing assumptions. *J. Multivar. Anal.*, 30:124-136.
- [39] Yakowitz, S. (1993). Nearest neighbor regression estimation for null-recurrent Markov time series. *Stoc. Proc. Appl.*, 48:311-318.
- [40] Yakowitz, S., Györfi, L., Kieffer, J., and Morvai, G. (1997). Strongly-consistent nonparametric estimation of smooth regression functions for stationary ergodic sequences. Under revision, *J. Multivar. Anal.*.
- [41] Yakowitz, S. and Heyde, C. (1998). Long range dependency effects with implications for forecasting and queuing inference. Preprint, submitted for publication.

Chapter 12

Extracting Dynamical Behavior via Markov Models

Gary Froyland

ABSTRACT Statistical properties of chaotic dynamical systems are difficult to estimate reliably. Using long trajectories as data sets sometimes produces misleading results. It has been recognized for some time that statistical properties are often stable under the addition of a small amount of noise. Rather than analyzing the dynamical system directly, we slightly perturb it to create a Markov model. The analogous statistical properties of the Markov model often have “closed forms” and are easily computed numerically. The Markov construction is observed to provide extremely robust estimates and has the theoretical advantage of allowing one to prove convergence in the noise $\rightarrow 0$ limit and produce rigorous error bounds for statistical quantities. We review the latest results and techniques in this area.

12.1 Introduction and Basic Constructions

Suppose that we find ourselves presented with a discrete time¹ dynamical system, and we would like to perform some (mainly ergodic-theoretic) analysis of the dynamics. We are not concerned with the problem of embedding, nor with the extraction of a system from time series. We assume that we have been presented with a dynamical system and do not question its validity.

Any analysis of a dynamical system involving *average* quantities requires a reference measure with which to average contributions from different regions of phase space. Often the measure that one wishes to use is the probability measure described by the distribution of a typical long trajectory of the system; it is commonly called the *natural measure* or *physical measure* of the system.

¹Similar constructions for flows are possible by considering the “time- t ” map.

12.1.1 What Do We Do?

This chapter discusses a method of modeling the dynamics by a finite state Markov chain. Naturally such a model contains much less information than the original dynamical system. However, this simplification of the dynamics allows the exact computation of many properties of the Markov chain which correspond to important indicators and properties of the original dynamics: for example, finding invariant sets, obtaining invariant measures, calculating rates of mixing and the spectrum of eigenvalues of transfer operators, computing means and variances of recurrence times, and estimating Lyapunov exponents; all of these calculations are *exact*² for the Markov chain. We hope that although we are throwing away a lot of information in our Markov model, we retain the essential properties of the original system. The questions then are: (i) Are these quantities computed for the Markov model good estimators of the corresponding quantities for the original system, and (ii) how best to define these Markov models for various sorts of systems.

12.1.2 How Do We Do This?

We describe the fundamental construction of the modeling process. Consider a dynamical system (M, T) defined by a map $T : M \rightarrow M$, where M is a compact subset of \mathbb{R}^d . Partition the phase space into a finite number of connected sets $\{A_1, \dots, A_n\}$ with nonempty interior. Usually, this partition will take the form of a regular grid covering the phase space M . We now completely ignore any dynamics that occurs inside each individual partition set, and focus only on the coarse-grained dynamics displayed by the evolution of whole partition sets. To form our Markov model, we identify each set A_i with a state i of our n -state Markov chain. We construct an $n \times n$ transition matrix P , where the entry P_{ij} is to be interpreted as:

$$P_{ij} = \text{the probability that a typical point in } A_i \text{ moves into } A_j \quad (12.1)$$

under one iteration of the map T .

We now meet the notion of *typicality* and begin to impinge on ergodic-theoretic ideas. Leaving formality for the moment, we shall assume that the trajectories $\{x, Tx, T^2x, \dots\}$ of Lebesgue-almost-all initial points $x \in M$ have the same distribution on M . This distribution may be represented as a probability measure, denoted by μ . Now, in light of (12.1), the most natural definition of P_{ij} is

$$\tilde{P}_{ij} = \frac{\mu(A_i \cap T^{-1}A_j)}{\mu(A_i)} \quad (12.2)$$

²A further approximation must be introduced for the calculation of Lyapunov exponents.

since $\tilde{P}_{ij} = \text{Prob}_\mu\{Tx \in A_j | x \in A_i\}$. Unfortunately, the natural measure μ is usually unknown and for want of a better alternative, we compute a (slightly) different transition matrix using normalized Lebesgue measure m instead of μ :

$$P_{ij} = \frac{m(A_i \cap T^{-1}A_j)}{m(A_i)}. \quad (12.3)$$

Several numerical techniques have been put forward regarding the computation of P ; see Section 12.6.1.

12.1.3 Why Do We Do This?

The alternative to Markov modeling of the dynamics via some coarse graining is to simply simulate very long orbits of the dynamical system. For the purposes of this discussion, we restrict ourselves to the problem of approximating the probability measure μ that describes the asymptotic distribution of almost all trajectories.

Given a long orbit $\{x, Tx, \dots, T^{N-1}x\}$, one is implicitly approximating the long-term distribution μ by the finite-time distribution $\mu_N(x) := \frac{1}{N} \sum_{k=0}^{N-1} \delta_{T^k x}$, where δ_x is the delta-measure at $x \in M$. This is certainly a simple way to compute an approximate invariant measure (or long-term distribution), as it does not involve any complicated matrix constructions; one just iterates one's map. There are however, drawbacks to this simple approach. It is possible that orbits display transient (nonequilibrium) behavior for lengthy periods of time before settling into a more regular (statistically speaking) mode. Thus, by following a simulated orbit for a finite time, there is the risk that one is only observing this transient behavior and not the true asymptotic behavior of the system. There is also the problem of computer roundoff; try to find the long-term distribution of the tent map or the circle doubling map by iterative simulation (all trajectories are attracted to 0 in finite time). These are extreme cases, but the potential compounding inaccuracy of long computer-generated orbits should not be forgotten. Let's be generous though, and assume that our approximation $\mu_N(x)$ actually does (weakly) converge to μ . How fast does this happen? What is the convergence rate with respect to the length of the orbit N ? Can one produce rigorous error bounds for the difference between the distributions $\mu_N(x)$ and μ ? For the most part, the answer to each of these questions is "We don't know yet". In toy cases, one can produce extremely crude probabilistic *lower* bounds for the error, of the form C/\sqrt{N} , but this is not really satisfactory.

Our method of Markov modeling attempts to overcome all of these difficulties. Transient effects are completely removed as we model the system as a Markov chain and its asymptotic behavior may be computed *exactly*. Computer roundoff is not so much of a problem, as we are now only computing a *single iterate* of the dynamics rather than a compounding sequence

of iterations. The constructions also permit a rigorous treatment of questions like rates of convergence, error bounds, and even just whether convergence occurs at all. Finally, from the practical point of view, the method of Markov coarse graining is often very efficient computationally, producing better answers in less time.

This discussion has primarily been aimed at the computation of invariant measures, but it also applies to the computation of other dynamical indicators such as the rate of decay of correlations, Lyapunov exponents, and statistics of recurrence times. The rate of decay of correlations and the isolated spectrum of transfer operators in particular are notoriously difficult to estimate via iterative techniques.

12.2 Objects and Behavior of Interest

In this section we give a summary of the definitions and properties of the objects we wish to estimate. Additional applications are outlined in Section 12.5.

12.2.1 Invariant Measures

The fundamental object for ergodic-theoretic concepts is an *ergodic invariant probability measure* μ . One should think of $\mu(B)$ as representing the proportion of “mass” that is contained in the subset $B \subset M$. The *invariance condition* of μ is $\mu = \mu \circ T^{-1}$. This is a generalized “conservation of mass” equality; the mass distribution described by μ is preserved under the action of T , in the same way that area-preserving maps preserve Lebesgue measure. *Ergodicity* is the measure-theoretic equivalent of topological transitivity. If two sets A and B both have positive μ -measure, then there is an $N \geq 0$ such that $\mu(A \cap T^{-N}B) > 0$; in other words, any region of positive μ -mass may be evolved forward to intersect any other region of positive μ -mass, with the intersection also having positive mass. Ergodic systems are indecomposable in the sense that μ -almost all starting points produce orbits with the same asymptotic distribution. The Birkhoff theorem [4, 52] tells us that the frequency with which orbits of μ -almost all starting points visit a set B is equal to the μ -measure of B ; formally,

$$\lim_{N \rightarrow \infty} (1/N) \#\{0 \leq k \leq N-1 : T^k x \in B\} = \mu(B) \quad (12.4)$$

for μ -almost all $x \in M$. However, as in the case of dissipative systems, there may be a thin invariant attracting set Λ with $\mu(\Lambda) = 1$, but $m(\Lambda) = 0$. Equation (12.4) gives us no information about orbits starting off this invariant set.

Definition 12.1. *An ergodic invariant probability measure is called a natural or physical measure, if*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} f(T^k x) \rightarrow \int_M f d\mu \tag{12.5}$$

for all continuous $f : M \rightarrow \mathbb{R}$ and Lebesgue-almost-all $x \in M$.

An alternative way of phrasing the definition of a physical measure is to state that the measure $\mu_N(x)$ of Section 12.1.3 converges weakly to μ as $N \rightarrow \infty$ for Lebesgue-almost-all $x \in M$. When talking about physical measures, Equation (12.4) may be strengthened to:

Corollary 12.1.

$$\lim_{N \rightarrow \infty} \frac{1}{N} \#\{0 \leq k \leq N - 1 : T^k x \in B\} = \mu(B) \tag{12.6}$$

for any subset $B \subset M$ with $\mu(\partial B) = 0$, and for Lebesgue-almost-all $x \in M$.

Corollary 12.1 says that “randomly chosen” initial conditions will, with probability one, produce trajectories that distribute themselves according to the natural measure μ . Deterministic chaotic systems typically have infinitely many ergodic invariant measures (for example, a convex combination of δ measures on a periodic orbit), but only one physical measure.

Example 12.1. The linear mapping of the torus $T : \mathbb{T}^2 \rightarrow \mathbb{T}^2$ defined by $T(x, y) = (2x + y, x + y) \pmod{1}$ has infinitely many ergodic invariant measures. For example, $\delta_{(0,0)}$ and $m =$ Lebesgue measure. Of all the ergodic invariant probability measures, only m satisfies (12.5).

12.2.2 Invariant Sets

We approach invariant sets in a rather roundabout way; we will describe them as the support of the physical measure, $\text{supp } \mu$, where $\text{supp } \mu$ is the smallest closed set having μ measure of 1 (equivalently, $x \in \text{supp } \mu$ iff every open neighborhood of x has positive μ measure). It is easy to show that $T(\text{supp } \mu) = \text{supp } \mu$ and $\text{supp } \mu \subset T^{-1}(\text{supp } \mu)$.

The reason for choosing $\text{supp } \mu$ as the distinguished invariant set, rather than some invariant set defined via topological conditions is as follows. Let x be an arbitrary point in $\text{supp } \mu$, and $B_\epsilon(x)$ be an open neighborhood of size ϵ about x . By (12.5), orbits of *Lebesgue-almost-all initial points* visit $B_\epsilon(x)$ infinitely often and with a positive frequency of visitation. Meanwhile regions away from $\text{supp } \mu$ are visited with frequency 0. As we wish to find the invariant set which appears on computer simulations of trajectories, it makes sense to consider only the support of μ and not some larger topologically invariant set.

12.2.3 Decay of Correlations

The measure-theoretic analogy of topological mixing is (measure-theoretic) mixing. We say that a system (T, μ) is *mixing* if $\mu(A \cap T^{-N}B) \rightarrow \mu(A)\mu(B)$ as $N \rightarrow \infty$ for any pair of measurable subsets $A, B \subset M$. This condition says that the probability that a point x lies in a set B at time $t = 0$ and then moves to a set A at time $t = N$ (for large N) is roughly the product of the measures of the sets A and B . That is, for large N , the two events $\{x \in B\}$ and $\{T^N x \in A\}$ become statistically independent, or decorrelated. For dynamical systems with some degree of smoothness, this loss of dependence is often studied via the formula:

$$C_{f,g}(N) := \left| \int f \circ T^N \cdot g \, d\mu - \int f \, d\mu \cdot \int g \, d\mu \right|, \quad (12.7)$$

where $f \in L^\infty(M, m)$ and $g : M \rightarrow \mathbb{R}$ has some smoothness properties. If one thinks of the functions $f, g : M \rightarrow \mathbb{R}$ as “physical observables” (output functions giving numerical information on some physical quantities of the system), then $C_{f,g}(N)$ quantifies the correlation between observing g at time t and f at time $t + N$. If $f = g$, we obtain what is commonly known as the *autocorrelation function*. For many chaotic systems, it is observed that $C_{f,g}(N) \rightarrow 0$ at a geometric rate, and it is of interest to estimate this rate. The rate of decay can be interpreted variously as providing information on how quickly the system settles into statistically regular behavior, how quickly transient behavior disappears, and how quickly physical observables become decorrelated. For all of these interpretations, the physical measure μ is central as it provides the reference measure that describes statistical equilibrium for the system. Correlation decay has strong connections with *transfer operators* (or *Perron-Frobenius operators*) and the spectrum and corresponding eigenfunctions of these operators. We postpone further discussion of these objects until Section 12.3.3.

12.2.4 Lyapunov Exponents

The Lyapunov exponents of a (piecewise differentiable) dynamical system describe the average local rates of expansion and contraction in phase space that are felt along typical trajectories. We may define them as the various values of λ that are observed³ in the following expression, as v is varied over \mathbb{R}^d

$$\lambda := \lim_{N \rightarrow \infty} \frac{1}{N} \log \|D_x T^N(v)\|. \quad (12.8)$$

As our dynamical system satisfies Equation (12.5), it is immediate that Lebesgue-almost-all trajectories will produce the same d values of λ , and

³For a d -dimensional system, λ may only take on at most d different values.

these values will be called *the* Lyapunov exponents⁴ of T .

Later we will be discussing Lyapunov exponents of random matrix products. These can be defined in an analogous way.

12.2.5 Mean and Variance of Return Times

The statistics of return times of orbits to subsets of phase space have not received a great deal of attention in the dynamical systems literature. A recent exception is [59], where regularity properties of return times are used to prove existence of physical measures and characterize rates of decay of correlations as exponential or algebraic. Suppose that one is given a subset $B \subset M$ with $\mu(B) > 0$. Every time a typical trajectory enters B , we note the number of iterations since the orbit last entered B . In this way, we obtain a sequence of times t_0, t_1, \dots such that $T^{t_i}x \in B$, $i \geq 0$. We may now ask what the mean and variance of this sequence of times are. Formally, we define a function $R : M \rightarrow \mathbb{Z}^+$ by $R(x) = \inf\{k \geq 1 : T^k x \in B\}$, and use this return time function to define an induced map $\tilde{T}x = T^{R(x)}x$, where $\tilde{T} : B \rightarrow B$. It is straightforward to show that $\mu|_B$ (defined by $\mu|_B(C) = \mu(C)/\mu(B)$ for $C \subset B$) is \tilde{T} -invariant. Therefore, we can define the mean return time

$$\mathbb{E}_{\mu|_B}(R) = \int_B R(x) d\mu|_B(x), \quad (12.9)$$

and the variance of the return times as

$$\text{var}_{\mu|_B}(R) = \mathbb{E}_{\mu|_B}(R^2) - \mathbb{E}_{\mu|_B}(R)^2. \quad (12.10)$$

By Kac's theorem [34, 52], $\mathbb{E}_{\mu|_B}(R) = 1/\mu(B)$ provided μ is ergodic. A corresponding simple formula for the variance is not known.

12.3 Deterministic Systems

In the first of two parts, we focus on deterministic systems. We show how each of the objects outlined in Section 12.2 will be approximated using the Markov model. We begin each subsection by simply outlining the constructions and computations one performs in practice. At the end of each subsection, we state situations where rigorous results are obtained. While the Markov systems appear to retain many of the salient features of the original system, the smooth dynamics is completely lost, so it is not at all straightforward to prove strong approximation results.

⁴Lyapunov exponents depend very much on a reference measure (in this case the physical measure) in order to be defined.

12.3.1 Basic Constructions

The region M is covered with a collection of connected sets $\mathfrak{P}_n = \{A_{n,1}, \dots, A_{n,n}\}$ with nonempty interiors. This covering has the properties that (i) $\bigcup_{i=1}^n A_{n,i} = M$ and (ii) $\text{Int } A_{n,i} \cap \text{Int } A_{n,j} = \emptyset$ for $i \neq j$, where $\text{Int } A$ denotes the interior of A . We think of each A_i as being a closed set, so that we do not have a partition in the strict sense, as the sets in the covering share boundaries.

We construct the matrix⁵

$$P_{n,ij} = \frac{m(A_{n,i} \cap T^{-1}A_{n,j})}{m(A_{n,i})}$$

as in (12.3).

Clearly, the quality of our Markov model depends heavily on the choice of partition, and it stands to reason that finer partitions produce better models (this will soon be made more precise). We will frequently produce a sequence of Markov models, each constructed from successively finer partitions, until we are satisfied with the accuracy of the estimates produced by the model. Sometimes one can produce better models by clever refinement strategies; these are discussed in Section 12.6.2.

12.3.2 Invariant Measures and Invariant Sets

ROUGH IDEA: *The invariant density of the Markov chain approximates the physical invariant measure of T .*

REQUIRED COMPUTATION: *Calculate the fixed left eigenvector of P .*

We fix our partition $\{A_{n,1}, \dots, A_{n,n}\}$ and calculate P_n and p_n , normalizing p_n so that $\sum_{i=1}^n p_{n,i} = 1$. The value $p_{n,i}$ is the weight given to state i by the stationary distribution of the Markov chain. Since state i represents the set $A_{n,i}$ in our smooth space, we define an approximate invariant measure μ_n by assigning $\mu_n(A_{n,i}) = p_{n,i}$. Within the set $A_{n,i}$, we distribute the mass in any way we like. A common method is simply to spread the weight $p_{n,i}$ uniformly over $A_{n,i}$, so that formally,

$$\mu_n(B) := \sum_{i=1}^n \mu_n(A_{n,i}) \frac{m(A_{n,i} \cap B)}{m(A_{n,i})}. \quad (12.11)$$

As we increase the number of partition sets n through refinement, some of these refined sets will be given zero measure by μ_n as an indication that trajectories spend no, or almost no, time in these regions. Thus we expect

⁵This construction is often called Ulam's method, as it was first proposed in [56] to use the matrix P_n to approximate invariant densities of interval maps.

that $\text{supp } \mu_{n'} \subset \text{supp } \mu_n$ for $n' > n$ (this can be made rigorous if T is a topological attractor [46]).

Further, since M is compact, the space of Borel probability measures on M (denoted $\mathcal{M}(M)$) is compact with respect to the weak topology. We may therefore continue this refinement procedure “forever” and extract a limiting measure μ^* as

$$\mu^* = \lim_{n \rightarrow \infty} \mu_n, \quad (12.12)$$

taking a convergent subsequence if necessary to define the limit. It is always assumed that $\max_{1 \leq i \leq n} \text{diam } A_{n,i} \rightarrow 0$.

We have the following results concerning these approximations.

Theorem 12.1. *Suppose that $T : M \rightarrow M$ is continuous and has a physical measure μ . Let μ_n denote the approximate invariant measures and μ^* a weak limit of this sequence. Denote by S the intersection $\bigcap_{n=n_0}^{\infty} \text{supp } \mu_n$. Then*

1. μ^* is T -invariant.
2. $\text{supp } \mu^* \subset \text{supp } \mu_n$ for all $n \geq 0$, and therefore $\text{supp } \mu^* \subset S$.
3. $\text{supp } \mu \subset \text{supp } \mu_n$ for all $n \geq 0$, and therefore $\text{supp } \mu \subset S$.

Proof:

1. By noticing that the Markov models are small random perturbations [39, 41] of T , it is relatively simple to prove that μ^* is T -invariant [14, 26].
2. This is a straightforward consequence of weak convergence.
3. This is proven in [45, 46].

□

The first result of Theorem 12.1 says that weak limits of the sequence of numerically computed measures μ_n are T -invariant. In other words, we are in fact approximating *some* T -invariant probability measure. However, there is the question of which invariant measure this is, as μ^* may be distinct from μ , the physical measure we wish to approximate. In this very general formulation, it is (as yet⁶) not possible to say that $\mu^* = \mu$.

Parts (ii) and (iii) of Theorem 12.1 say that at least the supports of the computed measures μ_n do contain the supports of both μ^* and μ , so that

⁶A sign which may be taken as promise, or simply a state of ignorance, is that the author does not know of a continuous dynamical system T with physical measure μ for which $\mu^* \neq \mu$ (using reasonably “regular” partitions and refinements which keep the partition sets as approximate d -dimensional “cubes” of approximately the same shape and size).

by stopping our computation at some finite n , we are not “losing” regions that are given positive measure by the physical measure μ .

Example 12.2. (The stiletto map). We introduce the Stiletto map [54] $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$.

$$T(x, y) = ((x + 1/3) \exp(-3x + 2) - 1/3 + y, 3x/10). \quad (12.13)$$

This map seems to possess chaotic dynamics on a fractal attracting set; see Figure 12.1. By selecting $M \subset \mathbb{R}^2$ to be a sufficiently small neighborhood of the




FIGURE 12.1. Trajectory of length 10^4 generated by the Stiletto map.

observed attracting set, it numerically appears that Lebesgue-almost-all $x \in M$ exhibit the invariant distribution described by the density of points in Figure 12.1. It is assumed that this distribution of points (in the infinite limit) describes the physical measure μ . We construct a Markov model using 1215 partition sets, where the sets are rectangles of equal size. The support of the resulting approximate invariant measure is shown in Figure 12.2, and the approximation itself is shown in Figure 12.3. We have used a relatively low number of partition sets for ease of viewing. Even for this crude model, there is good agreement between the distributions in Figures 12.1 and 12.3.

Rigorous Results

To state that $\mu^* = \mu$, we require further restrictions on the map T and the partition \mathfrak{P}_n ; in general, even the question of *existence* of a physical measure




FIGURE 12.2. The support of an approximate invariant measure ($\text{supp } \mu_{1215}$) shown as boxes.

for a given nonlinear system is still open. In one dimension, there is the classical existence result of Lasota and Yorke [43] that states that physical measures (probability measures with bounded densities) exist for piecewise C^2 interval maps $T : [0, 1] \rightrightarrows$ with $\inf_x |T'(x)| > 1$. Li [44] first proved that under the additional constraint that $\inf_x |T'(x)| > 2$, these invariant densities could be approximated using Ulam’s method in the sense that $\|\mu - \mu_n\|_{L^1} \rightarrow 0$ as $n \rightarrow \infty$. Since the publication of [44], there have been many variations of this basic result. In the setting of [44], Keller [36] proved $\|\mu - \mu_n\|_{L^1} = O(\log n/n)$. Recent results (often under additional “onteness” assumptions) have focused on explicit error bounds for the difference $\|\mu - \mu_n\|_{L^1}$; [17, 35, 49].

For higher-dimensional uniformly expanding systems, very roughly speaking, the papers of Boyarsky and Gora [27] and Ding [12] mirror those of [43] and [44]. There are several additional technical constraints on the map T and the partitions \mathfrak{P}_n that we do not discuss. Again under some ontteness conditions, Murray [48] applies the methods of [49] to provide error bounds for $\|\mu - \mu_n\|_{L^1}$.

For uniformly hyperbolic systems, the author [14, 18] shows that $\mu_n \rightarrow \mu$ weakly (resp. in L^1) when the physical measure μ is singular (resp. absolutely continuous), provided that the partitions \mathfrak{P}_n are Markov partitions.

A combination of theory and numerics [16] suggests that the convergence rate $O(\log n/n)$ holds in reasonable generality for systems with good mixing properties.




FIGURE 12.3. Representation of an approximate invariant (physical) measure (μ_{1215}) for the Stiletto map; darker regions indicate higher density and more mass.

12.3.3 Decay of Correlations and Spectral Approximation

ROUGH IDEA: The spectrum of the matrix P approximates the spectrum of the Perron-Frobenius operator.

REQUIRED COMPUTATION: Calculate the eigenvalues of P .

We begin by noting that we have an alternative formulation of (12.7) in terms of the Perron-Frobenius operator⁷ $\mathcal{P} : L^1(M, m) \rightarrow L^1(M, m)$.

Lemma 12.1. Let \mathcal{F} be a class of real-valued functions preserved⁸ by \mathcal{P} . Let $\sigma(\mathcal{P})$ denote the spectrum of \mathcal{P} when considered as an operator on \mathcal{F} , and set $r = \sup\{|z| : z \in \sigma(\mathcal{P}) \setminus \{1\}\}$. Then there is a constant $C < \infty$ such that $C_{f,g}(N) \leq Cr^N$ if $g \in \mathcal{F}$ and $f \in L^\infty$.

This result says that we may bound the rate of decay of correlations by the maximal nonunit spectral value for the operator \mathcal{P} ; see Figure 12.4 (upper right) for the typical spectral plot we have in mind. The Ulam matrix P_n may be thought of as a projection of \mathcal{P} onto a finite-dimensional space.

⁷See [6, 42] for definitions and properties of the Perron-Frobenius operator.

⁸For example, if T is C^γ , then $C^{\gamma-1}(M, \mathbb{R})$ is preserved by \mathcal{P} , and if $T : [0, 1] \rightarrow [0, 1]$ is a Lasota-Yorke map, then functions of bounded variation are preserved by \mathcal{P} .

Naively, then, we may think that the spectrum of the matrix P_n will approximate the spectrum of the Perron-Frobenius operator. This would be very useful, as it is simple to compute the spectrum of P_n since the matrix is very sparse and there are numerical routines to compute only eigenvalues which are large in magnitude (these are the ones that are principally of interest). Furthermore, the eigenfunctions of \mathcal{P} corresponding to large eigenvalues are also of interest, as they indicate the sorts of mass distributions that approach the equilibrium distribution (the physical measure) at the slowest possible rate. Perhaps we can also approximate these slowly mixing distributions (eigenfunctions of \mathcal{P}) as eigenvectors of the matrix P . If we can, there is the question of what these slowly mixing distributions represent. One generally thinks that the rate of mixing is determined by expansion properties of the dynamical system; that is, the more expansive (or more “chaotic”), the greater the rate of decay. But the existence of distributions which mix at a rate *slower* than that dictated by the minimal expansion rate of the system presents a seemingly paradoxical situation. Arguments in [7, 9, 11] suggest that these distributions describe “macroscopic structures” embedded within the dynamics, which exchange mass very slowly and work against the chaoticity.

Rigorous Results

There are two situations where these ideas can be made rigorous. First, for one-dimensional maps, there is the recent result of Keller and Liverani [37].

Theorem 12.2. *Let BV denote the space of functions of bounded variation on $[0, 1]$. Suppose $T : [0, 1] \rightarrow [0, 1]$ is an expanding, piecewise C^2 interval map, with $\inf_{x \in [0,1]} |T'(x)| = \alpha > 1$. Then isolated eigenvalues of $\mathcal{P} : BV \rightarrow BV$ outside the disk $\{|z| \leq 1/\alpha\}$ and the corresponding eigenfunctions are approximated by eigenvalues and the corresponding eigenvectors of P_n (eigenfunction convergence in the L^1 sense). The convergence rate for the eigenvectors to the eigenfunctions for eigenvalues $z \in (\alpha, 1)$ is $O(n^{-r})$, where $0 < r(z) < 1$, while the eigenvector approximating the invariant density ($z = 1$) converges like $O(\log n/n)$.*

Example 12.3. (The double wigwam map). We introduce the map $T : [0, 1] \rightarrow [0, 1]$ defined by

$$T(x) = \begin{cases} -2x + 1 - \sin(4\pi x)/30, & 0 \leq x < 1/4 \\ 3(x - 1/4) + 1/4 + \sin(4\pi(x - 1/4))/10, & 1/4 \leq x < 1/2 \\ 3(x - 1/2) - \sin(4\pi(x - 1/2))/10, & 1/2 \leq x < 3/4 \\ -2(x - 1) - \sin(-4\pi)/30, & 3/4 \leq x \leq 1, \end{cases}$$

the graph of which is shown in the upper-left frame of Figure 12.4. Theorem 12.2 will be used to show that the spectrum of $\mathcal{P} : BV \rightarrow BV$ contains a non-trivial isolated eigenvalue, and therefore a rate of decay *slower* than that

prescribed by the minimal expansion rate. As T is a *Lasota-Yorke* map, the classical result of [43] tells us that T possesses an invariant density (this is the density of the physical measure). The result of Li [44] tells us that this invariant density may be approximated by eigenvectors of the Ulam matrices. The bottom-left frame of Figure 12.4 shows a plot of an approximation of the invariant density using an equipartition of $[0, 1]$ into 512 sets. The upper-right frame shows the spectrum of the resulting 512×512 matrix. The large dotted circle denotes $\{|z| = 1\}$, and the dash-dot inner circle shows the upper bound for the essential spectral radius for this map $\{|z| = 0.6325\}$, $0.6325 = 1/\inf_{x \in [0,1]} |T'(x)|$. The cross shows an eigenvalue of P_{512} that is clearly outside this inner region and therefore corresponds to a true isolated eigenvalue of \mathcal{P} . The eigenfunction for this isolated eigenvalue is plotted in the lower-right frame.

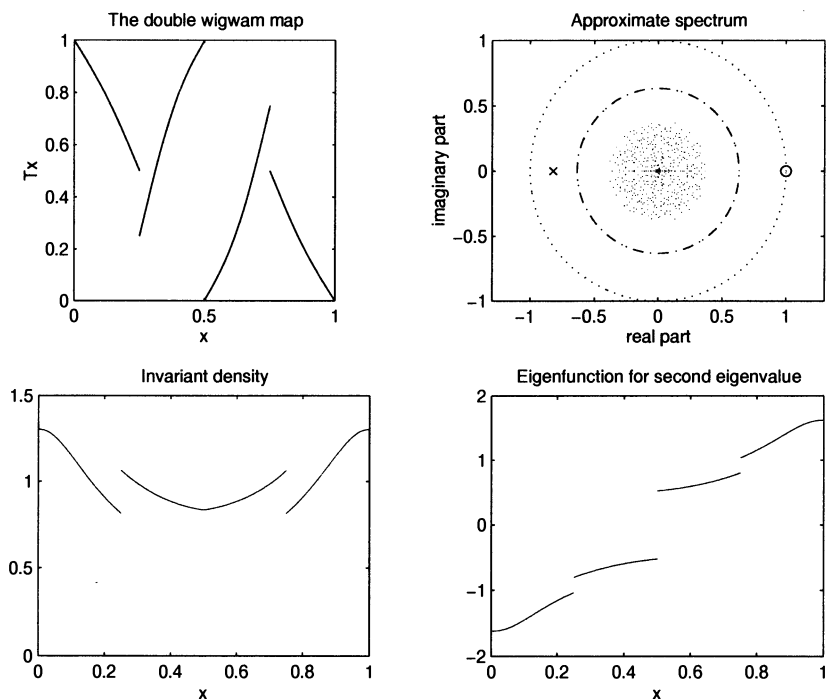


FIGURE 12.4. (upper left): Graph of T ; (upper right): Spectrum of 512×512 transition matrix P_{512} , the small circle represents the eigenvalue 1, the small cross represents another isolated eigenvalue; (lower left): Plot of the invariant density of T (the eigenfunction for the eigenvalue 1); (lower right): Plot of the eigenfunction for the second isolated eigenvalue.

What about higher-dimensional systems? For uniformly hyperbolic systems, a standard technique is to factor out the stable directions and con-

sider only the action of T along unstable manifolds W^u . This induces an expanding map $T_E : W^u \curvearrowright$ with corresponding Perron-Frobenius operator \mathcal{P}_E . We have the following result.

Theorem 12.3. ([15]). *Let $T : M \curvearrowright$ be $C^{\gamma+1}$, $0 < \gamma \leq 1$, uniformly hyperbolic, and possess a nice⁹ Markov partition. Construct P_n by setting \mathfrak{P}_n to be a refinement of this Markov partition, and consider \mathcal{P}_E to act on the function space $C^\gamma(W^u, \mathbb{R})$. Isolated eigenvalues of $\mathcal{P}_E : C^\gamma(W^u, \mathbb{R}) \curvearrowright$ and the corresponding eigenfunctions are approximated by eigenvalues and the corresponding eigenvectors of P_n (eigenfunction convergence in the smooth C^γ sense). The rate of convergence of both the eigenvalues and the eigenvectors to the isolated eigenvalues and corresponding eigenfunctions of \mathcal{P} is $O(1/n^r)$, where $0 < r < 1$ depends¹⁰ only on maximal and minimal values of the derivative of T in unstable directions, and convergence of the eigenfunctions (including the invariant density) is with respect to the stronger smooth norm.*

In the hyperbolic case, a bound for the rate of decay for the full map T may be extracted from the rate of decay for the induced expanding map T_E (see [15]). The case of uniformly expanding T is a simple special case of Theorem 12.3.

12.3.4 Lyapunov Exponents and Entropy

ROUGH IDEA #1: *Lyapunov exponents may be calculated by averaging local rates of expansion according to the physical measure (a spatial average).*

ROUGH IDEA #2: *The Lyapunov exponents of the Markov model approximate the Lyapunov exponents of T .*

ROUGH IDEA #3: *The local stretching rates and dynamics may be encoded in the matrix P to provide Lyapunov exponent and entropy estimates.*

We begin by recalling that for one-dimensional systems, the expression (12.8) may be rewritten as

$$\lambda = \int_M \log |T'(x)| \, d\mu(x) \tag{12.14}$$

by a straightforward application of (12.5) with $f(x) = \log |T'(x)|$. Thus, once we have an estimate of the physical measure μ , it is easy to compute

⁹See [18] for a definition of nice.

¹⁰Let $1/\ell$ (resp. $1/L$) denote the minimal (resp. maximal) stretching rate of T in unstable directions. Then $r = \log(\ell)/2 \log(L)$ and convergence is in the $\|\cdot\|_{\gamma/2}$ norm.

an approximation of λ via

$$\lambda_n := \int_M \log |T'(x)| d\mu_n(x) \approx \sum_{i=1}^n \log |T'(x_{n,i})| \cdot p_{n,i}, \quad (12.15)$$

where $x_{n,i} \in A_{n,i}$ (for example, $x_{n,i}$ could be the midpoint of the interval $A_{n,i}$). The error bounds for $\|\mu - \mu_n\|_1$ (when available) immediately translate into rigorous bounds for the error $|\lambda - \lambda_n|$. We now turn to multidimensional systems.

Approach #1

The direct use of the physical measure for Lyapunov exponent computation may be extended to higher-dimensional systems by rewriting (12.8) as

$$\lambda = \int_M \log \|D_x T(w_x)\| d\mu(x) \quad (12.16)$$

where $\{w_x\}_{x \in M}$ is a family of unit vectors in \mathbb{R}^d satisfying the identity $D_x T(w_x) = w_{Tx}$; different families yield the different Lyapunov exponents (see [21, 22, 53] for details). For the remainder of this section, we consider the problem of finding the largest Lyapunov exponent λ^1 ; the remaining exponents may then be found via standard methods involving exterior products. We denote the vector field corresponding to λ^1 by $\{w_x^1\}$. The vector w_x^1 is the eigenvector of the limiting matrix $\Lambda_x := \lim_{N \rightarrow \infty} ((D_x T^{-N})^\top (D_x T^{-N}))^{1/2N}$ corresponding to the smallest eigenvalue (in magnitude) [53]. One may approximate the vector w_x^1 by computing the smallest eigenvector of the matrix $\Lambda_{N,x} := (D_x T^{-N})^\top (D_x T^{-N})$ for some small finite N . For $N = 7$, the approximate vector field $\{w_{N,x}^1\}$ for the Stiletto map is shown in Figure 12.5. Thus we may compute an approximation to (12.16) by

$$\lambda_{n,N}^1 = \sum_{i=1}^n \log \|D_{x_{n,i}} T(w_{N,x_{n,i}}^1)\| \cdot p_{n,i} \quad (12.17)$$

where $w_{N,x}^1$ denotes the eigenvector obtained from $\Lambda_{N,x}$; see Table 12.1.

TABLE 12.1. Lyapunov exponent estimates for the Stiletto map using the method of [22] and the approximate invariant measure of Figure 12.3.

N	1	2	3	4	5	6	7	8
$\lambda_{1215,N}^1$	0.6826	0.4107	0.4458	0.4536	0.4530	0.4528	0.4529	0.4528




FIGURE 12.5. Approximate vector field of the “Lyapunov directions” $\{w_x^1\}$ corresponding to the largest Lyapunov exponent. In this two-dimensional example, w_x^1 should be tangent to the unstable manifold passing through x .

A recent method put forward by Aston and Dellnitz [2] notes that

$$\lambda^1 = \inf_{N \geq 1} \frac{1}{N} \int_M \log \|D_x T^N\| \, d\mu(x). \tag{12.18}$$

They therefore propose the approximation

$$\lambda_{n,N}^1 := \frac{1}{N} \sum_{i=1}^n \log \|D_{x_{n,i}} T^N\| \cdot p_{n,i}. \tag{12.19}$$

TABLE 12.2. Lyapunov exponent estimates for the Stiletto map using the method of [2].

N	2	4	8	16	32	64	128	256
$\lambda_{1215,N}^1$	0.6303	0.4530	0.4476	0.4322	0.4435	0.4391	0.4436	0.4441

In practice, $\lambda_{n,N}^1$ greatly overestimates λ^1 , so to speed up convergence, one defines $\lambda_{n,M}^{1'} := 2\lambda_{n,2^M} - \lambda_{n,2^{M-1}}$, $M \geq 1$; the values $\lambda_{n,M}^{1'}$ are observed to have better convergence properties; see Table 12.2.

Approach #2

Through our Markov modeling process, we have approximated the dynamics of T as a large Markov chain governed by P_n . To each state i in the Markov chain, we may associate the Jacobian matrix $D_{x_{n,i}}T$, where $x_{n,i}$ denotes the center point (for example) of the partition set $A_{n,i}$. We now consider the Lyapunov exponents of this Markov chain; as we move from state to state along a random orbit of the chain, we multiply together the matrices we have assigned to these states. This produces a random composition of matrices, and the theory of Lyapunov exponents for random matrix products is well developed (see Section 12.3.4 [1], for example). Continuing our theme of “the deterministic dynamics is well approximated by the Markov model”, we compute the top Lyapunov exponent of the Markov model and use this as an approximation of λ^1 . The top Lyapunov exponent for this Markov chain is given by an equation of the form:

$$\lambda_n^1 := \sum_{i=1}^n \left(\int_{\mathbb{R}P^{d-1}} \log \|D_{x_{n,i}}T(v)\| d\xi_{n,i}^1(v) \right) \cdot p_{n,i} \quad (12.20)$$

where $\xi_{n,i}^1$ is a probability measure on $\mathbb{R}P^{d-1}$ ($d - 1$ -dimensional real projective space, or “the space of directions in \mathbb{R}^d ”); see [20, 21] for details.

TABLE 12.3. Lyapunov exponent estimates for the Stiletto map using the method of [21].

“Resolution of ξ^1 ”	10	20	40	60	80	100	200
$\lambda_{1215,N}^1$	0.4322	0.4001	0.4419	0.4637	0.4408	0.4458	0.4456

Whereas the vector $w_{x_{n,i}}^1$ indicates a single direction in which to measure the stretching caused by $D_{x_{n,i}}T$, the measure $\xi_{n,i}^1$ indicates a distribution of directions in which to measure the stretching. This distribution is essential, as the vectors w_x^1 often vary within partition sets (for example, near the “toe” area of the Stiletto attractor where the unstable manifold bends sharply), so it is necessary to “average” these directions *within* partition sets, rather than take a single direction as in Approach #1. Roughly speaking, the distribution $\xi_{n,i}^1$ can be thought of as a histogram of the vectors w_x^1 , $x \in A_{n,i}$. For reasons of space, we refer the reader to [20, 21], in which the details of the calculation of λ_n^1 are spelled out.

Rigorous Results (Approach #3)

The two preceding approaches are not rigorous. Approach #1 is not rigorous because we do not know that $\mu_n \rightarrow \mu$ (although numerically this appears to happen). Approach #2 additionally suffers from the possible sensitivity of the Lyapunov exponents to perturbations of the system – for

example, the perturbation we used to create the Markov model – however, this sensitivity is also rarely observed numerically. In the uniformly expanding or uniformly hyperbolic case, if we use a Markov partition to construct our transition matrix, one can prove convergence of Lyapunov exponent estimates to the true value, and additionally, obtain rigorous estimates of the metric entropy and escape rate / pressure of the system.

Theorem 12.4. ([18]). *Construct $Q_n = m(A_{n,i} \cap T^{-1}A_{n,j})/m(A_{n,j})$ using a nice Markov partition. Let ϱ_n denote the largest eigenvalue of Q_n , and v_n the corresponding right eigenvector. Construct the stochastic matrix*

$$P_{n,ij} = Q_{n,ij}v_{n,j}/\varrho_n v_{n,i},$$

and compute the fixed left eigenvector p_n of P_n . Define μ_n as in (12.11) with $\mu_n(A_{n,i}) = p_{n,i}$. Define

$$\lambda_n := - \sum_{i,j=1}^n p_{n,i} P_{n,ij} \log Q_{n,ij} \tag{12.21}$$

$$h_n := \log \varrho_n + \lambda_n \tag{12.22}$$

Then as $n \rightarrow \infty$, $\mu_n \rightarrow \mu$, $\lambda_n \rightarrow \sum_{\lambda^{(i)} > 0} \lambda^{(i)}$ (the sum of the positive Lyapunov exponents), $h_n \rightarrow h_\mu(T)$ (the measure-theoretic entropy of T with respect to μ), and $\varrho_n \rightarrow P(T)$ (the topological pressure of T). Convergence rates are also available.

12.3.5 Mean and Variance of Return Times

ROUGH IDEA: The mean and variance of return times calculated for the Markov model approximate those of T .

REQUIRED COMPUTATIONS: Calculate the fixed left eigenvector of P , and solve a linear equation of the form $Ax = b$.

We have the following abstract result:

Theorem 12.5. *Let $T : X \rightarrow X$ preserve an ergodic invariant measure μ . Let $B \subset X$, with $0 < \mu(B) < 1$, and set $B^c = X \setminus B$.*

(i)

$$\mathbb{E}_{\mu|_B}(R) = 1/\mu(B), \tag{12.23}$$

(ii)

$$\text{var}_{\mu|_B}(R) = \frac{1 - \mu(B)}{\mu(B)} (2\mathbb{E}_{\mu|_{B^c}}(R) - 1/\mu(B)). \tag{12.24}$$

Part (i) is due to Kac [34] and (ii) to Blum and Rosenblatt [5] (though we have used the version given in [13]; see also [58]). Theorem 12.5 reduces the problem of calculating $\text{var}_{\mu|_B}(R)$ to a calculation of $\mathbb{E}_{\mu|_{B^c}}(R)$, the expected first absorption time into B for points in B^c . The calculation of first absorption times is simple to do when the dynamical system is a finite state Markov chain.

Let B be a subset of the phase space M . This set will be covered by a collection of partition sets $A_{n,i_1}, \dots, A_{n,i_q}$; in practice, one obtains more accurate results if it can be arranged so that B is exactly a union of some collection of partition sets.

We now apply Theorem 12.5 to our Markov model¹¹, setting B to be the collection of states $\{i_1, \dots, i_q\}$. Using (12.23) and (12.24) it is possible to calculate *exact* values for the mean and variance of the return times to the collection of states $\{i_1, \dots, i_q\}$. Appealing then to our guiding principle that our Markov model approximates our original map T , we take the mean and variance of the return times to the states $\{i_1, \dots, i_q\}$ as the approximate mean and variance for the return times to the set $B \subset M$.

We now outline the necessary computations for our Markov model. For simplicity, we assume that the states of our Markov chain have been reordered, so that the states $\{i_1, \dots, i_q\}$ now have labels $\{n - q + 1, \dots, n\}$; for the remainder of this section, the matrix P_n will denote this reordered matrix. To calculate the mean of the recurrence time (denoted \mathcal{M}_n):

1. Calculate the invariant density p_n for the Markov chain governed by P_n .
2. Set $\mathcal{M}_n := 1 / \sum_{i=n-q+1}^n p_{n,i}$.

To calculate the expected absorption time and the variance of the recurrence time (denoted \mathcal{A}_n and \mathcal{V}_n respectively):

1. Write P_n in the block form

$$P_n = \begin{pmatrix} Q_n & U_n \\ V_n & Y_n \end{pmatrix} \tag{12.25}$$

where the matrix Q_n is an $(n - q) \times (n - q)$ matrix giving the transition probabilities between states in our Markov model *not* corresponding to the set $B \subset M$.

2. Calculate the solution τ_n to the linear equation (I_{n-q} is the $(n - q) \times (n - q)$ identity matrix)

$$(I_{n-q} - Q_n)\tau_n = (1 \quad 1 \quad \dots \quad 1)^\top, \tag{12.26}$$

¹¹Formally, we set $X = \Omega = \{(\omega_0, \omega_1, \dots) : P_{\omega_i, \omega_{i+1}} > 0, \omega_i \in \{1, \dots, n\}, i \geq 0\}$, $T = \sigma : \Omega \rightarrow \Omega$, the left shift on Ω , and $\mu = \mathbb{M}$, the Markov measure defined on cylinder sets by $\mathbb{M}([\omega_i, \dots, \omega_{i+t}]) = p_{\omega_i} P_{\omega_i, \omega_{i+1}} \cdots P_{\omega_{i+t-1}, \omega_{i+t}}$. Then $B = [\omega_{i_1}] \cup \dots \cup [\omega_{i_q}]$.

3. Set $\mathcal{A}_n := (\sum_{i=1}^{n-q} p_{n,i} \tau_{n,i}) / (\sum_{i=1}^{n-q} p_{n,i})$.

4. Set $\mathcal{V}_n := (\mathcal{M}_n - 1)(2\mathcal{A}_n - \mathcal{M}_n)$.

See [19] for further details.

The number $\tau_{n,i}$ is an approximation of the average time required for a point in $A_{n,i}$ to move into B and is often of interest in itself.

Example 12.4. (The bouncing ball). We study a two-dimensional map of a cylinder $T : S^1 \times \mathbb{R} \rightarrow S^1 \times \mathbb{R}$ that describes the evolution of a ball bouncing on a sinusoidally forced table. We set

$$T(\phi, v) = (\phi + v, \alpha v - \gamma \cos(\phi + v)) \tag{12.27}$$

where $\phi \in [0, 2\pi)$ represents the phase of the table at impact, $v \in \mathbb{R}$ the velocity of the ball just after impact with the table, and T represents the evolution from one impact to the next; see [28] for details. We set $\alpha = 0.5$ and $\gamma = 10$ for the remainder of this example. Figure 12.6 shows a typical orbit of the system, and Figure 12.7 shows an approximation of the “physical” invariant measure μ ; again, there is good agreement between the two distributions. We suppose that we are




FIGURE 12.6. Plot of orbit of length 50,000 for the bouncing ball map.

interested in the time between successive impacts where the velocity of the ball is very low; that is, we create a time series by counting the time intervals between instances when the ball leaves the table with a velocity of magnitude less than




FIGURE 12.7. Approximate invariant measure for the bouncing ball map using 16,566 partition sets; darker regions indicate higher density.

TABLE 12.4. Estimates of the mean and variance of return times to the set $B = S^1 \times [-1.25, 1.25]$.

Number of partition sets n	1654	5086	16566
Mean \mathcal{M}_n	10.33	9.97	9.85
Root Variance $\sqrt{\mathcal{V}_n}$	9.82	9.45	9.33

1.25. Thus $B = S^1 \times [-1.25, 1.25]$ in the earlier notation. Performing the analysis described earlier, Table 12.4 shows the results for various partition refinements.

Compare these values with 9.830 ± 0.021 and 9.373 ± 0.032 , the mean and variance, respectively, obtained directly from ten orbits of length 10^6 (plus/minus one standard deviation of the ten values obtained).

The result of the calculation of τ_n is shown in Figure 12.8. It is clear that there is a sharp divide between areas which return to low velocities relatively quickly (the very dark strips in Figure 12.8) and those areas that take longer to return. A histogram plot of τ reveals that if a point does not return very quickly to B , then it takes a much longer time.




FIGURE 12.8. Approximate absorption times into the set $S^1 \times [-1.25, 1.25]$; faster absorption is indicated by darker shading.

12.4 Random Systems

We now discuss extensions of our Markov modeling process to random dynamical systems. Instead of having a single deterministic mapping $T : M \rightarrow M$, we now have a collection of mappings $\{T_1, \dots, T_r\}$, where $T_k : M \rightarrow M$ for $k = 1, \dots, r$. A random orbit $\{x_N\}_{N=0}^\infty$ is defined by setting

$$x_N = x_N(k_{N-1}, \dots, k_0, x_0) := T_{k_{N-1}} \circ \dots \circ T_{k_1} \circ T_{k_0} x_0, \quad N \geq 1, \quad (12.28)$$

where the indices $k_0, k_1, \dots \in \{1, \dots, r\}$ are generated by a stationary stochastic process. We will be considering two situations; namely where the indices are generated by I.I.D. processes and by Markov processes.

The former case corresponds to the situation where at each time step, a map is selected at random (according to some fixed probability distribution on the numbers $\{1, \dots, r\}$), independently of any previously applied maps. We shall say that the probability of selecting the map T_k at any given time is $w_k \geq 0$; naturally $\sum_{k=1}^r w_k = 1$. The probability of the sequence $T_{k_{N-1}} \circ \dots \circ T_{k_0}$ occurring is simply $w_{k_{N-1}} \cdots w_{k_0}$. By extending this product to infinite sequences of indices, we obtain a probability measure \mathbb{P} on $\{1, \dots, r\}^{\mathbb{Z}^+}$. The monograph [40] lays down much of the basic theory concerning such dynamical systems, including generalizations of standard dynamical systems quantities to the random situation.

In the Markov case, the probability with which a map is selected at each

time step depends only on the previously applied map¹². We shall say that the probability of selecting a map T_l given that the map T_k was selected at the last time step, is $W_{kl} \geq 0$; we require that $\sum_{l=1}^r W_{kl} = 1$, and so W is a stochastic matrix governing the Markov chain that produces our random sequence of indices. The probability of the sequence $T_{k_{N-1}} \circ \dots \circ T_{k_0}$ occurring is $w_{k_0} W_{k_0 k_1} \dots W_{k_{N-2} k_{N-1}}$; by extending this to infinite sequences, we obtain a probability measure (a Markov measure) on $\{1, \dots, r\}^{\mathbb{Z}^+}$, which we also denote by \mathbb{P} .

Examples 12.5

1. Consider the bouncing ball map of the last section. Suppose that our ball is nonuniform, and that one side is more “springy” than the other. Sometimes, the ball will land on the springy side, and sometimes it will land on the not-so-springy side. Which side the ball lands on determines the value of α , and so at each time step there is a random choice of α , and therefore an application of either $T_{\alpha_{\text{springy}}}$ or $T_{\alpha_{\text{not-so-springy}}}$. We’ll return to this example later.
2. A set of maps $\{T_1, \dots, T_r\}$ could arise as perturbations of a single map T via $T_k x := Tx + \epsilon_k$, where $\epsilon_k \in \mathbb{R}^d$ is a perturbation. We choose a probability vector (w_1, \dots, w_r) where the value w_k represents the probability of our map T encountering the perturbation ϵ_k . A random I.I.D. composition of the $\{T_k\}$ models a deterministic system subjected to small I.I.D. perturbations.
3. Random dynamical systems can also arise in the context of dynamical systems with inputs. The effect of an input is essentially to produce different dynamics (in other words, a different map T_k) at the time step in which it occurs. If the model is truly random, these inputs could occur according to an I.I.D. process or Markov process. However, more structured sets of inputs can also be modeled by Markov processes, for example, where a randomly selected input triggers a fixed sequence of inputs before another random input is selected.

We now define what is meant by an invariant measure for our random system.

Definition 12.2. Let $\Omega = \{1, \dots, r\}^{\mathbb{Z}^+}$, and for $\omega = (\omega_0, \omega_1, \omega_2, \dots) \in \Omega$, define the left shift $\sigma : \Omega \rightarrow \Omega$ by $(\sigma\omega)_i = \omega_{i+1}$. The probability measure \mathbb{P} on Ω introduced earlier is σ -invariant. Define the skew product $\tau : \Omega \times M \rightarrow \Omega \times M$ by $\tau(\omega, x) = (\sigma\omega, T_{\omega_0}x)$; our random orbits $\{x_N\}_{N \geq 0}$ may be written as $x_N = \text{Proj}_M(\tau^N(\omega, x_0))$, where Proj_M denotes the canonical projection onto M .

¹²If one desires to treat Markov processes with longer memory, they may be written in terms of a first-order Markov chain in the standard way.

We will say that a probability measure μ on M is an invariant measure for our random system, if there exists a τ -invariant probability measure $\tilde{\mu}$ on $\Omega \times M$ such that

1. $\tilde{\mu}(E \times M) = \mathbb{P}(E)$ for all measurable $E \subset \Omega$, and
2. $\tilde{\mu}(\Omega \times B) = \mu(B)$ for all measurable $B \subset M$.

Definition 12.3. A probability measure μ is called a natural or physical measure for a random dynamical system if μ is defined as $\mu(B) = \tilde{\mu}(\Omega \times B)$ where $\tilde{\mu}$ is a τ -invariant probability measure satisfying

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} f(\tau^k(\omega, x)) \rightarrow \int_{\Omega \times M} f(\omega, x) d\tilde{\mu}(\omega, x) \tag{12.29}$$

for all continuous $f : \Omega \times M \rightarrow \mathbb{R}$, and $\mathbb{P} \times m$ almost all $(\omega, x) \in \Omega \times M$.

Remark 12.1. If we choose the continuous test function f in (12.29) to be independent of ω , then we have the simple consequence that:

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{j=0}^{N-1} f(T_{k_j} \circ \dots \circ T_{k_0} x) \rightarrow \int_M f d\mu \tag{12.30}$$

for Lebesgue-almost-all $x \in M$ and \mathbb{P} -almost-all random sequences of maps.

By setting $f(x) = \chi_A(x)$, where $A \subset M$ is such that $\mu(\partial A) = 0$, then

$$\frac{1}{N} \text{card}\{0 \leq j \leq N - 1 : T_{k_j} \circ \dots \circ T_{k_1} \circ T_{k_0} x \in A\} \rightarrow \mu(A),$$

again for Lebesgue almost all $x \in M$ and \mathbb{P} -almost-all random sequences.

In rough terms, this says that if you plot the points in a random orbit defined by (12.28), then for Lebesgue-almost-all starting points x_0 and \mathbb{P} -almost-all random sequences of maps, one obtains the same distribution of points. From the point of view of analysing the average behavior of the random system, this is the correct distribution to approximate. The physical measure μ is usually *not* invariant under *any* of the individual transformations T_k , in the sense that $\mu \circ T_k^{-1} \neq \mu$. However, μ is “invariant on average”, which by heavily abusing notation may be written as $\mathbb{E}(\mu \circ T_k^{-1}) = \mu$. In the I.I.D. case, this formula is entirely accurate, as there the invariance condition is simply $\sum_{k=1}^r w_k \mu \circ T_k^{-1} = \mu$.

12.4.1 Basic Constructions

We must construct a transition matrix

$$P_n(k) = \frac{m(A_{n,i} \cap T_k^{-1} A_{n,j})}{m(A_{n,i})} \tag{12.31}$$

for each of the maps T_k .

Remark 12.2. *An alternative definition of the matrix $P_n(k)$ is as follows. Within each set $A_{n,i}$ select a single point $a_{n,i}$. Then set*

$$P'_n(k) = \begin{cases} 1, & \text{if } T_k a_{n,i} \in A_{n,j}, \\ 0, & \text{otherwise.} \end{cases} \quad (12.32)$$

Clearly, the computational effort involved in the numerical construction of $P'_n(k)$ is less than that of $P_n(k)$ in (12.31), especially in higher dimensions (in the Tips and Tricks section, we discuss other numerical methods of computing $P_n(k)$). We do not recommend using $P'_n(k)$ for deterministic systems, as the results are usually very poor. However, for random systems, one can still obtain quite good results with the cruder approximation of (12.32).

How these matrices are combined depends on whether the stochastic process is I.I.D. or Markov.

I.I.D. Case

In the I.I.D. case, we set

$$P_n = \sum_{k=1}^r w_k P_n(k). \quad (12.33)$$

Markov Case

In the Markov case, let W be the transition matrix for the Markov process that generates the random sequence of indices for the maps $\{T_k\}$.

Now set

$$S_n = \begin{pmatrix} W_{11} P_n(1) & W_{12} P_n(1) & \cdots & W_{1r} P_n(1) \\ W_{21} P_n(2) & W_{22} P_n(2) & \cdots & W_{2r} P_n(2) \\ \vdots & \vdots & \ddots & \vdots \\ W_{r1} P_n(r) & W_{r2} P_n(r) & \cdots & W_{rr} P_n(r) \end{pmatrix}. \quad (12.34)$$

In both the I.I.D. and Markov cases, the matrices P_n and S_n may be thought of as finite-dimensional projections of an “averaged” Perron-Frobenius operator; see [17] for details. With either of these two matrices, one may apply the methods described in Section 12.3 to estimate invariant objects such as invariant measures, invariant sets, Lyapunov exponents, and recurrence times.

12.4.2 Invariant Measures

I.I.D. Case

We calculate the fixed left eigenvector p_n of P_n as constructed in (12.33), and normalize so that

$$\sum_{i=1}^n p_{n,i} = 1. \quad (12.35)$$

Set $\mu_n(A_{n,i}) = p_{n,i}$ and define the approximate invariant measure as in (12.11).

Markov Case

We calculate the fixed left eigenvector of S_n and denote this as $s_n = [s_n^{(1)} | s_n^{(2)} | \cdots | s_n^{(r)}]$ where each $s_n^{(k)}$, $k = 1, \dots, r$ is a vector of length n , and $\sum_{k=1}^r \sum_{i=1}^n s_{n,i}^{(k)} = 1$. Define the approximate invariant measure as

$$\mu_n(A_{n,i}) = \sum_{k=1}^r s_{n,i}^{(k)}. \quad (12.36)$$

We now use (12.11) again to define a measure on all of M .

Results parallel to those of Theorem 12.1 hold for our random systems.

Proposition 12.1. *Suppose that each $T_k : M \curvearrowright$ is continuous and the resulting random dynamical system has a physical measure μ (in the weaker sense where only (12.30) need hold, rather than (12.29)). Let $\{\mu_n\}$ denote a sequence of approximate invariant measures as defined in either (12.35) or (12.36), and let μ^* be a weak limit of this sequence. Denote by S the intersection $\bigcap_{n=n_0}^{\infty} \text{supp } \mu_n$. Then the conclusions of Theorem 12.1 hold.*

Proof: One first requires the facts that the matrices (12.33) and (12.34) represent a finite-dimensional approximation of an appropriately averaged Perron-Frobenius operator; this is detailed in [17]¹³. With this established, the proofs run along the same lines as the deterministic case. \square

Example 12.6. (The (nonuniform) bouncing ball)

We now suppose that our bouncing ball has gone soft on one side, so that sometimes we register a value of $\alpha = 0.1$, rather than the original value of $\alpha = 0.5$. We assume that every time it lands on the soft side, it will surely land on the good side next time, while if it lands on the good side, it has a 50/50 chance of landing on the soft side next time. The situation we describe is a Markov random

¹³The form of (12.34) is slightly different to the matrix given in [17] as we have performed a similarity transformation on the latter to yield a more intuitive representation.

composition of two mappings $T_{\alpha=0.5}$ and $T_{\alpha=0.1}$. The transition matrix for this Markov chain is

$$W = \begin{pmatrix} 1/2 & 1/2 \\ 1 & 0 \end{pmatrix},$$

where $\alpha = 0.5$ is identified with state #1 and $\alpha = 0.1$ with state #2. We construct S_n as in (12.34), and compute the approximate invariant measure as in (12.36); see Figure 12.10. Again, there is good agreement between the two figures.




FIGURE 12.9. Plot of orbit of length 50,000 for the random bouncing ball map

Rigorous Results

In certain situations, one is able to obtain rigorous upper bounds for the difference between μ_n and μ .

The first of these is where the random system *contracts* phase space on average. Typical examples of such systems are the iterated function systems (IFS's) of Barnsley [3] and co-workers, where a finite number of mappings are selected using an I.I.D. law to create fractal images. Suppose that T_k is selected with probability w_k , and define $s_k = \max_{x,y \in M} \|T_k x - T_k y\|/\|x - y\|$ as the Lipschitz constant for T_k ; then set $s = \sum_{k=1}^r w_k s_k$. It is straightforward to show that if $s < 1$, then this random dynamical system has a unique invariant measure, the support of which is a fractal




FIGURE 12.10. Approximate invariant measure for the random bouncing ball map using 18 512 partition sets.

set. Furthermore, one has the bound¹⁴ [20] (see also [55])

$$d_H(\mu, \mu_n) \leq \left(\frac{1+s}{1-s} \right) \max_{1 \leq i \leq n} \text{diam}(A_{n,i}), \tag{12.37}$$

where d_H is the natural metric generating the weak topology on measures, defined by

$$d_H(\nu_1, \nu_2) = \sup \left\{ \left| \int_M h \, d\nu_1 - \int_M h \, d\nu_2 \right| ; h : M \rightarrow \mathbb{R} \text{ has Lipschitz constant } 1 \right\}.$$

If a uniform grid is used, this bound may be improved by a factor of 2. Similar results hold for Markov compositions. Ulam-type methods of approximating invariant measures of IFS's are also discussed in [51].

The second situation is where the dynamical system is *expanding* on average. This setting is more complicated as the random system may have infinitely many invariant measures, and it is important to show that Ulam's method approximates the physical measure (in this expanding case, the physical measure will have a bounded density). In the case of I.I.D. random

¹⁴This rigorous result holds even when the crude approximation of Remark 12.2 is used.

dynamical systems on the interval $[0, 1]$ where each map T_k is a *Lasota-Yorke map*, it is known that a bounded invariant density exists provided that $\sum_{k=1}^r w_k / |T'_k(x)| < 1$; see [50]. Under some additional conditions, it is shown in [17] that (i) the random system (either I.I.D. or Markov) possesses a unique bounded invariant density, and (ii) that the Ulam estimates μ_n converge to the physical measure μ (which has a bounded density). In addition, convergence rates of $O(\log n/n)$ for the difference $\|\mu - \mu_n\|_{L^1}$ are proven, and if each T_k is a C^2 map of the circle S^1 , rather than of the interval $[0, 1]$, explicitly calculable numerical bounds for the error $\|\mu - \mu_n\|_{L^1}$ are given. In the future, we will no doubt see extensions of these results to higher dimensions.

12.4.3 Lyapunov Exponents

The random version of (12.8) is

$$\lambda := \lim_{N \rightarrow \infty} \frac{1}{N} \log \|D_{x_{N-1}} T_{k_{N-1}} \circ \dots \circ D_{x_1} T_{k_1} \circ D_x T_{k_0}(v)\|. \quad (12.38)$$

Often, the same value of λ is obtained for \mathbb{P} -almost-all random sequences, Lebesgue-almost-all $x \in M$, and for every $v \in \mathbb{R}^d$. We denote this value by λ^1 .

Things are very simple in the case of one-dimensional systems driven by an I.I.D. process. In this case, the expression (12.38) may be alternately expressed as

$$\lambda = \sum_{k=1}^r w_k \int_M \log |T'_k(x)| \, d\mu(x) \quad (12.39)$$

by a straightforward application of (12.29) with $f(\omega, x) = \log |T'_{\omega_0}(x)|$. Thus, once we have an estimate of the physical measure μ , it is easy to compute an approximation of λ via

$$\begin{aligned} \lambda_n &:= \sum_{k=1}^r w_k \int_M \log |T'_k(x)| \, d\mu_n(x) \\ &\approx \sum_{k=1}^r w_k \sum_{i=1}^n \log |T'_k(x_{n,i})| \cdot p_{n,i}, \end{aligned} \quad (12.40)$$

where $x_{n,i} \in A_{n,i}$ (for example, the midpoint of the interval $A_{n,i}$).

Rigorous Results

For random systems, we adopt Approach #2 of Section 12.2.3, as the other two approaches are not so helpful in the random situation. Here we only briefly describe the calculation of Lyapunov exponents of I.I.D. random

dynamical systems where the Jacobian matrices are constant. This situation arises when each of the mappings T_k are affine (as with most IFS's). Equation (12.38) now becomes independent of the x_0, \dots, x_{N-1} , and is a function only of the sequence k_0, \dots, k_{N-1} ; we are essentially dealing with an I.I.D. random matrix product.

Suppose that M is two-dimensional so that our Jacobian matrices are 2×2 matrices. We need to define a probability measure ξ on the angle space $\mathbb{RP}^1 \cong [0, \pi)$ (ξ is a relative of the probability measure alluded to in Section 12.3.4). Each matrix DT_k (note independence of x) defines a map from $[0, \pi)$ to itself via $\angle(v) = \angle(DT_k(v))$, where $\angle(v)$ is the angle between v and some fixed reference vector in \mathbb{R}^2 . To simplify notation, we will identify a vector $v \in \mathbb{R}^2$ and its angle with respect to some fixed reference vector. The Jacobian matrix DT_k will then be thought of as an action on the space of angles $[0, \pi)$.

Example 12.7. Suppose $DT_k = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}$, and $v = (1, 1)$. We identify v with the angle $\pi/2$ (this is the angle v makes with the reference vector $(1, 0)$). Then $DT_k(v) = (1, 3)$, and we identify this vector with the angle $\tan^{-1} 3$. By a slight abuse of notation, we may write $DT_k(\pi/2) = \tan^{-1}(3)$, and in this way we consider DT_k to be an action on the set of angles $[0, \pi)$.

The probability measure ξ on \mathbb{RP}^1 that we seek should satisfy:

$$\xi(E) = \sum_{k=1}^r w_k DT_k^{-1}(E) \tag{12.41}$$

for every measurable subset $E \subset [0, \pi)$. This is because of the following result.

Theorem 12.6. ([25]). *Suppose that each $d \times d$ matrix DT_k , $k = 1, \dots, r$ is nonsingular, and that the only subspace of \mathbb{R}^d that is invariant under all of the DT_k is the trivial subspace. Then with probability one,*

$$\lambda^1 = \lim_{N \rightarrow \infty} \frac{1}{N} \log \|DT_{k_{N-1}} \circ \dots \circ DT_{k_0}(v)\|$$

for every $v \in \mathbb{R}^d$. Furthermore,

$$\lambda^1 = \lambda(\xi) := \sum_{k=1}^r w_k \int_{\mathbb{RP}^{d-1}} \log \|DT_k(v)\| d\xi(v),$$

where ξ is any probability measure satisfying (12.41).

We approximate a measure ξ satisfying (12.41) in essentially the same way that we have already used for measures on M . Partition $[0, \pi)$ into

a finite collection of intervals $E_1, \dots, E_m \subset [0, \pi)$, and define an $m \times m$ stochastic matrix by:

$$D_{m,gh}(k) = \frac{m(E_g \cap DT_k^{-1}(E_h))}{m(E_g)}. \quad (12.42)$$

Alternatively (*ala* Remark 12.2), one could choose a collection of points e_1, \dots, e_m such that $e_g \in E_g$ and set

$$D_{m,gh}(k) = \begin{cases} 1, & \text{if } DT_k(e_g) \in E_h \\ 0, & \text{otherwise,} \end{cases} \quad (12.43)$$

or use the other suggestions in Section 12.6.1. One now computes the fixed left eigenvector of the matrix $D_m = \sum_{k=1}^r w_k D_m(k)$; we denote this eigenvector by ξ_m . Selecting points e_1, \dots, e_m as before, we define an approximation of λ^1 as

$$\lambda_m^1 := \sum_{k=1}^r w_k \sum_{g=1}^m \log \|D_m(k)(e_g)\|, \quad (12.44)$$

where, in the expression $\|D_m(k)(e_g)\|$, e_g is a unit vector in the direction represented by e_g , and we measure the length of the vector $D_m(k)(e_g)$ (this is the factor by which $D_m(k)$ stretches vectors in the direction of e_g). The preceding constructions may be generalized to higher dimensions; see [20] for details. We have summarized the simplest situation here; the treatment of Markov random matrix products, and I.I.D. random nonlinear dynamical systems may be found in [20].

12.4.4 Mean and Variance of Return Times

To estimate the mean and variance of return times, we again construct a finite Markov model, and calculate the mean and variance of return times to a suitable set of states. In the I.I.D. case we can define a Markov model using (12.33), and proceed as for deterministic systems.

In the Markov case, we use (12.34), and produce a left eigenvector s_n of S_n such that $\sum_{k=1}^r \sum_{i=1}^n s_{n,i}^{(k)} = 1$. When writing S_n in the block form of (12.25), recall that each partition set $A_{n,i}$ corresponds to r states of the Markov chain governed by S_n . With this in mind, one may substitute S_n and s_n into the algorithm described in Section 12.3.5. It is also possible to consider situations where the set B depends on the map T_k which is currently applied [19].

Example 12.8. (The nonuniform bouncing ball (cont...))

We return to the random dynamical system of Example 12.6 and compute the return times to low impact velocity configurations described by the set $S^1 \times [-1.25, 1.25]$.

TABLE 12.5. Estimates of the mean and variance of return times to the set $B = S^1 \times [-1.25, 1.25]$.

Number of partition sets n	1738	5488	18512
Mean \mathcal{M}_n	11.18	10.93	10.98
Root Variance $\sqrt{\mathcal{V}_n}$	10.73	10.45	10.49

Compare these values with 11.05 ± 0.03 and 10.57 ± 0.03 , the mean and variance respectively, obtained directly from 10 simulated random orbits of length 10^6 (plus/minus one standard deviation of the 10 values obtained). This example is one situation where we would benefit by using one of the alternate partition selection techniques described in Section 12.6.2.

12.4.5 Advantages for Markov Modeling of Random Dynamical Systems

To close this section, we discuss two further advantages of Markov modeling over trajectory simulation that are not present in the deterministic case.

The first of these concerns the accuracy of the approximations. For deterministic systems, the two competing approaches (simulating long orbits and coarse graining) both have their inaccuracies. The iterative approach of following long orbits (let's assume that we can do perfect computations) has the problem of temporal deviation from equilibrium behavior. That is, we should really have orbits of *infinite* length, but instead we have orbits of *finite* length whose statistical behavior is not the same. In contrast, with the Markov modeling approach, we can *exactly* compute the long term behavior of our model, but we compute the long term behavior of an *approximation* of the system, rather than of the true system. Turning now to random systems, the iterative approach is fighting errors on two fronts: namely, the deviation from equilibrium in phase space mentioned earlier, and additionally, the deviation of the distribution of the finite length random sequence of indices k_0, \dots, k_{N-1} from its equilibrium distribution \mathbb{P} . On the other hand, the Markov modeling approach completely eliminates errors from the random fluctuations by averaging them out through the expectations performed in its construction. Thus our Markov models for random dynamical systems do not suffer from the inaccuracies due to random fluctuations, and are therefore (heuristically at least) more accurate; this is borne out numerically in Lyapunov exponent computations [20].

The second advantage lies in the flexibility of the Markov modeling approach regarding the underlying stochastic process. Suppose that we wish to study a *family of systems* which use the same maps T_1, \dots, T_r , but a different distribution \mathbb{P} (in the bouncing ball example, this would amount to varying the probabilities with which impacts occur on the soft and hard

sides). Most of the computational effort goes into constructing the (fixed) transition matrices $P_n(k)$, $k = 1, \dots, r$, while the ancillary calculations involving eigenvectors and so on, are relatively cheap. Thus, we may perform analyses on a *whole family of systems* very quickly, by reusing most of the original constructions. In contrast, if we were to use the direct method of simulating long orbits, then entirely new orbit calculations would be required for each new set of probabilities.

12.5 Miscellany

We briefly outline some other applications and techniques related to our approach of Markov modeling. Unless otherwise stated, we refer only to deterministic systems.

Global Attractors

Related partition-based methods may be used to approximate the global attractor for a given subset of phase space. If $B \subset M$, then the global attractor of B is defined by $G = \bigcap_{k \geq 0} T^k(B)$. Methods of computing an (in principle¹⁵) rigorous box covering of the global attractor are detailed in [8]. Bounds for the Hausdorff distance between the approximate covering and the global attractor are given for uniformly hyperbolic diffeomorphisms.

Using similar techniques, box coverings for global attractors $G(\omega)$ of individual sample paths ω of random dynamical systems have been studied [38].

Work is in progress on covering “averaged” global attractors of random systems; such global attractors contain all possible random orbits so $G_{av} = \bigcup_{\omega} G(\omega)$.

Noisy Systems

A popular idea is to “noise up” a system by considering the transformation $x \mapsto Tx + \epsilon$, where the perturbation $\epsilon \in \mathbb{R}^d$ is chosen in a uniform distribution from some small ball centered around 0. This may be viewed as defining a random dynamical system where the collection of maps $T_\epsilon = Tx + \epsilon$ is applied in an I.I.D. fashion with equal probability. The Perron-Frobenius operator $\mathcal{P}_\epsilon : L^2(M, m) \rightarrow L^2(M, m)$ for this random perturbation has the desirable property that it is a compact operator under very mild conditions on T , and this greatly simplifies the ergodic theoretic analysis. For example, it is relatively easy to show that this noisy system has a unique invariant probability density, in contrast to the purely deterministic case. The finite-state Markov modeling may now be applied to the perturbed system, and various convergence results proven concerning the approximation of invariant

¹⁵When combined with Lipschitz estimates for the map [33].

measures and invariant sets; see [11]. This setting forms the basis of the thesis [32], where the merits of alternative partitioning methods are also considered.

Rotation Numbers

The approximation of rotation numbers of orientation preserving C^2 circle diffeomorphisms using Ulam constructions is described in [57].

Topological Entropy

It is possible to obtain rigorous upper bounds for the topological entropy of T with respect to a fixed (coarse) partition. All orbits of T are possible under the Markov model; however, the converse is not true. In this sense, the Markov model is more “complex” from the orbit-generating viewpoint. However, as the partitions are refined and our Markov model becomes more accurate, these extra orbits are successively eliminated, so that our upper bounds become increasingly sharp [24].

Spectra of “Averaged” Transfer Operators for Random Systems

One may also attempt to garner dynamical information from the spectrum and eigenvectors of the matrices (12.33) and (12.34), in analogy to the deterministic case. This is work in progress.

12.6 Numerical Tips and Tricks

We discuss methods of computing the transition matrix and of partition selection. Most transition matrix computations in this chapter have used the GAIO software package, available on request from

<http://www.upb.de/math/~agdellnitz/gaio/>

Algorithms 2 and 3 of Section 12.6.1 and 1–4 of Section 12.6.2 in this chapter have been coded in this software.

12.6.1 Transition Matrix Construction

Techniques for the computation of the transition matrix may be split into three main classes; namely “exact” methods, Monte-Carlo/Imaging methods, and an exhaustive method of approximation.

1. “Exact” methods: For one-dimensional systems, it is often possible to construct the transition matrix exactly. If the map is locally one-to-one on each partition set, then only the inverse images of the endpoints of each set need be calculated.

If the inverse images are difficult to obtain, an alternative is to compute the matrix

$$P'_{n,ij} = \frac{m(TA_i \cap A_j)}{m(TA_i)}, \quad (12.45)$$

which in one dimension again requires only the computation of forward images of partition endpoints.

The matrix P'_n is not useful theoretically because (i) forward images of sets may not be measurable, while inverse images (T continuous) of Borel measurable sets are always measurable, and (ii) the standard form (12.3) arises as a discretization of the Perron-Frobenius operator for T , while (12.45) does not. If T is linear on each A_i , then $P_{ij} = P'_{ij}$ for $j = 1, \dots, n$; this forward-imaging exact computation was carried out for two-dimensional piecewise linear reconstructions in [23]. Otherwise, the difference between P and P' is governed by the second derivative of T and the diameter of the partition sets. We do not recommend using forward-imaging for maps with very large second derivatives.

2. *Monte-Carlo / Imaging of test points*: The most popular method is the so-called Monte-Carlo approach [31]. To compute P_{ij} , one randomly selects a large number of points $\{a_1, \dots, a_N\} \subset A_i$ and sets $P_{ij} \approx \#\{a \in \{a_1, \dots, a_N\} : T(a) \in A_j\}/N$. A similar approach is to choose a uniform grid of test points within each partition set and perform the same calculation. My personal feeling is that the latter approach is better as the uniform grid more reliably approximates Lebesgue measure. Any number of variations on the selection of test points can be used, though Monte-Carlo and uniform grids are the most common.
3. *Exhaustion*: A recent approach [29] is to rigorously approximate the transition probability by a means of *exhaustion* akin to the exhaustion methods of Eudoxus. To compute the Lebesgue measure of the portion of A_i that is mapped into A_j , one repeatedly refines the set A_i until it is known (via Lipschitz estimates on the map) that a refined subset of A_i is mapped *entirely inside* A_j . In this way, the set A_i is repeated broken down into small pieces which map entirely inside A_j , with this process terminating at the desired level of precision.

12.6.2 Partition Selection

This section is devoted to suggesting methods of producing better Markov models via smarter partition selection. That is, how should one choose partitions to best capture the dynamics of the system.

Of course, if a Markov partition is available, this is clearly the best choice. However, we are assuming that this is not the case, and we are left with

the decision of how to construct a suitable “grid”. For the most part, we consider partition selection where the criteria for a good partition is that it produces a good estimate of the physical measure (at least a better estimate than a uniform grid would produce). Of course, often we don’t know what the physical measure is, so this mostly restricts rigorous numerical testing to one-dimensional systems. Nevertheless, we outline three main approaches, and suggest heuristically when they may be useful. In all cases, one selects an initial coarse partition, computes the invariant measure for the Markov model, and on the basis of information contained in the invariant measure of the current model, a choice is made on which partition sets to refine and which to not refine.

1. *Standard approach*: Refine any partition sets which are currently assigned nonzero measure.
2. *Equal mass approach*: Refine any partition sets which are assigned a measure greater than $1/n$, where n is the current number of partition sets; see [10].

The rationale behind this is that one should focus more finely on regions where there is large mass. In this sense, the method is not only targeted at obtaining more accurate estimates of the invariant measure, but also more accurate modeling of the *dynamics* of the system; this has been demonstrated for the estimation of return times in [19]. This method is particularly suited to systems which possess a singular physical measure, as in dissipative chaotic systems. However, because of the nonuniform refinement (the minimal ratio of cell sizes is almost always at least 2 if refinement is done by “halving” a set), it often performs worse than the standard method in cases where the physical measure is smooth. For this approach to be useful, the ratio $\sup_{x \in \text{supp } \mu_n} \phi_n(x) / \inf_{x \in \text{supp } \mu_n} \phi_n(x)$ should be much larger than 2 for all $n \geq 0$ (ϕ_n is the density of the approximate measure μ_n).

3. *High derivative approach*: Let c_i denote the “center point” of a partition set A_i . Refine partition sets where the value of

$$E_i := m(A_i) \text{diam}(A_i) \times \max \left\{ \frac{|p_i/m(A_i) - p_j/m(A_j)|}{|c_i - c_j|} : A_j \text{ is a neighbor of } A_i \right\}, \tag{12.46}$$

is greater than $(1/n) \sum_{i=1}^n E_i$; [30].

The expression that is maximized is meant to be an approximation of the derivative of the invariant density ϕ on A_i in the “direction of” A_j . In [30], one assumes that the physical measure is smooth; therefore, if the current estimate of the invariant measure has adjacent sets given very different measures, there must be an error in this region,

so one refines these sets to obtain better estimates. The number E_i is intended to approximate the error incurred on the partition set A_i .

An alternative viewpoint is as follows. In [16], it is noted that the matrix $\tilde{P}_{ij} := \mu(A_i \cap T^{-1}A_j)/\mu(A_i)$ is an optimal approximation in the sense that the fixed left eigenvector \tilde{p} of \tilde{P} assigns exactly the correct weights to the partition sets; that is, $\tilde{p}_i = \mu(A_i)$ (this approach is also followed in [35]). The difference between P and \tilde{P} is essentially given by how “non-Lebesgue-like” the measure μ is *within each partition set*; roughly speaking, how “non-constant” the distribution of μ is within partition sets. One may try to reduce¹⁶ the error $\|\mu - \mu_n\|_1$ by creating a partition which produces a transition matrix P similar to that of the special matrix \tilde{P} . Such an analysis also leads to the error minimization criterion (12.46).

The high derivative method is targeted specifically toward more accurate estimates of the physical measure. It often performs better than the equal mass approach for maps with smooth densities.

4. *Large difference approach* [32]: One refines all partition sets and constructs a temporary transition matrix P_{temp} and invariant measure p_{temp} for the refined partition. This refined invariant measure is compared with the invariant measure p_{old} and only sets in the old partition for which the measure according to p_{temp} and p_{old} is very different are split up. The transition matrix P_{temp} is now discarded. This approach is based on a standard method of numerical analysis.

Comparisons of the three alternative approaches are detailed in [32].

Acknowledgments

We thank Oliver Junge for helpful discussions and critical comments during the preparation of this chapter, and Oliver Junge and Stefan Sertl for assistance with the use of the GAIO package. The list of references is not meant to be exhaustive and is far from being so. GF apologizes for any oversights.

This research was supported by the Deutsche Forschungsgemeinschaft under Grant De 448/5-4.

¹⁶Such a “distortion reducing” approach is also discussed in [47] for one-dimensional maps (in particular, the logistic family $T_a x = ax(1-x)$), and a relative of the *Equal mass approach* is advocated as a means to make P better approximate \tilde{P} .

References

- [1] Ludwig Arnold. *Random Dynamical Systems*. Springer monographs in mathematics. Springer, Berlin, 1998.
- [2] Philip J. Aston and Michael Dellnitz. The computation of Lyapunov exponents via spatial integration with application to blowout bifurcations. *Computer methods in applied mechanics and engineering*, 170:223–237, 1999.
- [3] Michael F. Barnsley. *Fractals everywhere*. Academic Press, Boston, 1988.
- [4] George D. Birkhoff. Proof of the ergodic theorem. *Proceedings of the National Academy of Sciences of the USA*, 17:656–60, 1931.
- [5] J. R. Blum and J. I. Rosenblatt. On the moments of recurrence time. *Journal of Mathematical Sciences*, 2:1–6, 1967.
- [6] Abraham Boyarsky and Pawel Góra. *Laws of Chaos – Invariant Measures and Dynamical systems in One Dimension*. Probability and Its Applications. Birkhäuser, Boston, 1997.
- [7] Michael Dellnitz, Gary Froyland, and Stefan Sertl. On the isolated spectrum of the Perron-Frobenius operator. *Nonlinearity*, 13(4):1171–1188, 2000.
- [8] Michael Dellnitz and Andreas Hohmann. A subdivision algorithm for the computation of unstable manifolds and global attractors. *Numerische Mathematik*, 75:293–317, 1997.
- [9] Michael Dellnitz and Oliver Junge. Almost invariant sets in Chua’s circuit. *International Journal of Bifurcation and Chaos*, 7(11):2475–2485, 1997.
- [10] Michael Dellnitz and Oliver Junge. An adaptive subdivision technique for the approximation of attractors and invariant measures. *Computing and Visualization in Science*, 1:63–68, 1998.
- [11] Michael Dellnitz and Oliver Junge. On the approximation of complicated dynamical behavior. *SIAM Journal on Numerical Analysis*, 36(2):491–515, 1999.
- [12] Jiu Ding and Aihui Zhou. Finite approximations of Frobenius-Perron operators. a solution of Ulam’s conjecture to multi-dimensional transformations. *Physica D*, 92:61–68, 1996.
- [13] H. Francke, D. Plachky, and W. Thomsen. A finitely additive version of Poincaré’s recurrence theorem. In N. Christopeit, K. Helmes, and M. Kohlmann, editors, *Stochastic Differential Systems - Proceedings of the 3rd Bad Honnef Conference, June 3–7, 1985*, Lecture Notes in Control and Information Science, Berlin, 1985. Springer-Verlag.
- [14] Gary Froyland. Finite approximation of Sinai-Bowen-Ruelle measures of Anosov systems in two dimensions. *Random & Computational Dynamics*, 3(4):251–264, 1995.
- [15] Gary Froyland. Computer-assisted bounds for the rate of decay of correlations. *Communications in Mathematical Physics*, 189(1):237–257, 1997.
- [16] Gary Froyland. Approximating physical invariant measures of mixing dynamical systems in higher dimensions. *Nonlinear Analysis, Theory, Methods, & Applications*, 32(7):831–860, 1998.
- [17] Gary Froyland. Ulam’s method for random interval maps. *Nonlinearity*, 12(4):1029–1052, 1999.
- [18] Gary Froyland. Using Ulam’s method to calculate entropy and other dynamical invariants. *Nonlinearity*, 12:79–101, 1999.

- [19] Gary Froyland and Kazuyuki Aihara. Estimating statistics of neuronal dynamics via Markov chains. To appear in *Biological Cybernetics*.
- [20] Gary Froyland and Kazuyuki Aihara. Rigorous numerical estimation of Lyapunov exponents and invariant measures of iterated function systems and random matrix products. *Int. J. Bifur. Chaos*, 10(1):103–122, 2000.
- [21] Gary Froyland, Kevin Judd, and Alistair I. Mees. Estimation of Lyapunov exponents of dynamical systems using a spatial average. *Physical Review E*, 51(4):2844–2855, 1995.
- [22] Gary Froyland, Kevin Judd, Alistair I. Mees, and Kenji Murao. Lyapunov exponents and triangulations. In *Proceedings of the 1993 International Symposium on Nonlinear Theory and its Applications, Hawaii, December 1993*, volume 1, pages 281–286, 1993.
- [23] Gary Froyland, Kevin Judd, Alistair I. Mees, Kenji Murao, and David Watson. Constructing invariant measures from data. *International Journal of Bifurcation and Chaos*, 5(4):1181–1192, 1995.
- [24] Gary Froyland, Oliver Junge, and Gunter Ochs. Rigorous computation of topological entropy with respect to a finite partition. Submitted.
- [25] Harry Furstenberg and Yuri Kifer. Random matrix products and measures on projective spaces. *Israel Journal of Mathematics*, 46(1–2):12–32, 1983.
- [26] Pawel Góra. On small stochastic perturbations of mappings of the unit interval. *Colloquium Mathematicum*, 49(1):73–85, 1984.
- [27] Pawel Gora and Abraham Boyarsky. Absolutely continuous invariant measures for piecewise expanding C^2 transformations in R^N . *Israel Journal of Mathematics*, 67(3):272–286, 1989.
- [28] John Guckenheimer and Philip Holmes. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*, volume 42 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 1983.
- [29] R. Guder, M. Dellnitz, and E. Kreuzer. An adaptive method for the approximation of the generalized cell mapping. *Chaos, Solitons, and Fractals*, 8(4):525–534, 1997.
- [30] Rabbijah Guder and Edwin Kreuzer. Adaptive refinement of the generalized cell mapping. Preprint. 1998.
- [31] Fern Y. Hunt. A Monte Carlo approach to the approximation of invariant measures. *Random & Computational Dynamics*, 2(1):111–133, 1994.
- [32] Oliver Junge. *Mengenorientierte Methoden zur numerischen Analyse dynamischer Systeme*. PhD thesis, University of Paderborn, Paderborn, 1999.
- [33] Oliver Junge. Rigorous discretization of subdivision techniques. In *Proceedings of Equadiff '99, Berlin, August 1999.*, 1999.
- [34] M. Kac. On the notion of recurrence in discrete stochastic processes. *Bulletin of the American Mathematical Society*, 53:1002–1010, 1947.
- [35] Michael Keane, Rua Murray, and Lai-Sang Young. Computing invariant measures for expanding circle maps. *Nonlinearity*, 11(1):27–46, 1998.
- [36] Gerhard Keller. Stochastic stability in some chaotic dynamical systems. *Monatshefte für Mathematik*, 94:313–353, 1982.
- [37] Gerhard Keller and Carlangelo Liverani. Stability of the spectrum for transfer operators. Preprint, 1998.

- [38] Hannes Keller and Gunter Ochs. Numerical approximation of random attractors. Institut für Dynamische Systeme, Universität Bremen, Report Nr. 431, August 1998.
- [39] R.Z. Khas'minskii. Principle of averaging for parabolic and elliptic differential equations and for Markov processes with small diffusion. *Theory of Probability and its Applications*, 8(1):1–21, 1963.
- [40] Yuri Kifer. *Ergodic Theory of Random Transformations*, volume 10 of *Progress in Probability and Statistics*. Birkhäuser, Boston, 1986.
- [41] Yuri Kifer. *Random Perturbations of Dynamical Systems*, volume 16 of *Progress in Probability and Statistics*. Birkhäuser, Boston, 1988.
- [42] Andrzej Lasota and Michael C. Mackey. *Chaos, Fractals, and Noise. Stochastic Aspects of Dynamics*, volume 97 of *Applied Mathematical Sciences*. Springer-Verlag, New York, second edition, 1994.
- [43] Andrzej Lasota and James A. Yorke. On the existence of invariant measures for piecewise monotonic transformations. *Transactions of the American Mathematical Society*, 186:481–488, 1973.
- [44] Tien-Yien Li. Finite approximation for the Frobenius-Perron operator. A solution to Ulam's conjecture. *Journal of Approximation Theory*, 17:177–186, 1976.
- [45] Walter M. Miller. Stability and approximation of invariant measures for a class of nonexpanding transformations. *Nonlinear Analysis*, 23(8):1013–1025, 1994.
- [46] Walter M. Miller and Fern Y. Hunt. Approximation of attractors for finite dimensional maps by Ulam's method. Preprint, May 1996.
- [47] Rua Murray. Adaptive approximation of invariant measures. Preprint. 1998.
- [48] Rua Murray. Existence, mixing and approximation of invariant densities for expanding maps on R^d . Preprint. 1998.
- [49] Rua Murray. Approximation error for invariant density calculations. *Discrete and Continuous Dynamical Systems*, 4(3):535–557, 1998.
- [50] S. Pelikan. Invariant densities for random maps of the interval. *Transactions of the American Mathematical Society*, 281(2):813–825, 1984.
- [51] Mario Peruggia. *Discrete Iterated Function Systems*. A K Peters, Wellesley, 1993.
- [52] Karl Petersen. *Ergodic Theory*, volume 2 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1983.
- [53] David Ruelle. Ergodic theory of differentiable dynamical systems. *IHES Publications Mathematiques*, 50:275–320, 1979.
- [54] Leonard A. Smith. The maintenance of uncertainty. In G. Cini Castagnoli and A. Provenzale, editors, *Past and Present Variability of the Solar-Terrestrial system: Measurement, Data Analysis and Theoretical Models*, volume CXXXIII of *Proceedings of the International School of Physics*, pages 177–246, Amsterdam, 1997. Italian Physical Society.
- [55] Jaroslav Stark. Iterated function systems as neural networks. *Neural Networks*, 4:679–690, 1991.
- [56] S. M. Ulam. *Problems in Modern Mathematics*. Interscience, 1964.
- [57] Bodo Werner and N. Nicolaisen. Discretization of circle maps. *Zeitschrift für Angewandte Mathematik und Physik*, 49(6):869–895, 1998.
- [58] J. Wolfowitz. The moments of recurrence time. *Proceedings of the American Mathematical Society*, 18:613–614, 1967.
- [59] Lai-Sang Young. Recurrence times and rates of mixing. *Israel Journal of Mathematics*, 110:153–188, 1999.

Part III

Methods and Applications

Chapter 13

Formulas for the Eckmann-Ruelle Matrix

Timothy D. Sauer

ABSTRACT Determination of the local linearization information from experimental dynamical data is a key step in the methodology of attractor reconstruction. Because the dimension of the reconstruction space is typically higher than the original phase space dimension, some of the information in the reconstructed Jacobian, which we call the Eckmann-Ruelle matrix, reflects details of the embedding rather than the underlying dynamics. We establish formulas for the expected values of the entries of the Eckmann-Ruelle matrix, in both the presence and absence of observational noise.

13.1 Introduction

Mathematical modeling of a dynamical system from first principles can fail for many different reasons. In some situations basic principles are not sufficiently well known; in others the laws are known but there are too many free variables to analyze or simulate accurately. It has become standard practice in such cases to take dynamical measurements from the system and attempt to rebuild the system from these measured signals. For linear systems this is the goal of system identification methods in modern signal processing. For nonlinear systems, the geometry of the dynamics in the state space is no longer trivial, and plays a pronounced role. The fact that this geometry can be recovered from observations, and in some cases from a single scalar time series or spike train, is the focus of fundamental yet surprising theoretical work [10, 1, 19, 6, 15, 14].

Successful analysis of the original system in an alternate state space, often called the *reconstruction* or *embedding* space, is the basis of several techniques used for the exploitation of chaotic dynamics from measured data, including fixed point detection, noise reduction, control of chaos, and Lyapunov exponent calculation [8, 9]. These techniques share a crucial need for accurate first order descriptions of the reconstructed dynamics, which is represented by derivative information in the reconstruction space.

The usual representation of this information is in the form of an $m \times m$ matrix corresponding to each reconstructed state, where m is the dimension of the reconstruction space. The integer m is commonly called the *embed-*

ding dimension. This matrix is designed to be the best local approximation of the Jacobian of the time- T map in the m -dimensional reconstruction space. The time- T map represents the action of the reconstructed dynamics under a particular fixed time interval T . In most applications, the embedding dimension is greater than the dimension of the underlying dynamics, and a significant problem arises with the definition of this matrix. Because the reconstructed map is defined only on a lower-dimensional subset of the reconstruction space, partial derivatives cannot be taken in all m directions, and the Jacobian as traditionally defined does not exist. This does not mean that an $m \times m$ “best linearization” matrix cannot be defined and determined by least squares or other means. However, the matrix may lack the usual properties of a Jacobian.

There are two ways to work around the unfortunate fact that the linearization matrix is not a Jacobian. First, one may try to restrict attention to the infinitesimal directions in reconstruction space that do come from the underlying measured dynamics. Second, one may choose a way to calculate the $m \times m$ local linearization matrix and proceed as if it were a Jacobian, isolating and later discarding the computational artifacts that arise from this assumption. Which of the two approaches is more agreeable depends on the application. The first approach has been taken by many workers in the area [2, 18, 3, 13, 4, 20]. The second approach was originally suggested by Eckmann and Ruelle [6, 5] and Sano and Sawada [12] for the purpose of extracting Lyapunov exponents from experimental data. In this application, the $m \times m$ matrices representing local linearizations are successively multiplied together along the trajectory in reconstruction space, (as Jacobians would if the underlying dynamics were known) for the purpose of calculating the average stretching rates that Lyapunov exponents measure. Eckmann and Ruelle propose approximating the matrices from data using ordinary least squares. We will call the local linearization matrix defined in this way the *Eckmann-Ruelle matrix*, or simply the ER-matrix. Computational artifacts are unavoidable when taking this approach. For example, if the goal is to extract Lyapunov exponents, the method returns m exponents even when the underlying dynamics has a lower dimension. These extra exponents are known in the literature as *spurious* Lyapunov exponents [11, 16].

This chapter investigates the consequences of the second approach and summarizes recent work on describing the ER matrix in terms of the original dynamics and the properties of the reconstruction. The entries of the ER matrix also depend critically on the level of noise with which the measurements are collected, so the description of the ER matrix will be in terms of the expected values of the matrix entries, averaged over typical noise realizations. In particular, there are two useful asymptotic formulas, which apply in the cases of “low” and “medium” noise, respectively.

13.2 Attractor Reconstruction

Attractor reconstruction consists of observing the current state x of a system by a number m of measurements. Assume that we are trying to gather information on dynamics on an n -dimensional state space R^n . Denote by f the map describing the dynamics, so that state x moves deterministically to state $f(x)$ during a fixed time interval. We will denote by $P(x) = (p_1(x), \dots, p_m(x))$ an m -dimensional vector, where p_1, \dots, p_m are measurement functions. Under certain genericity conditions (see for example [15]) the image under P of the attractor states in m -dimensional Euclidean space can be shown to be topologically equivalent to the original dynamical phase space attractor. The set of image vectors can be studied for geometrical and dynamical properties, leading to a variety of methodologies for analyzing and exploiting the system, as mentioned earlier.

Early studies treated the static properties of the dynamical attractor S , such as fractal dimension [7]. In the seminal article [6], Eckmann and Ruelle suggested reconstructing the *dynamics* on $P(S)$. Assuming that P is a one-to-one correspondence on S (that $x \neq y$ implies $P(x) \neq P(y)$), there is an induced map $F = PfP^{-1}$ on $P(S)$ which maps $P(x)$ to $P(f(x))$. This leads to the following commuting diagram of maps

$$\begin{array}{ccccc}
 & & F & & \\
 & P(S) & \longrightarrow & R^m & \\
 P & \uparrow & & \uparrow & P \\
 & S & \longrightarrow & R^n & \\
 & & f & &
 \end{array} \tag{13.1}$$

where F is defined on the subset $P(S)$ of R^m . Then the map F on the reconstructed phase space attractor $P(S)$ is well-defined and can be learned from time-ordered measurements on the state space using statistical techniques. A key goal of applications is to learn the dynamical reconstruction map F accurately enough that local linearization information can be ascertained. In other words, we would like to ascertain the linear map $A(x)$ which makes the diagram

$$\begin{array}{ccccc}
 & & A(x) & & \\
 DP(x) & \begin{array}{c} R^m \\ \uparrow \\ R^n \end{array} & \longrightarrow & \begin{array}{c} R^m \\ \uparrow \\ R^n \end{array} & DP(f(x)) \\
 & & & & Df(x)
 \end{array} \tag{13.2}$$

commute at the tangent space level, for each state x . If F were defined on an m -dimensional domain around $P(x)$, the Jacobian DF at $P(x)$ would exist, and setting $A = DF$, (13.2) would follow immediately from (13.1) by the chain rule of calculus. However, in the common case that the domain $P(S)$ of F has dimension less than m , a choice must be made in how to define

the replacement for the Jacobian DF to make (13.2) hold. The solution of [6] is to use the best local linear approximation to F at $P(x)$, in the sense of least squares. Henceforth we will refer to this choice as the *Eckmann-Ruelle matrix* at x (or at $P(x)$). (Strictly speaking, the “Eckmann-Ruelle matrix” is an abuse of terminology, since the matrix depends on the choices of bases in the domain and range tangent spaces. More properly, it is the “Eckmann-Ruelle linear function”, which has a matrix representation in any given set of bases.) With the Eckmann-Ruelle matrix $A(x)$ defined in the limit as the neighborhood over which the least squares optimization is done shrinks to x , and assuming $m > n$, it can be shown that the diagram (13.2) commutes.

13.3 The Noiseless Case

In the absence of noise, we can often explicitly determine the Eckmann-Ruelle matrix. To simplify the analysis, first assume that the dynamical system f is a one-dimensional map and that the image under the measurement function P is a curve in R^m . To model the local linear behavior of F at a typical $P(x)$, note that any point on the curve near $P(x)$ can be written as $P(x+h)$ for some h . Since by definition $FP = Pf$, the Taylor expansion of the nearby point

$$P(x+h) = P(x) + hP'(x) + \frac{h^2}{2}P''(x) + \dots + \frac{h^m}{m!}P^{(m)}(x) + O(h^{m+1}) \quad (13.3)$$

must be mapped by F to

$$Pf(x+h) = Pf(x) + h(Pf)'(x) + \frac{h^2}{2}(Pf)''(x) + \dots + \frac{h^m}{m!}(Pf)^{(m)}(x) + O(h^{m+1}). \quad (13.4)$$

Here we are using the notation $P^{(m)}(x)$ to denote the m th derivative of the vector-valued function P evaluated at x . The m -dimensional vectors $P'(x), \dots, P^{(m)}(x)$ will span R^m in generic situations.

As shown in [16], because (13.3) and (13.4) hold for several different h , each power series term in (13.3) must map to the term of corresponding degree in (13.4). In general, let A be an $m \times m$ matrix for which (1) $A(\sum_{j=1}^m b_{ij}\vec{u}_j) = \sum_{j=1}^m b_{ij}\vec{v}_j$ for $i = 1, \dots, m$ and (2) $B = (b_{ij})$ is a nonsingular matrix. Then $A\vec{u}_j = \vec{v}_j$ for $j = 1, \dots, m$. Using this fact with B equal to the nonsingular matrix with rows $[h, \frac{1}{2}h^2, \dots, \frac{1}{m!}h^m]$ for m distinct nonzero values of h , we conclude that in the limit as $h \rightarrow 0$, the Eckmann-Ruelle matrix $A(x)$ multiplies the vector $P^{(i)}(x)$ to $(Pf)^{(i)}(x)$ for $i = 1, \dots, m$. This fact motivates defining the m -dimensional basis $\vec{u}_i = P^{(i)}(x), i = 1, \dots, m$ as the *canonical embedding coordinate basis* at $P(x)$, and the corresponding basis $\vec{v}_i = P^{(i)}(f(x))$ at $Pf(x)$. This choice

of basis simplifies the representation of the Eckmann-Ruelle matrix. For $m = 2$ the matrix $M(x)$ maps

$$\begin{aligned} \vec{u}_1 &\rightarrow (Pf)'(x) = P'(f(x))f'(x) = f'(x)\vec{v}_1 \\ \vec{u}_2 &\rightarrow (Pf)''(x) = P'(f(x))f''(x) + P''(f(x))f'(x)^2 \\ &= f''(x)\vec{v}_1 + f'(x)^2\vec{v}_2 \end{aligned} \tag{13.5}$$

and so in this basis the Eckmann-Ruelle matrix is upper-triangular:

$$A(x) = \begin{pmatrix} f'(x) & f''(x) \\ 0 & f'(x)^2 \end{pmatrix}. \tag{13.6}$$

In the low noise limit, attempts to reconstruct the best linear approximation of F at $P(x)$ will converge to a matrix that is equivalent (by similarity transformation) to the Eckmann-Ruelle matrix $A(x)$ in (13.6), as the neighborhood size goes to zero.




FIGURE 13.1. Matrix elements a_{11} (bottom trace) and a_{12} (top trace) of the ER matrix reconstructed from data in embedding dimension 2, at the fixed point $(3/4, 3/4)$ of the $3 \sin x$ map, plotted as a function of the measurement noise amplitude δ . The two sides of the figure represent the two alternative formulas. The left side of the diagram corresponds to the noisefree formula (13.8), which predicts values $a_{11} \approx 1.86$ and $a_{12} \approx 7.43$. The right side shows a plateau over two noise decades for which the formula (13.23) predicts expected values -1.55 and 0.79 . Error bars denote standard error. The neighborhood size is $\epsilon = .01$.

Let us specialize to the one-dimensional map $f(x) = 3 \sin x$ to illustrate this result with a computer experiment. The map f has an unstable fixed point at approximately $x_0 = 2.2789$. A simple delay-coordinate reconstruction function P can be defined by $P(x) = (f(x), x)$. The 2×2

Eckmann-Ruelle matrix $A(x_0)$ in these coordinates is of form

$$A(x_0) = \begin{pmatrix} a_{11} & a_{12} \\ 1 & 0 \end{pmatrix}. \quad (13.7)$$

where a_{11} and a_{12} best fit $f^2(x) = a_{11}f(x) + a_{12}x + c$ for x near x_0 . Formula (13.6) gives the matrix $A(x_0)$ represented in the domain basis $(P'(x_0), P''(x_0)) = ((f'(x_0), 1), (f''(x_0), 0))$ and range basis $(P'(x_1), P''(x_1)) = ((f'(x_1), 1), (f''(x_1), 0))$, where $x_1 = f(x_0)$. In the case of a fixed point, $x_1 = f(x_0) = x_0$. Changing to the elementary basis $((1, 0), (0, 1))$ yields

$$\begin{aligned} & \begin{pmatrix} f'(x_1) & f''(x_1) \\ 1 & 0 \end{pmatrix} \begin{pmatrix} f'(x_0) & f''(x_0) \\ 0 & f'(x_0)^2 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ \frac{1}{f''(x_0)} & -\frac{f'(x_0)}{f''(x_0)} \end{pmatrix} \\ &= \begin{pmatrix} f'(x_0) + \frac{f''(x_1)}{f''(x_0)}f'(x_0)^2 & -\frac{f''(x_1)}{f''(x_0)}f'(x_0)^3 \\ 1 & 0 \end{pmatrix}. \end{aligned} \quad (13.8)$$

Evaluating at the fixed point x_0 gives approximately

$$\begin{pmatrix} 1.8557 & 7.4274 \\ 1.0000 & 0.0000 \end{pmatrix}. \quad (13.9)$$

Figure 13.1 shows the result of a computer experiment in which the standard least squares approach [6, 12] is applied to embedded data in a neighborhood of $P(x_0)$ to approximate the 2×2 Eckmann-Ruelle matrix $A(x_0)$ in elementary coordinates. The method of ordinary least squares is used to fit a first-order model of the map which moves the reconstructed state $P(x_0)$ to its reconstructed image. The entries a_{11} and a_{12} of $A(x_0)$ are plotted as a function of observational noise added to the embedded data. For very low noise, the derivation gives essentially the correct answer. For noise greater than 10^{-6} , the ER matrix entries are not well approximated by the formula. In this noise range, the calculations we did earlier with terms of size h^2 are no longer valid and we must re-derive the matrix using the noise terms.

13.4 The Medium Noise Case

In the case of dynamics measured with observational noise, it is possible to carry out a direct analytic calculation of the expected values of the matrix entries that are reconstructed by the Eckmann-Ruelle procedure, as usually implemented using least squares fits. Although the formulas we derive for the ER matrix converge to the noiseless values determined earlier as the noise approaches zero, the convergence occurs only for extremely small noise (namely, significantly smaller than the square of the size of the neighborhood over which the least squares approximation is done). This is

shown in Figure 13.1. For noise of a more practical size, the expected value of the Jacobian entries plateau to values quite different from those for the noiseless case.

We will call noise of variance δ^2 “low” if $\delta \ll \epsilon^2$, “medium” if $\epsilon^2 \ll \delta \ll \epsilon$, and “high” if $\delta \approx \epsilon$. Jacobian reconstruction in the low noise case was discussed earlier. The high noise case, with random fluctuations as large as the neighborhood size, is unlikely to reveal significant local linearization information. In [17] an explicit calculation of the expected values of the Eckmann-Ruelle matrix in the “medium noise” case was carried out for some representative cases. As above, denote the deterministic m -dimensional embedding vector by $P(x) = (p_1(x), \dots, p_m(x))$, where p_1, \dots, p_m are measurement functions. A measurement with observational Gaussian noise is expressed as $P(x) + \delta N(0, 1)$, where δ^2 is the variance of the noise and $N(0, 1)$ denotes the standard normal distribution. The Eckmann-Ruelle matrix $A(x_0)$ at reconstructed state $P(x_0)$ is the linear part of the linear-plus-constant function which best fits the mapping $(p_1(x) + y_1, \dots, p_m(x) + y_m)$ to $(p_1(f(x)) + z_1, \dots, p_m(f(x)) + z_m)$ in a small neighborhood N of $P(x_0)$, where the y_k, z_k are independent Gaussian random variables in $\delta N(0, 1)$.

An ordinary least squares fit for the i th row (a_{i1}, \dots, a_{im}) of the ER-matrix $A(x_0)$ reduces to choosing parameters $a_{i1}, \dots, a_{im}, c_i$ that minimize

$$\int_{\mu} \left[\sum_{k=1}^m a_{ik} (p_k(x) + y_k) + c_i - p_i(f(x)) - z_i \right]^2 dx \tag{13.10}$$

where μ denotes the natural measure of the attractor restricted to states x in a small neighborhood of x_0 . A Taylor expansion around x_0 allows us to write

$$\begin{aligned} & \sum_{k=1}^m a_{ik} (p_k(x) + y_k) + c_i - p_i(f(x)) - z_i \\ &= \sum_{k=1}^m a_{ik} p_k(x_0) + c_i - p_i f(x_0) \\ &+ \left[\sum_{k=1}^m a_{ik} \nabla p_k(x_0) - \nabla (p_i f)(x_0) \right] \cdot (x - x_0) + \sum_{k=1}^m a_{ik} y_k - z_i \\ &+ O(|x - x_0|^2). \end{aligned} \tag{13.11}$$

According to the medium noise assumption $\epsilon^2 \ll \delta$, we can neglect the last term of (13.11). The parameter c_i is a free variable in the first term, so that it remains to minimize the variance of the middle term of the right-hand side of (13.11), for x chosen from the probability distribution of points in the neighborhood of x_0 . Assume that the distribution of x is uniform in a spherical neighborhood centered at x_0 of radius ϵ/l , where l is a Lipschitz

constant representing the expansiveness of P in the neighborhood of x_0 . In most cases, we expect that deviations from uniformity will be on the order of other terms we have neglected. Defining $h = x - x_0$, to minimize the expectation

$$E\{|| (A(DP) - D(Pf))h + Ay - z ||^2\} \tag{13.12}$$

from (13.11), we set the partial derivatives of (13.12) with respect to a_{ij} equal to zero. The derivatives pass inside the expectation integral, yielding the matrix equation

$$\begin{aligned} 0 &= E\{[(A(DP) - D(Pf))h + Ay - z][DP h + y]^T\} \\ &= E\{[(A(DP) - D(Pf))h][DP h]^T\} + E\{Ayy^T\} \\ &= [A(DP) - D(Pf)]E\{hh^T\}(DP)^T + \delta^2 A \end{aligned} \tag{13.13}$$

where all derivatives are evaluated at x_0 . (Here we have used the fact that for a matrix C and vector w , the gradient of the dot product $(Cw)^T(Cw)$ with respect to the matrix entries c_{ij} , $1 \leq i, j \leq m$, is the matrix $2Cww^T$.) Using the uniformity assumption, we calculate $E\{hh^T\} = (\epsilon^2/3nl^2)I$, where I denotes the $n \times n$ identity matrix, and so

$$0 = [A(DP) - D(Pf)](DP)^T + \frac{3n\delta^2 l^2}{\epsilon^2} A. \tag{13.14}$$

Therefore the expected values of the entries of the ER matrix are given as the solution A to the $m \times m$ matrix equation

$$A[(DP)(DP)^T + \rho I] = D(Pf)(DP)^T, \tag{13.15}$$

where $\rho = 3nl^2(\delta/\epsilon)^2$. Equation (13.15) can be thought of as a set of “normal equations” for the ER matrix. The m^2 scalar equations making up this matrix equation are

$$\sum_{k=1}^m a_{ik} \nabla p_k(x_0) \cdot \nabla p_j(x_0) + \rho a_{ij} = \nabla(p_i f)(x_0) \cdot \nabla p_j(x_0) \tag{13.16}$$

for $1 \leq i, j \leq m$.

To illustrate the use of formula (13.15), we consider the case of a one-dimensional map f . Then $n = 1$, $(DP)^T(DP)$ is a scalar, and one checks by direct substitution that

$$A = [\rho + (DP)^T(DP)]^{-1} D(Pf)(DP)^T \tag{13.17}$$

solves (13.15). The expected value of the ij th entry of the reconstructed Eckmann-Ruelle matrix A is therefore

$$a_{ij} = \frac{(p_i f)'(x_0)p'_j(x_0)}{3l^2(\delta/\epsilon)^2 + \sum_{k=1}^m p'_k(x_0)^2}. \tag{13.18}$$

In the medium noise case where $\delta \ll \epsilon$, we may neglect the first term in the denominator which leaves

$$a_{ij} = \frac{(p_i f)'(x_0) p'_j(x_0)}{\sum_{k=1}^m p'_k(x_0)^2}. \tag{13.19}$$

The important special case of time delay coordinates corresponds to $p_k = g f^{m-k}$ for a scalar observation function g . Here f^{m-k} denotes the $(m-k)$ -fold composition of the map f . In this special case all but the top row of the reconstructed Jacobian is trivial (for $i > 1$, row i is all zeros except for $a_{i,i-1} = 1$). For delay coordinates, (13.19) translates to

$$a_{1j} = \frac{(g f^m)'(x_0) (g f^{m-j})'(x_0)}{\sum_{k=1}^m (g f^{m-k})'(x_0)^2}, \tag{13.20}$$

which can be rewritten as

$$\frac{g'(x_m) f'(x_{m-1}) \cdots f'(x_0) g'(x_{m-j}) f'(x_{m-j-1}) \cdots f'(x_0)}{\sum_{k=1}^m [g'(x_{m-k}) f'(x_{m-k-1}) \cdots f'(x_0)]^2}, \tag{13.21}$$

where we have denoted the f -trajectory of x_0 by $x_k = f^k(x_0)$ for all k .

Formulas for the reconstructed Jacobian in the case where $\delta = 0$, the noiseless case, were given in the previous section. They differ from the equation (13.15), which is also a limit as $\delta \rightarrow 0$ but restricted to the case $\epsilon^2 < \delta$, which allowed us to neglect the higher-order terms of (13.11). In the development of the previous section, the absence of noise meant that the higher order terms could not be neglected, leading to a very different formula for the noiseless case.

To clarify this difference, we compare the two noise cases for the example of the previous section, the ER-matrix at the fixed point of the one-dimensional map $f(x) = 3 \sin x$. Again assume delay coordinates are used with observation function $g(x) = x$ and embedding dimension $m = 2$. Evaluating (13.8) in the noiseless case yielded (13.9). In the medium noise case, formula (13.21) results in the ER-matrix

$$\begin{pmatrix} \frac{f'(x_0)^2 f'(x_1)}{1+f'(x_0)^2} & \frac{f'(x_0) f'(x_1)}{1+f'(x_0)^2} \\ 1 & 0 \end{pmatrix}. \tag{13.22}$$

At the fixed point $x_0 \approx 2.2789$ of $f(x) = 3 \sin x$, (13.22) evaluates to

$$\begin{pmatrix} -1.5452 & 0.7920 \\ 1.0000 & 0.0000 \end{pmatrix} \tag{13.23}$$

for the ER-matrix. Both formulas are consistent with the underlying dynamics. The matrices in (13.9) and (13.23) share the eigenvector $(.8899, -.4561)$ corresponding to the eigenvalue -1.9511 coming from the underlying one-dimensional fixed point. However, they differ in the details of the

embedding, and therefore in the other, spurious, eigenvalue. Figure 13.1 shows computational results of least squares fits for the entries of the matrices in the zero noise limit, and medium noise range, respectively. The top rows of the two matrices in (13.9) and (13.23) can be verified in the left side, respectively right side of the figure.

13.5 Spurious Lyapunov Exponents

As discussed in the introduction, a side-effect of attractor reconstruction when the embedding dimension m is greater than the dimension n of the original dynamics is the production of $m - n$ spurious Lyapunov exponents. With formulas for the Eckmann-Ruelle matrix we can hope to explain the origin of the spurious exponents. In the noiseless case, for some values of m and n , there are simple formulas for the spurious exponents; in the noisy case they are more complicated.

Calculation of the Lyapunov exponents of the reconstruction map F on R^m is equivalent to computing the growth rates of the singular values of the Eckmann-Ruelle matrix for F^k . To find the best linear approximation of F^k at x , the Eckmann-Ruelle matrices are composed to yield $A(f^{k-1}(x)) \cdots A(f(x))A(x)$, the matrix representing local dynamics. We begin with noiseless observations. As an example, in the $m = 2, n = 1$ case of (13.6), there is a consistent set of tangent space bases along the trajectory for which the matrix representation of the ER-matrix is upper-triangular. This allows a simple derivation of the m growth rates of the $m \times m$ ER-matrices, using the following well-known fact. Assume that A_k is an $m \times m$ upper triangular matrix for $k = 1, 2, \dots$, and define $S_n = A_n \cdots A_1$. Assume that the magnitudes of the entries of A_k are bounded independent of k , and that the diagonal entries of S_n have geometric growth rates $\gamma_1, \dots, \gamma_m$ as $n \rightarrow \infty$. (The geometric growth rate of a sequence of real numbers $\{r_k\}$ is $\gamma = \lim_{k \rightarrow \infty} \frac{\ln |r_k|}{k}$.) Then there exist vectors $\vec{v}_1, \dots, \vec{v}_m$ such that $\|S_n \vec{v}_i\|$ has growth rate γ_i for each $i = 1, \dots, m$. In our context, the A_k are Eckmann-Ruelle matrices along an orbit, and the growth rates γ_i are apparent (i.e., computed) Lyapunov exponents. It follows from (13.6) that calculation will yield the correct Lyapunov exponent $\lambda = \lim_{k \rightarrow \infty} \frac{1}{k} \ln f'(x_k) \cdots f'(x_1)$ as well as a spurious exponent 2λ . This can be verified from numerical computations shown on the left side of Figure 13.2.

For $n = 1$ and larger m , the story is similar. Extending the calculation (13.5) to a reconstruction in m dimensions yields an $m \times m$ upper-triangular Eckmann-Ruelle matrix $M(x)$ whose diagonal entries are $f'(x), f'(x)^2, \dots$,




FIGURE 13.2. Lyapunov exponents from a reconstruction of the one-dimensional map $f(x) = 3 \sin x$ in R^2 , with observation function $P(x) = (f(x), x)$. We used the standard procedure for such a calculation, making least squares fits of the Eckmann-Ruelle matrices over neighborhoods of constant size, and approximating the Lyapunov exponents on the basis of these reconstructed local linear dynamics using a running QR -factorization [6, 12, 5]. As the level of observational noise approaches zero, the apparent Lyapunov exponents converge to 0.46 and 0.92, consistent with (13.6). The medium range for noise corresponds to approximately $10^{-4} < \delta < 10^{-2}$ on the right side of the figure, where the Lyapunov exponents predicted by (13.25) are the original Lyapunov exponent 0.46 and the spurious exponent -1.06 .

$f'(x)^m$. For example, the $m = 3$ analogue of the matrix in (13.6) is

$$M(x) = \begin{pmatrix} f'(x) & f''(x) & f'''(x) \\ 0 & f'(x)^2 & 3f'(x)f''(x) \\ 0 & 0 & f'(x)^3 \end{pmatrix}. \quad (13.24)$$

Due to the upper-triangular form of the matrices in canonical embedding coordinates, it is possible to see that Lyapunov exponents in this basis (and therefore in any basis) will be $\lambda, \dots, m\lambda$. Further examples of spurious Lyapunov exponents for larger n in the noiseless case are given in [16].

In order to study the spurious Lyapunov exponents that arise along a chaotic trajectory observed with noise, it is more convenient to change to a new coordinate system which for the linear mapping of (13.22) is given by the basis $\{(f'(x_0), 1)^T, (1, 0)^T\}$ in the domain and $\{(f'(x_1), 1)^T, (1, 0)^T\}$ in

the range. The matrix in the new coordinates is in upper-triangular form:

$$\begin{pmatrix} 0 & 1 \\ 1 & -f'(x_1) \end{pmatrix} \begin{pmatrix} \frac{f'(x_0)^2 f'(x_1)}{1+f'(x_0)^2} & \frac{f'(x_0)f'(x_1)}{1+f'(x_0)^2} \\ 1 & 0 \end{pmatrix} \begin{pmatrix} f'(x_0) & 1 \\ 1 & 0 \end{pmatrix} \\ = \begin{pmatrix} f'(x_0) & 1 \\ 0 & \frac{-f'(x_1)}{1+f'(x_0)^2} \end{pmatrix}. \quad (13.25)$$

Since the range basis becomes the domain basis on the next iterate, we get the product of Eckmann-Ruelle matrices over several iterations by multiplying the upper-triangular matrices together. By the fact about growth rates mentioned above, the Lyapunov exponents formed by the Eckmann-Ruelle matrices will be the trajectory averages of log absolute value of the two diagonal elements of the upper-triangular matrices.

The first Lyapunov exponent will be the trajectory average of $\log |f'(x_k)|$, the Lyapunov exponent λ of the current f trajectory. The second Lyapunov exponent will be the average over the trajectory of

$$\log \frac{|f'(x_{k+1})|}{1 + f'(x_k)^2} = \log |f'(x_{k+1})| - \log(1 + f'(x_k)^2). \quad (13.26)$$

The first part tends to λ , so the second Lyapunov exponent is λ minus the trajectory average of

$$\log(1 + f'(x)^2). \quad (13.27)$$

In particular, this shows that the spurious Lyapunov exponent is negative, and independent of the noise level δ as long as δ is in the medium range. Returning to the simple example of the $3 \sin x$ map, one can determine the average of $\log(1 + f'(x)^2)$ over a typical trajectory as ≈ 1.52 , resulting in a predicted spurious Lyapunov exponent of $0.46 - 1.52 \approx -1.06$. This prediction agrees with the standard Lyapunov exponent calculation using the Eckmann-Ruelle orthogonalization procedure, shown in Figure 13.2, over the medium noise range $.0001 < \delta < .01$.

13.6 Concluding Remarks

As mentioned in the introduction, the approach of this chapter is one way to acquire local linearization information, given the problem that the embedding dimension is greater than the original dimension of the dynamics. The other approach is to try to determine the tangent directions of the dynamics in reconstruction space and try to avoid artifacts like spurious Lyapunov exponents in that way. Techniques exist to locally project the dynamical structure in embedding space to these special directions and determine the appropriate linearizations there. Neither approach is foolproof.

The formulas given in this chapter for the Eckmann-Ruelle matrix are a beginning, but they need further work to yield easily computable formulas for all possible combinations of embedding dimension m and dynamics dimension n . Regarding the noisy case, formula (13.15) does apply for all m and n but writing explicit formulas like (13.19) for the entries has only been done for small n . More precise information may help decide the truth of the conjecture that if the ER-matrices are determined by least squares, in the noisy case (medium noise), spurious Lyapunov exponents are negative. For example, in the case $n = 1$, $m = 2$ this follows from (13.26).

In the noiseless case, formulas for upper-triangular ER-matrices that allow the true and spurious Lyapunov exponents to be read easily are known only for pairs m, n where

$$m = \sum_{d=1}^k \binom{d+n-1}{n-1}$$

for some k . This includes all m if $n = 1$. If $n = 2$, it includes $m = 2$, $m = 2 + 3$, $m = 2 + 3 + 4$, etc. For example, it is shown in [16] that if we denote the true Lyapunov exponents by λ_1 and λ_2 in the case $n = 2$, $m = 5$, the spurious Lyapunov exponents will be $2\lambda_1$, $\lambda_1 + \lambda_2$, and $2\lambda_2$.

Finally, it is clear, because of the error-in-variables problem and other reasons, that ordinary least squares is a suboptimal method for determining Eckmann-Ruelle matrices from data. However, it appears that practitioners will continue to use it until there is a demonstrably superior alternative. It is possible that the correct alternative will be specific to the application the user has in mind.

References

- [1] D. Aeyels. Generic observability of differentiable systems. *SIAM J. Control and Optimization*, 19:595–603, 1981.
- [2] D. S. Broomhead, R. Jones, and G. P. King. Topological dimension and local coordinates from time series. *J. Phys. A*, 20:L563–L569, 1987.
- [3] R. Cawley and G. H. Hsu. Local-geometric projection method for noise reduction in chaotic maps and flows. *Phys. Rev. A*, 46:3057–3082, 1992.
- [4] A. G. Darbyshire and D. S. Broomhead. Robust estimation of tangent maps and Lyapunov spectra. *Physica D*, 89:287–305, 1996.
- [5] J.-P. Eckmann, S. O. Kamphorst, D. Ruelle, and S. Ciliberto. Liapunov exponents from time series. *Phys. Rev. A*, 34:4971–4979, 1986.
- [6] J.-P. Eckmann and D. Ruelle. Ergodic theory of chaos and strange attractors. *Rev. Mod. Phys.*, 57:617–656, 1985.
- [7] P. Grassberger and I. Procaccia. Measuring the strangeness of strange attractors. *Physica D*, 9:189–208, 1983.

- [8] H. Kantz and T. Schreiber. *Nonlinear Time Series Analysis*. Cambridge University Press, 1997.
- [9] E. Ott, T. Sauer, and J. A. Yorke, editors. *Coping with chaos; Analysis of chaotic data and the exploitation of chaotic systems*. Wiley & Sons Inc., New York, 1994.
- [10] N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw. Geometry from a time series. *Phys. Rev. Lett.*, 45:712–716, 1980.
- [11] U. Parlitz. Identification of true and spurious Lyapunov exponents from time series. *Int. J. of Bifurcation and Chaos*, 2:155–165, 1992.
- [12] M. Sano and Y. Sawada. Measurement of Lyapunov spectrum from a chaotic time series. *Phys. Rev. Lett.*, 55:1082–1085, 1985.
- [13] T. Sauer. A noise reduction method for signals from nonlinear systems. *Physica D*, 58:193–201, 1992.
- [14] T. Sauer. Reconstruction of dynamical systems from interspike intervals. *Phys. Rev. Lett.*, 72:3811–3814, 1994.
- [15] T. Sauer, J. A. Yorke, and M. Casdagli. Embedology. *J. Stat. Phys.*, 65:579–616, 1991.
- [16] T. D. Sauer, J. A. Tempkin, and J. A. Yorke. Spurious Lyapunov exponents in attractor reconstruction. *Phys. Rev. Lett.*, 81(20):4341–4344, 1998.
- [17] T. D. Sauer and J. A. Yorke. Reconstructing the jacobian from data with observational noise. *Phys. Rev. Lett.*, 83(7):1331–1334, 1999.
- [18] R. Stoop and J. Parisi. Calculation of Lyapunov exponents avoiding spurious elements. *Physica D*, 50:89–94, 1991.
- [19] F. Takens. Detecting strange attractors in turbulence. In D. A. Rand and L.-S. Young, editors, *Dynamical Systems and Turbulence (Warwick 1980)*, volume 898 of *Lecture Notes in Mathematics*, pages 366–381. Springer-Verlag, Berlin, 1980.
- [20] A. Schenk zu Schweinsberg, T. Ritz, U. Dressler, B. Hübinger, R. Doerner, and W. Martienssen. Quasicontinuous control of a bronze ribbon experiment using time-delay coordinates. *Phys. Rev. E*, 55(3):2145–2158, 1997.

Chapter 14

Noise and Nonlinearity in an Ecological System

Paul A. Dixon¹

Maria J. Milicich

George Sugihara

ABSTRACT In this chapter, a case study from marine ecology is presented in which the application of techniques from nonlinear time series analysis is shown to provide insight into the interplay between stochastic physical forcing and nonlinear biological response in a natural system. Specifically, the replenishment of a population of reef fishes is analyzed in detail, and important nonlinearities are demonstrated in the processes underlying variability in the supply of larval propagules to the reef. This information is used to guide the construction of a series of models which attempt to forecast larval supply from readily measured physical variables. The most successful models are those that account for the nonlinearities in the response of larvae to their physical environment. Such models provide better forecasts than can be achieved with conventional linear techniques and identify processes hidden to linear analysis. The importance of understanding the interplay between noise and nonlinearity in ecological systems is discussed.

14.1 Introduction

A fundamental and largely unresolved problem in ecology is the relative contribution of exogenous and endogenous processes to temporal variability in the size of natural populations. Motivations for solving this problem are case-specific and range from purely scientific interest to the need for sensible management decisions, but the ultimate goal is generally the same: To make predictions about future states of natural systems from measurable initial conditions. To do this requires information about the identities of both the biological and physical variables that comprise the system and the functions that describe their interactions. Broadly, there are two approaches to obtaining such information. One can proceed from the bottom-up by first writing equations believed to capture the essence of the system's dynamics and then fitting the free parameters of the resulting

¹Author for correspondence.

model from observational data. This necessitates sufficient prior knowledge to choose appropriate state variables and the forms of the functions which relate them. The alternative is the top-down approach, for which is required the ability to identify relevant variables and their interrelations directly from observations, is required to deduce the underlying processes at work. In either case the exercise is closed and judged a success when satisfactory predictions can be made and the ultimate limits to predictability are understood.

The ubiquitous observation is that natural populations do not remain constant over time, but rather change in what often seems a random fashion. The manner in which population biologists have thought about and attempted to explain this fact has also changed over time, in a pattern characteristic of ecology as a whole: first one set of processes, external forcing, was thought of primary importance (e.g., Nicholson 1933[28], Andrewartha and Birch 1954[1], Lack 1954[16])², followed by a shift in emphasis to a second set, internal dynamics, (e.g., May 1974[23], Schaffer 1985[33], Schaffer and Kot 1986[34]), and finally, when neither exclusive view proved satisfying, it became fashionable to acknowledge that both must play a role (e.g., Ellner and Turchin 1995[9]). However, tacit acknowledgement that a synthetic view is necessary is not the same as actually producing an effective synthesis. The aim here is to demonstrate that tools from nonlinear time series analysis, when applied judiciously, can help provide such a comprehensive understanding of the interplay between physical forcing and biological response, that is, between external and internal dynamics in ecological systems.

This chapter will describe in detail a case study (Dixon et al. 1999[7]) that is representative of a broad problem in fisheries biology: understanding and ultimately forecasting variability in fish stocks where population change is largely decoupled from population size. Fisheries assessors use various models to try to predict future replenishment from reproductive stock size (Rothschild 1986[31]), but such fitted curves have been plagued by extreme variance in recruitment, and explaining the scatter in these relationships has remained a largely unrealized goal throughout this century. A useful metaphor for those unfamiliar with these issues is to think of the life history of marine fishes as an input-output system (Fig. 14.1). Many marine species are characterized by a bipartite life history, in which larval forms possess behavioral and morphological characteristics associated with the challenges of a planktonic existence that may be completely

²In fact these authors were at the center of an earlier controversy in ecology, concerning agents of population control; that is, whether climatic (external) or biotic (internal) factors kept populations bounded in spite of their propensity for exponential increase. While the authors referenced here disagreed on this point, ecologists of the era did in general agree that, left alone, a population's internal dynamics would lead to equilibrium behavior, and accordingly that variability was indicative of external forcing.

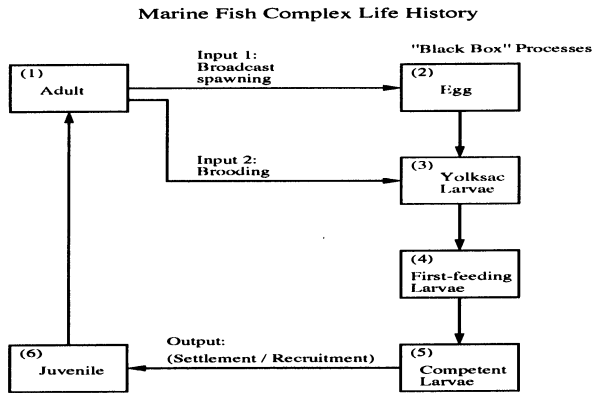


FIGURE 14.1. The bipartite life history common to many marine organisms, viewed as an input-output system. Density-dependence between population size and population change - that is, the relationship between the inputs and outputs of the system - may be interrupted by processes affecting mortality of eggs, larvae, and (in some cases) juveniles during their time in the plankton.

dissimilar from those of juveniles and adults. Investigating feedback control in a population is then equivalent to attempting to relate the system's output, recruitment³, to its input, the magnitude of adult spawning. Unfortunately, the larval phase in this schematic, which has the potential to decouple the inputs and outputs, is an essentially unobservable black box. Thus, while some researchers have attributed the unexplained variability to sampling error and the misuse of analytical techniques (e.g., Caputi 1988[2], Hilborn and Walters 1992[12]), others have suggested that unpredictable and high levels of mortality during the intervening larval phase will necessarily obscure the stock-recruitment relationship (e.g., Hjort 1914[13], Lasker 1981,[17], Cushing 1982[5]). Because an understanding of the fate of larval stages has remained elusive, it has proven difficult to evaluate the relative contributions of these sources of error.

The next two sections will adopt a dynamical systems approach to retrospectively investigate processes at work during the larval phase. The system that will be considered is Lizard Island, Australia, part of the Northern Great Barrier Reef, and the focus will be on a common and abundant family of reef fishes, the pomacentrids or damselfishes. Section 14.2 will begin by

³The term "recruitment" is often ambiguous and may be used in the literature to refer to any one of several points in a species life history (Richards and Lindeman 1987[30]). This chapter will focus on a population of reef fishes, and will therefore equate recruitment to settlement - the metamorphosis from a planktonic larval to a sedentary juvenile form. Many pelagic fish stocks possess a planktonic juvenile phase as well, which may be equally important in decoupling the relationship between population size and population change (Peterman et al. 1988[29]).

analyzing time series of spawning output, supply of mature larvae to the reef and subsequent juvenile abundance for one species of pomacentrids (boxes 2, 4, and 5 in Fig. 14.1) and will provide evidence for nonlinearity in the processes underlying variability in larval supply. This information will then be used in Section 14.3 to guide a search for the identities of the most important physical variables that drive this variability. After these variables have been identified, the functions which relate them to larval supply will be deduced. The result is a nonparametric, nonlinear model that affords better forecasts than linear models, and identifies the importance of events early in larval life which are hidden to linear analysis. This chapter will conclude by placing these results in the general context of understanding the interplay between stochastic physical forcing and nonlinear biological response in ecological systems.

14.2 Univariate Analysis: Nonlinearity in the Larval Phase

Pomacentrus amboinensis, the Ambon damsel, is a common shallow-water damselfish found in the eastern Indian and western Pacific Ocean, including the Great Barrier Reef. A unique set of time series are available for this species which describe contemporaneous temporal variability at three stages in life history: spawning output, larval supply, and the abundance of newly settled juveniles (Milicich et al. 1992[26], Meekan et al. 1993[24]). Briefly, spawning was estimated from daily visual census of eggs in nests; observations were made for 168 days over two consecutive spawning seasons. Juvenile abundance was monitored for 166 days over two seasons by scuba divers swimming visual transects over patch reefs. Larval supply was measured using two replicate light traps, which fished for three hours a night for 280 nights over three consecutive seasons. Raw data are given in Figure 14.2.

Three forecasting tests will now be applied to these data. The first two result in distinct pieces of evidence for nonlinearity in the processes affecting variability at the end of the larval phase: an enhanced success of nonlinear forecasting algorithms in predicting the data and the presence of structure in the residuals from linear forecasts of larval supply. Results of the third test, which examines the rate at which forecastability declines as a function of prediction time, give evidence for the stochastic nature of the processes underlying variability in larval supply.

14.2.1 *S-Map Analysis*

The first test for nonlinearity is an examination of the relative success of linear and nonlinear forecasting algorithms in predicting the data. For this,




FIGURE 14.2. The replenishment of *P. amboinensis*. Top panel gives spawning output measured by the number of eggs released from the nest, middle panel gives the supply of mature larvae to the reef, and bottom panel the number of juveniles on the reef. All measurements were made daily.

S-maps were used to make predictions (Sugihara 1994[36]). The data were embedded using lagged coordinates⁴. For the embedded time series $\mathbf{X}_t \in \mathbb{R}^{m+1}$, where the constant term in Eq. 14.2 below is given by $\mathbf{X}_t(0) \equiv \mathbf{1}$, and the time series value T_p steps forward is $\mathbf{X}_{t+T_p}(1) = \mathbf{Y}_t$, forecasts at T_p were given by:

$$\hat{\mathbf{Y}}_t = \sum_{j=0}^m C_t(j) \mathbf{X}_t(j). \quad (14.1)$$

For each predictee, \mathbf{X}_t , singular value decomposition was used to solve for C using the rest of the data set as follows:

$$\mathbf{B} = \mathbf{A} \mathbf{C} \quad (14.2)$$

where

$$\mathbf{B}_i = w(\|\mathbf{X}_i - \mathbf{X}_t\|) \mathbf{Y}_i \quad (14.3)$$

$$\mathbf{A}_{ij} = w(\|\mathbf{X}_i - \mathbf{X}_t\|) \mathbf{X}_i(j) \quad (14.4)$$

and

$$w(d) = \exp^{-\theta d/\bar{d}}. \quad (14.5)$$

The degree of nonlinearity in these maps is thus controlled by the tuning parameter θ , where $\theta = 0$ gives a global linear map, and theta increasingly positive yields increasingly local, nonlinear forecasts. Forecasts were made out-of-sample through leave-one-out cross-validation (that is, all vectors whose coordinates included any of the coordinates of the predictee were

⁴For the results given here, the time lag was held constant at one day; results were robust to the specific choice.

eliminated from model fitting) and were made one time step ahead, for the next day's value. Figure 14.3 plots forecastability as a function of model nonlinearity and embedding dimension (number of reconstructed variables used). The abscissa gives model nonlinearity, and the ordinate forecast success measured both in terms of correlation coefficient between predicted and observed values and average arithmetic error of the predictions.

Two features of this plot are readily apparent. First, regardless of which metric of model performance is used, nonlinear models with an embedding dimension of two or greater significantly (Z -test; P less than .05) outperform linear maps for the larval supply time series, whereas for the spawning and juvenile data this is not the case. The underlying philosophy here is that the characteristics of the model which gives the best out-of-sample forecasts most closely reflect the nature of the underlying dynamical processes; accordingly, these results may be taken as preliminary evidence that there are important nonlinearities underlying the variability in larval supply of this species at this particular location. Additionally, the overall level of predictability is different for the three time series; it is generally much harder to predict larval supply than spawning output or juvenile abundance.

14.2.2 *Residual Delay Maps*

A second piece of evidence for nonlinearity in larval supply results from examining the residuals from linear model forecasts, through the method of residual delay maps (RDMs) (Sugihara et al. 1999[37]). The basic idea is that if the variability in a time series arises strictly from linear, stochastic processes, then the residuals from linear model forecasts should be randomly distributed around zero. Alternatively, if there is nonlinear structure of sufficiently low dimension to be detected with the available data, the functional form of the nonlinearity — the fashion in which the residuals depart from a random distribution centered around zero — will likely become apparent. To examine this, residual delay maps plot the average residual from a forecast at time $t+1$ as a function of the average value of the raw data at time t (hence the term “delay map,” as tomorrow's error is plotted as a function of today's value). This technique for detecting nonlinearity in observed data has the advantage of computational simplicity and relatively modest data requirements. To reduce the variance and clarify any systematic structure in the residuals that may exist, the raw values of the time series are first sorted and binned. For each bin, the average residual from a linear forecast for the points in that bin is calculated and plotted against the mean value of the points in the bin. Residual delay maps thus employ a two-way averaging procedure. Figure 14.4 gives the results of this test. To facilitate comparison, raw data were first divided by their standard deviations; this normalizes the magnitudes of the residuals across the three data sets, but does not affect the relative position of the points in the figure. The residuals from linear (AR3; the choice of the number of lags was

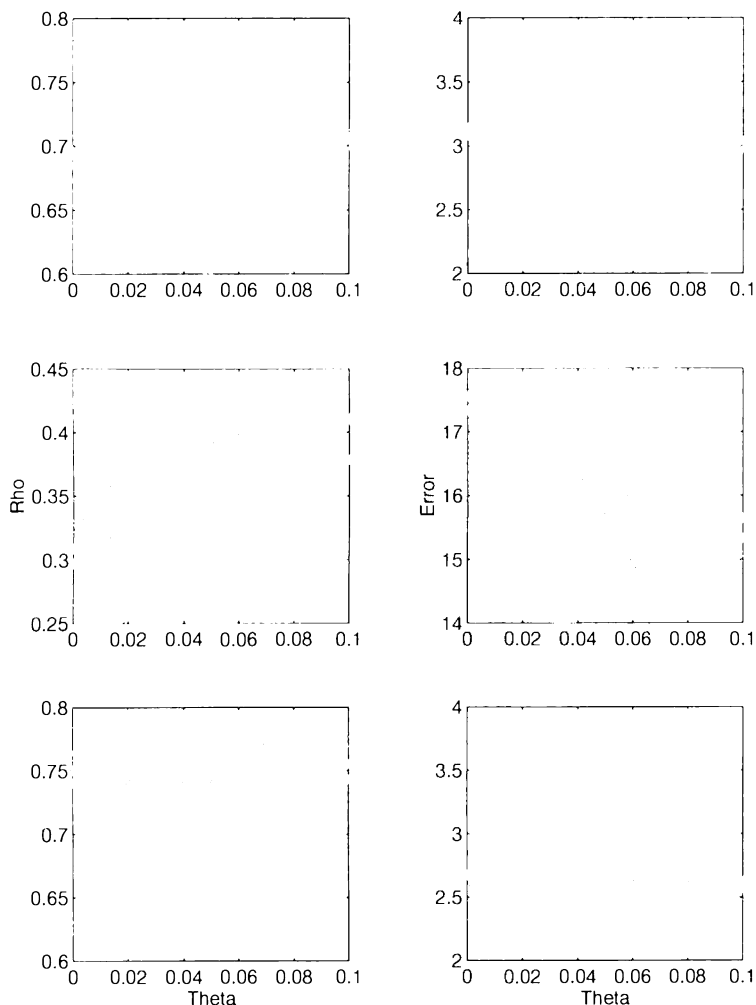


FIGURE 14.3. S-map analysis of spawning output (top), larval supply (middle), and juvenile abundance (bottom) for *P. amboinensis* at Lizard Island. The left panels give prediction success as a function of model nonlinearity as measured by the correlation coefficient between predicted and observed values, the right panels success as measured by the average error of the predictions. Circles: embedding dimension = 2; stars: $E = 3$; pluses: $E = 4$. Model nonlinearity is controlled by the tuning parameter θ , where $\theta = 0$ gives a global linear model, and θ increasingly positive gives increasingly nonlinear algorithms. Only for the larval supply data is a significant increase in predictability observed with nonlinear models.

not critical) forecasts of adult spawning output fluctuate evenly around zero, whereas the residuals for larval supply seem to display a systematic trend, with high values of larval supply today leading systematically to an overestimation of larval supply tomorrow. When juvenile abundance is considered, these patterns have broken down; there is one outlier residual, but no evidence for a systematic trend.

14.2.3 *Prediction Decay*

The final forecasting characteristic of interest is the rate at which predictability decays with increasing forecast horizon. Figure 14.5 (top panel) gives the results of this test for the *P. amboinensis* data. Again, forecasts are given by S-maps. For the spawning output and juvenile abundance time series, predictability decays steadily, in linear fashion, to essentially zero over five time steps (days) into the future. By comparison, larval supply data become unpredictable three days into the future. Beyond this basic difference, the test is ambiguous. It is not possible to label the rate of prediction decay as exponential (as opposed to linear) for the larval supply data, as there are only three points and the location of the second one (time horizon = two days) is sensitive to exactly how nonlinear a model one chooses. Indeed, this difference in predictability as a function of prediction time seems to reflect little more than the autocorrelation spectrum of the three data sets (Fig. 14.5, bottom panel). Because such strong short-term autocorrelation can potentially confound forecasting results, the analysis will now be extended to include formal hypothesis testing through the use of surrogate data.

14.2.4 *Surrogate Analysis*

There are a number of possible explanations as to why data may appear nonlinear to forecasting algorithms and related analytical techniques, some of which are trivial (e.g., the non-normality of the data itself may be responsible), and some of which are of real biological interest. The approach to distinguishing among these that will be adopted here is to proceed through the use of surrogate data (Theiler et al. 1992[39], Kantz and Schreiber 1997[14]). The idea is to create a large number of surrogate realizations of the data, each of which corresponds to a specific null hypothesis pertaining to the origin of the nonlinearity, and then to reanalyze the surrogates with the same forecasting techniques as were applied to the real data. By choosing a statistic which encapsulates the results of the analysis, a P-value can then be assigned to the null hypothesis from the percentage of the surrogates that yield a value of that statistic greater than or equal to that calculated from the original time series. Surrogate data techniques thus enable a transition from heuristic, descriptive evidence of nonlinearity to rigorous hypothesis testing.




FIGURE 14.4. Residual delay map analysis, for bin size = 10. Top panel: average error from a linear model of tomorrow's spawning output as a function of today's value. Middle: larval supply. Bottom: juvenile abundance. Only for the larval supply data is there a systematic departure from residuals evenly distributed around zero; following days of high larval supply, linear models consistently overestimate the true value.




FIGURE 14.5. Prediction decay (top panel) for *P. amboinensis* data. Larval supply (stars) is unpredictable past two days into the future, whereas prediction decay is slower for spawning output (circles) and juvenile abundance (pluses). This prediction characteristic reflects the autocorrelation spectra (bottom panel) of the data.

Figure 14.6 gives a sample surrogate time series for supply of larval pomacentrids to Lizard Island for each null hypothesis tested. In order of increasing complexity, these are:

Ho1 The data arise from an identical, independently distributed (I.I.D.) noise process.

Ho2 The data arise from a set of linear, stochastic processes.

Ho3 The data arise from a static, monotonic, nonlinear filter of a set of underlying linear, stochastic processes.

Creating surrogates is the most straightforward for the first null hypothesis: the time series is simply randomly shuffled. This is of interest because it tests whether the frequency distribution of the data itself is responsible for the observed nonlinearity. As has been seen, larval supply time series are episodic; daily values are not normally distributed. Indeed, this is the case for most biological properties in the ocean (Cassie 1963[3]). Because larval reef fishes have been demonstrated to be spatially coherent on scales of tens of kilometers (Doherty and Williams 1988[8]), one possible explanation for temporal variability in larval supply is that peaks in larval supply simply represent the chance advection of patches of larvae past the reef (Williams and English 1992[41]). An obvious question is thus whether the frequency distribution of the data alone is all that the forecasting tests have really measured. However, as will be seen in the analysis of the surrogates corresponding to Ho1 and Ho3 (described later), this is an insufficient explanation for the patterns observed. Variability alone does not necessarily imply nonlinearity.

Ho2 offers a direct test of the idea that variability in larval supply does not arise solely from the action of linear, stochastic processes. Here, surrogate data generation preserves the power spectrum of the data while randomizing the phases. Accordingly, the surrogates do not possess the same frequency distribution as the original time series.

Ho3 combines the features of the first two by preserving both the Fourier spectra and the frequency distribution of the real data. Analyzing these surrogates tests the idea that the nonlinearity is consistent with a monotonic, nonlinear transformation of an underlying set of linear, stochastic processes. This hypothesis is labeled “null” in deference to the dynamical systems community from whence these methods come, and for whom such an explanation for nonlinearity in a time series is not usually of great interest. However, to the ecologist, this explanation for the nonlinearity observed is a suggestive one, for a “nonlinear transformation of an underlying linear, stochastic process” is, as will be seen, a good starting point for a description of how biological entities respond to variability in their physical environment.

Here, R , the statistic that encapsulates the forecasting analysis, will be given by the difference in forecasting success between the linear and best

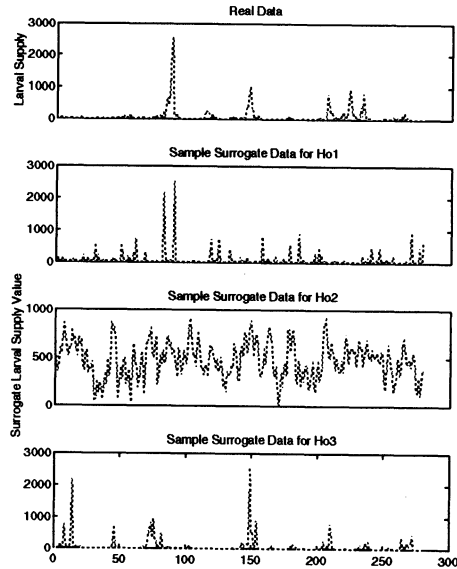


FIGURE 14.6. Raw data for supply of all pomacentrids to Lizard Island (top panel), and sample surrogate realizations corresponding to Ho1, Ho2, and Ho3 (top to bottom).

nonlinear algorithm, measured by the square of the correlation coefficient between predicted and observed values (percent variance explained). In principle, the choice of statistics is not critical; this particular R is intended to place the emphasis on successfully forecasting the times of greatest biological interest, the peaks in larval supply. The analysis to this point has centered on one species of pomacentrid; here, however, to decrease the number of zero values, larval supply of all pomacentrids will be considered. There is a strong correlation (ρ greater than 0.9) between supply of *P. amboinensis* and supply of all pomacentrids, so it seems unlikely that aggregating to a coarser degree of taxonomic resolution will create much difficulty.

Figure 14.7 gives graphical portraits of the distribution of R values for Ho1, Ho2, and Ho3, as a function of the embedding dimension. All three may be rejected for E greater than one. The non-normality of the data alone is not responsible for the improved forecasts seen with nonlinear models, and the data are not the result of a set of linear, stochastic processes. That Ho3 can also be rejected is intriguing, and, as will be seen in the next section, is probably due to the fact that, while variability in larval supply may originate from biological responses acting as a nonlinear transformation of an underlying set of linear, stochastic (physical) processes, there is no a priori reason why the transformation should be a monotonic one.

To this point, two suggestive pieces of evidence for nonlinearity of larval




FIGURE 14.7. Distribution of R-values corresponding to Ho1, Ho2, and Ho3 for supply of all larval pomacentrids to Lizard Island. All three hypotheses may be rejected with greater than 95% confidence for $E_i = 2$.

supply time series have been presented: enhanced out-of-sample forecasting success of nonlinear algorithms, and structure in the residuals from linear model forecasts. Additionally, the stochastic nature of larval supply has been demonstrated, as it is impossible to make predictions more than two days into the future. The picture that is beginning to emerge is that there are important nonlinearities in the processes at work in the larval phase, which suggests that a strong relationship between spawning patterns and patterns of juvenile abundance can no longer be expected. The next section will put this information to use in identifying the physical processes responsible for this variability in larval supply.

14.3 Multivariate Analysis: Forecasting Larval Supply

The discussion will now leave the confines of black-box modeling from reconstructed dynamical variables and begin forecasting larval supply directly from observations of the physical environment. The analysis to this point has suggested that to succeed in this task, forecasting models will have to take into account the nonlinearities in the interaction between larval fishes and their fluctuating physical environment. There are two potential aspects to this nonlinearity: the forms of the functions relating the biology to the physics (which need not be linear), and a nonlinear coupling between key physical variables in their total impact on larval supply. Whenever this second aspect is important, as will be the case here, it will not be possible to deduce the action of different physical variables of larval supply by considering them in isolation. Accordingly, many conventional techniques for analyzing ecological data, such as multivariate linear regression or principal component analysis, will be inapplicable to this problem. What is needed is a nonlinear approach.

14.3.1 Lunar Phase

The most obvious property of pomacentrid larval supply is the cyclic nature of the timing of supply peaks. This results from the entrainment of egg release from the nest to the lunar cycle, in conjunction with a relatively invariant larval duration, which otolith analysis indicates averages nineteen days (Wellington and Victor 1989[40], Meekan et al. 1993[24]). This deterministic feature of the data thus represents a logical place to begin constructing a model. The specific physical variable chosen here to represent the lunar phase is the percent illumination, calculated by multiplying the proportion of a full moon for each night by the elongation, or phase angle, for that night (Kopal 1971[15]). Statistics for all of the models that will be constructed in this section are summarized in Table 14.1.




FIGURE 14.8. Forecasting larval supply from lunar phase. Forecast success (Z-axis) is measured by the correlation coefficient between predicted and observed values, set to zero for negative values; it is shown as a function of lag and model nonlinearity. There is a peak in predictability at 19-20 days, in good agreement with the average age of the larvae, and there is little to be gained with nonlinear models.

Because the relationship between pomacentrid larval supply and lunar phase is well-established at Lizard Island (Milicich 1994[25]), this variable provides an opportunity to test a technique that will be used again shortly for another variable where the situation is not so clear. One standard (and simple) method for detecting cycles in ecological data is to calculate the autocorrelation at increasing time lags (Finerty 1980[10]); if a cycle is present, the autocorrelation will be a maximum when the time lag corresponds to the cycle's wavelength and a minimum at half the wavelength. A similar approach will be adopted here: Larval supply will be predicted from percent illumination at increasing time lags, from one to thirty days in the past. The twist is that, in addition to exploring predictability as a function of lag time, it will be evaluated with respect to model nonlinearity. Adding this dimension to the analysis results in the surface plot in Figure 14.8. As before, forecasts were given by S-maps, and success was measured by the correlation coefficient between predicted and observed values. Forecasts again were made out of sample by removing each point for which a forecast was being made from the model fitting procedure. As these larvae average nineteen days in age, the expectation is that there should be a peak in predictability at a corresponding nineteen-day lag in lunar phase. Also, because the analysis of spawning output in the previous section indicated that egg release from the nest is well-described by a linear model, it seems likely that for this variable, linear models should perform about as well as nonlinear ones. Both of these expectations are met; there is indeed a peak in predictability at the appropriate time lag, and predictability remains essentially constant as the model is made more nonlinear.

Because of the deterministic nature of the lunar cycle, there is a second peak in predictability, centered at approximately five days and it appears that predictability begins to improve again as the lag approaches thirty days. If peaks in larval abundance are associated with a full moon nineteen days prior, it stands to reason that they will also be associated with a new moon five days and thirty-three days prior. To test that this is actually what these results are capturing, it is therefore necessary to extract the function relating larval supply to lunar phase at each of these two lags. In this case, where only one variable is being considered by each model, such an extraction is straightforward. However, the technique which will be used for this will be introduced in the most general terms, to allow an easy extension to the higher dimensional cases that will follow.

To deduce the functional relationship between larval supply and the physical parameters of interest, a space is first constructed whose axes are given by the physical variables of interest, normalized by their standard deviations (to normalize the ranges of each variable). Each point in the larval supply time series is then mapped to its corresponding physical coordinates in this state space. A neighborhood average (of N neighbors) is then constructed for each point in the state space; that is, for every point, that point's N nearest neighbors in state space are identified, and the average value of each physical variable in the neighborhood calculated. The average value of larval supply for the points in the neighborhood is also determined, and, for each physical variable, a plot is made of the neighborhood averages of that variable against the neighborhood averages of larval supply (after multiplying by the appropriate standard deviations). The advantage of this technique is that, when more than one physical variable is considered, it becomes possible to smooth the data while preserving the non-additive sense of the interactions between the physical variables.

Results for the five- and nineteen-day lags of lunar phase are given in Figure 14.9. As expected, maximum larval abundance is associated with a new moon (zero percent illumination) for the five-day lag, and the full moon (100 percent illumination) for the nineteen-day lag. The next question is the degree to which modeling larval supply from lunar phase captures the dynamics of the real data. The nineteen-day-lag model gives a correlation coefficient of 0.3 (Table 14.1). Figure 14.10 gives another sense of model performance by plotting the actual predictions alongside the real values. In all four panels, the real data are plotted on the ordinate in the positive direction, and the forecasts (multiplied by minus one) in the negative. The upper-left panel gives the predictions from the linear model built on the five-day lag; the upper-right panel gives the same variable with a nonlinear model. The bottom panels are from corresponding models (linear and nonlinear) constructed from the nineteen-day lag. In all four cases, the deterministic component of larval supply gives a sense of the timing of the peaks, but as expected, there is no information in the phase of the moon about the fate of the larvae after they are born. It is thus not possible to




FIGURE 14.9. Response of larvae to lunar phase (measured by percent illumination). For the 19-day lag, larvae are positively associated with the full moon (percent illumination = 100); for the 5-day lag, the new moon (percent illumination = 0).

distinguish strong peaks from weak ones or from no peak at all, with just this variable.

14.3.2 Cross-Shelf Wind

For this, additional variables corresponding to other events in larval life need to be incorporated into the analysis. The next process that will be considered was motivated by a prior analysis of this system (Milicich 1994[25]), which suggested that the transport environment experienced by mature larvae returning to the reef might be of importance to the magnitude of larval supply.⁵ As no current meter data are available from this system, a three-day running average of recent cross-shelf wind speed (the onshore/offshore component of the wind) was selected as the physical variable most representative of surface current speed and direction. Figure 14.11 gives the results of adding this variable and again constructing a series of forecasting algorithms, ranging from global to local fits of the data. The top panel demonstrates that, considered linearly, there is no evidence that this variable is important, but when the models are made nonlinear, predictability begins to increase. This is as expected; these two effects occur at different times in larval life (the beginning and the end), and thus are coupled multiplicatively in their total impact. The middle panels give the predictions

⁵Indeed, exploring the transport of particles to and from reefs is a usual starting point in attempting to understand larval supply for vertebrates and invertebrates alike.

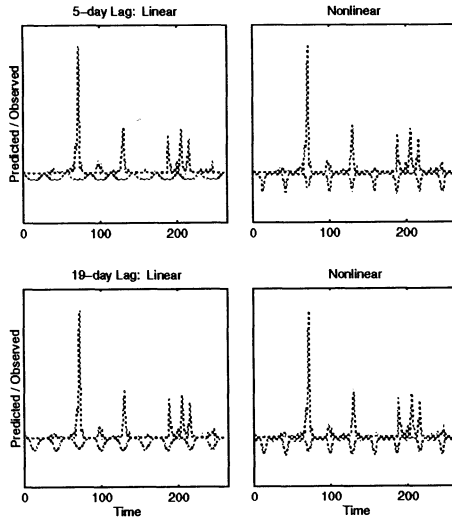


FIGURE 14.10. Predicting larval supply from the lunar phase. Upper panels: 5-day lag; lower panels: 19-day lag. Left panels: linear models; right panels: nonlinear models. Both linear and nonlinear models give a sense of the timing of larval supply peaks, but not the magnitude.

from linear and nonlinear models based on these two variables. Only in the nonlinear case does it begin to become possible to distinguish strong peaks from weaker ones. The bottom panels give the results of extracting, in the fashion detailed earlier, the functions relating larval supply to these two variables. As before, larval supply peaks are associated with a full moon nineteen days prior; they may now also be seen to be associated with weak, onshore winds (positive values of cross-shelf wind speed are onshore; negative values offshore). This is contrary to what would be expected if mature larvae act as passive particles (in which case stronger onshore winds, and associated greater onshore transport of surface waters, would be expected to increase larval supply), but is consistent with recent findings concerning the abilities of mature larval fishes to swim, aggregate, and generally promote their own orientation (e.g., Stobutzki and Bellwood 1994[35], Leis et al. 1996[21], Leis and Carson-Ewart 1997[20]).

14.3.3 Average Daily Wind

The final process that will be considered is the potential importance of wind stress on the first-feeding success of young larval fishes. The suggestion that such critical periods are important to subsequent recruitment has a long history in the fisheries literature (e.g., Hjort 1914[13]), and the




FIGURE 14.11. Modeling pomacentrid larval supply from lunar phase and recent cross-shelf (onshore) wind speed. Top panel: predictability improves as the model is made nonlinear (theta positive). Middle panels: functional responses. Larval supply is as before positively associated with a full moon 19 days prior, and also with weak, onshore winds (positive values = onshore; negative values = offshore). Bottom panels: model predictions, for linear (left) and nonlinear (right) algorithms. Adding in cross-shelf wind speed begins to give a sense of the magnitude of the peaks.




FIGURE 14.12. Investigating the importance of wind events early in larval life. There is a peak in predictability when wind stress lagged 16 days is incorporated into nonlinear models. It appears that first-feeding success is indeed important to subsequent recruitment in this system.

relationship between turbulence and feeding success has been the subject of much study over the past ten years, both theoretical and experimental (e.g., Rothschild and Osborne 1988[32], Davies et al. 1991[6], MacKenzie et al. 1994[22], Sundby and Fossum 1990[38], Cury and Roy 1990[4], Wroblewski et al. 1989[42]). Detecting such lagged effects from time series taken at a fixed point in space is, however, problematic, as there is no guarantee that the local history of the physical environment will be the same as that experienced by the larvae themselves, given their planktonic lifestyle. However, it seems at least plausible that such events could be demonstrated at Lizard Island, given the suggestion that retention mechanisms may work to keep larvae local to their native reef (Leis 1986[19]). The difficulty is that, because not all larvae are precisely the same age and because both the time required for yolk-sac resorption and the time to inevitable mortality in the absence of successful feeding are unknown, it is not obvious exactly which lag should be chosen. As a sanity check, and to be conservative, a wide range of lags (from one to thirty days) was therefore evaluated, each over a range of degrees of model nonlinearity. Clearly, if the modeling results indicate that the best lag corresponds to a time (for example) before the fish were even born, something is amiss. Results are given in Fig. 14.12. Considered linearly, there is little to distinguish among the lag choices. However, as the models are made nonlinear, a single peak in predictability beyond that observed for the bivariate nonlinear model emerges. Strikingly, this peak coincides very well with the age of the larvae. The pomacentrid larval duration averaged nineteen days; the peak in predictability is centered at sixteen days, three days after release from the nest.

The approach to variable selection seems to be yielding a reasonable result, but do the functional relationships contained in the model continue to




FIGURE 14.13. Functions relating larval supply to the physical environment extracted from the trivariate physical model. Top two panels are essentially as before; now (bottom panel) larval supply can be seen to be associated with intermediate wind speeds (centered on 6 meters per second) 16 days prior, when larvae are very young and presumably first-feeding.

make biological sense? Figure 14.13 gives these, extracted as before. Larval supply peaks remain associated with the full moon nineteen days prior, and recent weak, onshore winds. Additionally, it now appears that there is an approximately dome-shaped response to the sixteen-day lag wind stress, centered on about 6 meters per second average daily wind speed. This is in good agreement with the results of the previously mentioned theoretical and experimental studies, and is the first demonstration from field observations that such events are important enough that they can be used to help predict subsequent recruitment. Figure 14.14 gives model forecasts, plotted with the real data. Clearly there is much variability as yet unexplained, but it is possible to at least begin to predict the magnitude of the peaks in larval supply.




FIGURE 14.14. Output of the trivariate nonlinear physical model: the best predictions made here for larval supply in this system.

Variables	Lag	Model type	ρ	Err	RMS
		Linear (out-of-sample)	0.29	100.5	241.4
		Nonlinear S-map ($\theta = 0.6$)	0.30	103.1	241.0
Cross-shelf wind	1	Linear (out-of-sample)	0.34	109.1	237.4
		Nonlinear S-map ($\theta = 0.6$)	0.45	94.1	221.0
Cross-shelf wind	1	Linear (out-of-sample)	0.38	114.8	233.9
Daily wind	16	Nonlinear S-map ($\theta = 0.8$)	0.55	88.2	211.5

TABLE 14.1. Summary of model statistics. Lag in days; Rho gives correlation coefficient between predicted and observed values, Err is the average error of the predictions, and RMS is the root mean squared error from predictions.




FIGURE 14.15. Yearly resupply of larval pomacentrids simulated for ten years by running the nonlinear trivariate model on surrogate physical inputs. The 4-fold variability in replenishment during this time arises from stationary processes; there is no interannual variability in the physical forcing. It is the specific realizations of the stochastic environment that are important.

14.4 Conclusions

The goal of this chapter has been to explore in detail a case study where the application of nonlinear techniques leads to real gains in understanding the variability in a natural biological system. Although the forecast skill of the best model constructed in this example is still relatively weak, it is a good deal better than that which can be achieved with linear analysis. More importantly, the model identifies processes which have long been suspected of being important in driving change in fish populations, but whose operation in nature has historically been very difficult to demonstrate.

This work is part of an emerging trend in ecology to explore the interactions between noise and nonlinearity in population dynamics (Leirs et al. 1997[18], Grenfell et al. 1998[11], Myers et al. 1998[27], Dixon et al. 1999[7], Zimmer 1999[43]), and the results presented here and elsewhere raise a host of new questions for ecologists to consider. One issue in particular that may need further consideration is how to think about interannual variability. In situations such as that presented here, it is entirely possible for the interplay between noise and nonlinearity to cause a population whose variability is a product of stationary processes to itself appear non-stationary on interannual time scales. For example, Figure 14.15 gives the result of running the nonlinear trivariate physical model of larval supply for multiple years on surrogate physical inputs. Larval supply is summed over each year and plotted over time for ten years. Population replenishment varies by a factor of more than 3 from year to year, and yet this model is driven entirely by stationary physical processes. This is not to say that interannual variability in physical forcing is unimportant; the opposite is

very often the case. Nevertheless, in this example, the tempting conclusion that interannual variability in the physical environment was underlying the observed variability in population resupply would have been premature.

Some recent analyses of time series of natural populations have focused on understanding the density-dependent component of population change in the presence of environmental noise by exploring whether the internal dynamics themselves were chaotic or stable (Ellner and Turchin 1995[9]). The impact of external forcing was dealt with by subjecting nearby initial biological conditions to identical sequences of random shocks. In other words, the question that was asked was whether similar initial population states would diverge rapidly from each other if environmental variability was realized the same way each time. In the example given here, however, the biological variability results directly from different *realizations* of the stochastic forcing processes. If the goal is to predict variability from measurable initial conditions, the suggestion is that not only will the relevant physical variables will have to be identified (along with which are the favorable and unfavorable conditions), but also the sensitivity of natural populations to the temporal sequences of physical events will need to be understood. And, of course, in many cases, such sensitivity may be a function of population size itself. Clearly, a full understanding of the interplay between noise and nonlinearity in natural systems is a difficult challenge. The potential rewards are great, however, and this is likely to be an area where the analytical approaches that are the focus of this book will make strong contributions to the understanding of ecological systems in the future.

References

- [1] H. G. Andrewartha and L. C. Birch. *The Distribution and Abundance of Animals*. University of Chicago Press, 1954.
- [2] N. Caputi. Factors affecting the time series bias in stock-recruitment relationships and the interaction between time series and measurement error bias. *Canadian Journal of Fisheries and Aquatic Science*, 45, 1988.
- [3] R. M. Cassie. Microdistribution of plankton. *Oceanography and Marine Biology Annual Review*, 1:223–252, 1963.
- [4] P. Cury and C. Roy. Optimal environmental window and pelagic fish recruitment success in upwelling areas. *Canadian Journal of Fisheries and Aquatic Science*, 46, 1989.
- [5] D. Cushing. *Climate and Fisheries*. Academic Press, London, 1982.
- [6] C. Davies, G. Flierl, P. Wiebe, and P. Franks. Micropatchiness, turbulence, and recruitment in plankton. *Journal of Marine Research*, 49, 1991.
- [7] P.A. Dixon, M.J. Milicich, and G. Sugihara. Episodic fluctuations in larval supply. *Science*, 283:1528–1531, 1999.
- [8] P. J. Doherty and D.McB. Williams. The preplenishment of coral reef fishes. *Oceanography and Marine Biology*, 26:487–551, 1988.

- [9] S. Ellner and P. Turchin. Chaos in a noisy world: New methods and evidence from time-series analysis. *The American Naturalist*, 145:343–375, 1995.
- [10] J. P. Finerty. *The Population Ecology of Cycles in Small Mammals*. Yale University Press, New Haven, 1980.
- [11] B.T. Grenfell, K. Wilson, B.F. Finkenstadt, T.N. Coulson, S. Murray, S.D. Albon, J.M. Pemberton, T.H. Clutton-Brock, and M.J. Crawley. Noise and determinism in synchronized sheep dynamics. *Nature*, 394:674–677, 1998.
- [12] R. Hilborn and C. Walters. *Quantitative Fisheries Stock Assessment: Choice, Dynamics, and Uncertainty*. Chapman and Hall, London, 1992.
- [13] J. Hjort. Fluctuations in the great fisheries of northern europe viewed in the light of biological research. *Rapp. P.-V. Reun. Cons. Perm. Int. Explor. Mer.*, 20:1–228, 1914.
- [14] H. Kantz and T. Schreiber. *Nonlinear Time Series Analysis*. Cambridge University Press, New York, 1997.
- [15] Z. Kopal. *Physics and Astronomy of the Moon*. Academic Press, New York and London, 1971.
- [16] D.L. Lack. *The Natural Regulation of Animal Numbers*. Oxford Press, 1954.
- [17] R. Lasker. *Marine Fish Larvae*. University of Washington Press, Seattle, 1981.
- [18] H. Leirs, N.C. Stenseth, J.D. Nichols, J.E. Hines, R. Verhagen, and W. Verheyen. Stochastic seasonality and nonlinear density-dependent factors regulate population size in an african rodent. *Nature*, 389:176–180, 1997.
- [19] J.M. Leis. Vertical and horizontal distribution of fish larvae near coral reefs at lizard island, great barrier reef. *Marine Biology*, 90:505–516, 1986.
- [20] J.M. Leis and B.M. Carson-Ewart. In situ swimming speeds of the late pelagic larvae of some indo-pacific coral-reef fishes. *Marine Ecology Progress Series*, 159:165–174, 1997.
- [21] J.M. Leis, H.P.A. Sweatman, and S.E. Reader. What the pelagic stages of coral reef fishes are doing out in blue water: Daytime field observations of larval behavioural capabilities. *Marine and Freshwater Research*, 47:401–411, 1996.
- [22] B. MacKenzie, T. Miller, S. Cyr, and W. Leggett. Evidence for a dome-shaped relationship between turbulence and larval fish ingestion rates. *Limnology and Oceanography*, 39:1790–1799, 1994.
- [23] R.M. May. Biological populations with nonoverlapping generations: Stable points, stable cycles, and chaos. *Science*, 186:645–647, 1974.
- [24] M.G. Meekan, M.J. Milicich, and P.J. Doherty. Larval production drives temporal patterns of larval supply and recruitment of a coral reef damselfish. *Marine Ecology Progress Series*, 93:217–225, 1993.
- [25] M.J. Milicich. Dynamic coupling of reef fish replenishment and oceanographic processes. *Marine Ecology Progress Series*, 110:135–144, 1994.
- [26] M.J. Milicich, M.G. Meekan, and P.J. Doherty. Larval supply: A good predictor of recruitment of three species of reef fish (pomacentridae). *Marine Ecology Progress Series*, 86:153–166, 1992.
- [27] R.A. Myers, M. Gordon, J.M. Bridson, and M.J. Bradford. Simple dynamics underlie sockeye salmon (*oncorhynchus nerka*) cycles. *Canadian Journal of Fisheries and Aquatic Science*, 55:2355–2364, 1998.
- [28] A.J. Nicholson. The balance of animal populations. *Journal of Animal Ecology*, 2:132–178, 1933.

- [29] R.M. Peterman, M.J. Bradford, N.C.H. Lo, and R.D. Methot. Contribution of early life history stages to interannual variability in recruitment of northern anchovy (*Engraulis mordax*). *Canadian Journal of Fisheries and Aquatic Sciences*, 45:8–16, 1988.
- [30] W.J. Richards and K.C. Lindeman. Recruitment dynamics of reef fishes: Planktonic processes, settlement and demersal ecologies, and fishery analysis. *Bulletin of Marine Science*, 41:392–410, 1987.
- [31] B.J. Rothschild. *Dynamics of Marine Fish Populations*. Harvard University Press, Cambridge, 1986.
- [32] B.J. Rothschild and T.R. Osborne. Small-scale turbulence and plankton contact rates. *Journal of Plankton Research*, 10:465–474, 1988.
- [33] W.M. Schaffer. Order and chaos in ecological systems. *Ecology*, 66:93–106, 1985.
- [34] W.M. Schaffer and M. Kot. Chaos in ecological systems: The coals that newcastle forgot. *Trends in Ecology and Evolution*, 1:58–63, 1986.
- [35] I.C. Stobutzki and D.R. Bellwood. An analysis of the sustained swimming abilities of pre- and post-settlement coral reef fishes. *Journal of Experimental Marine Biology and Ecology*, 175:275–286, 1994.
- [36] G. Sugihara. Nonlinear forecasting for the classification of natural time series. *Philosophical Transactions of the Royal Society of London A*, 348:477–495, 1994.
- [37] G. Sugihara, M. Casdagli, E. Habjan, D. Hess, P. Dixon, and G. Holland. Residual delay maps unveil global patterns of atmospheric nonlinearity and produce improved local forecasts. *Proceedings of the National Academy of Science*, In press, 1999.
- [38] S. Sundby and P. Fossum. Feeding conditions of arcto-norwegian cod larvae compared with the Rothschild-Osborne theory on small-scale turbulence and plankton contact rates. *Journal of Plankton Research*, 12:1153–1162, 1990.
- [39] J. Theiler, B. Galdrikian, A. Longtin, S. Eubank, and J.D. Farmer. Using surrogate data to detect nonlinearity in time series. In M. Casdagli and S. Eubank, editors, *Nonlinear Modeling and Forecasting*, Santa Fe Institute Studies in the Sciences of Complexity, pages 163–188, 1992.
- [40] G.M. Wellington and B.C. Victor. Planktonic larval duration of one hundred species of Pacific and Atlantic damselfishes. *Marine Biology*, 101:557–567, 1989.
- [41] D. McB. Williams and S. English. Distribution of fish larvae around a coral reef: Direct detection of a meso-scale, multi-specific patch? *Continental Shelf Research*, 112:923–937, 1992.
- [42] J.S. Wroblewski, J.G. Richman, and G.L. Mellor. Optimal wind conditions for the survival of larval northern anchovy, *Engraulis mordax* - a modeling investigation. *Fishery Bulletin*, 87:387–398, 1989.
- [43] C. Zimmer. Life after chaos. *Science*, 284:83–86, 1999.

Chapter 15

Cluster-Weighted Modeling: Probabilistic Time Series Prediction, Characterization, and Synthesis

Bernd Schoner¹
Neil Gershenfeld

ABSTRACT Cluster-weighted modeling, a mixture density estimator around local models, is presented as a framework for the analysis, prediction and characterization of non-linear time series. First architecture, model estimation and characterization formalisms are introduced. The characterization tools include estimator uncertainty, predictor uncertainty, and the correlation dimension of the data set. In the second part of this chapter the framework is extended to synthesize audio signals and is applied to model a violin in a data-driven input-output approach.

15.1 Introduction

The list of time series worth forecasting is about as long as the first unsuccessful attempts to do so. It would be most helpful to know beforehand when a heart is about to stop beating, what the weather will be like tomorrow and when the stock market is going to crash. Unfortunately, these examples share the one property that they have nothing in common and that they don't fit into any familiar categories of system dynamics theory. Not only are they non-linear, non-Gaussian and non-stationary, they are essentially non-everything.

Linear systems theory has yielded a multitude of results that are widely applied in practically all engineering and scientific disciplines. The majority of signal processing, system engineering, control and characterization techniques rely on linear assumptions and use a theory that has matured in decades of research and implementations. However, the limitations of linear theory are clear: Non-linear behavior of any kind cannot be handled.

The reconstruction (embedding) theorem on the other hand provides

¹ Author for correspondence.

the theoretical means to handle highly non-linear behavior of arbitrary physical systems with hidden dynamics [22]. It shows that the system's state space can be mapped into a diffeomorphic space, constructed from any observable of the system, and that we can characterize the data with respect to dimensionality and dynamic behavior in the reconstructed space. The reconstruction theorem also detects low dimensional structure in a high dimensional data space, which lets us work in the space described by the effective degrees of freedom of a system, for example, a violin, rather than its countless mechanical degrees of freedom.

Unfortunately it turns out to be rather difficult to use a reconstructed state space to predict the output of a complex system. While low-dimensional systems are tractable (Fig.15.1), models become easily unstable given a complicated state space or an arbitrary prediction horizon. Driven systems should be easier to handle than autonomous systems. However, the model dimensionality of a driven system is significantly bigger, since input and output observables need to be considered at the same time[4]. The presence of noise in practically any real world system further complicates the embedding task. Due to these problems we end up with a fairly small number of examples where embedding, despite its theoretical promise, has been applied successfully to predict a signal.

In between linear and highly non-linear systems there is a large class of systems that are not easily classified as one or the other but combine characteristics from both worlds. The bow string interaction of a violin, for example, is strongly non-linear, since it transforms the slow actions of the player into a fast audio signal. At the same time the effect of the violin body is most efficiently described by a linear filter since there is only little non-linear effects in the bridge and body dynamics² [14]. Hence a violin combines linear and non-linear processing.

This chapter introduces cluster-weighted modeling (CWM) as a modeling tool that allows one to characterize and predict systems of arbitrary dynamic character. The framework is based on density estimation around Gaussian kernels which contain simple local models describing the system dynamics of a data subspace. In the extreme case where only one kernel is used the framework collapses to a simple model that is linear in the coefficients. In the opposite extreme it allows one to embed and forecast data that may be non-Gaussian, discontinuous, high-dimensional and chaotic. In between CWM covers a multitude of models, each of which is characterized by a different local model and state representation. We create globally non-linear models with transparent local structures through the embedding of past practice and mature techniques in the general non-linear framework.

²The exception from this is the famous Wolf tone, a tone that periodically collapses despite constant bowing. The phenomenon is particularly strong on the cello and is caused by a non-linear coupling between a string and a body mode [6].

The limitations of artificial neural networks (ANNs) have become apparent almost as quickly as their modeling power: Networks take long to converge, coefficients are only meaningful in the context of the entire model and failure and success of an architecture are unpredictable beforehand. More recently a new family of networks, which interpret data probabilistically, and are often represented in graphical networks, has been developed [3, 9, 11]. As a meta-class of models, graphical models are conceptually unbounded. They unify existing network architectures, for example, classical ANNs in a single theory [15] provide new insights and extensions to conventional networks and open up new application domains. Graphical models are also referred to as independence networks, since the graphical representation really describes dependence and independence among random variables. They are called Bayesian belief networks since dependencies between variables are expressed in terms of conditional probability functions that have implicit or explicit prior beliefs built into them. They are furthermore named influence diagrams since causal dependences between variables are clearly illustrated. “Influence” is meant probabilistically, which contains deterministic causality as a special case. Unfortunately graphical models lack a systematic search algorithm that maps a given problem into a network architecture. Instead, before the network parameters can be trained on new data, the architecture needs to be redesigned node by node from scratch.

Cluster-weighted modeling is a special case of a probabilistic model that gives up some of the generality of graphical models in favor of ease of use, a minimal number of hyper-parameters and a fast parameter search. It has been designed as an architecture that is as general as reasonably possible, but as specific to a particular application as necessary. We present a tool that allows us to do statistical time series analysis from a physicist’s perspective and at the same time allows us to solve complicated engineering problems, for example, the design of a digital musical instruments. As opposed to ANNs it provides transparent local structures and meaningful parameters, it allows one to identify and analyze data subspaces and converges quickly.

The first part of this chapter provides the basic architecture, estimation, and characterization tools of CWM. The second part is concerned with the problem of building a data-driven input-output model of a violin. The violin is a complex driven device that in its socio-cultural, artistic and physical subtlety is hardly matched by any other human artifact. At the same time the violin provides a very clear error metric in that the model is just as good as it sounds. From a non-linear dynamics and statistics viewpoint the violin is a paradigmatic object, since it shows non-linear and linear, stochastic and deterministic behavior at the same time.

15.2 Cluster-Weighted Modeling

15.2.1 Architecture

Cluster-weighted modeling (CWM) is an input-output inference framework based on probability density estimation of a joint set of input features and output target data. It is similar to mixture-of-experts type architectures [10] and can be interpreted as a flexible and transparent technique to approximate an arbitrary function. Unlike conventional Kernel-based techniques, CWM requires only one hyper-parameter to be fixed beforehand and provides data parameters such as the length scale (bandwidth) of the local approximation as an output rather than an input of the algorithm [5].

We start with a set of discrete or real-valued input features \mathbf{x} which may be measured features or components in a time-lagged embedding space, and an discrete or real valued output target vector \mathbf{y} . Given the joint input-output set $\{\mathbf{y}_n, \mathbf{x}_n\}_{n=1}^N$, the most general model infers the joint density $p(\mathbf{y}, \mathbf{x})$ of the data set, from which conditional quantities such as the expected \mathbf{y} given \mathbf{x} , $\langle \mathbf{y} | \mathbf{x} \rangle$, and the expected covariance of \mathbf{y} given \mathbf{x} , $\langle \mathbf{P}_y | \mathbf{x} \rangle$ can be derived.

We expand this joint density in clusters labeled c_m , each of which contains an input domain of influence, a local model, and an output distribution. In a first step the joint density is separated into an unconditioned cluster probability and a conditional probability of a data given a cluster, which is then further expanded into an input domain of influence and an output distribution,

$$\begin{aligned}
 p(\mathbf{y}, \mathbf{x}) &= \sum_{m=1}^M p(\mathbf{y}, \mathbf{x}, c_m) & (15.1) \\
 &= \sum_{m=1}^M p(\mathbf{y}, \mathbf{x} | c_m) p(c_m) \\
 &= \sum_{m=1}^M p(\mathbf{y} | \mathbf{x}, c_m) p(\mathbf{x} | c_m) p(c_m).
 \end{aligned}$$

Many problems require a distinction between slowly varying state variables describing the global boundary conditions and state of the system and quickly-varying variables describing the fast dynamics of the system. If this is the case we decompose \mathbf{x} into \mathbf{x}_s and \mathbf{x}_f and obtain for the density

$$p(\mathbf{y}, \mathbf{x}) = \sum_{m=1}^M p(\mathbf{y} | \mathbf{x}_f, c_m) p(\mathbf{x}_s | c_m) p(c_m), \quad (15.2)$$

where \mathbf{x}_s and \mathbf{x}_f may be identical, overlapping in some dimensions or completely distinct.

The input distribution is taken to be a Gaussian distribution,

$$p(\mathbf{x}|c_m) = \frac{|\mathbf{P}_m^{-1}|^{1/2}}{(2\pi)^{D/2}} e^{-(\mathbf{x}-\mu_m)^T \cdot \mathbf{P}_m^{-1} \cdot (\mathbf{x}-\mu_m)/2}, \quad (15.3)$$

where \mathbf{P}_m is the cluster-weighted covariance matrix in the feature space. It can be reduced to variances in each dimension, when computational complexity is an issue.

The output distribution is taken to be

$$p(\mathbf{y}|\mathbf{x}, c_m) = \frac{|\mathbf{P}_{m,y}^{-1}|^{1/2}}{(2\pi)^{D_y/2}} e^{-(\mathbf{y}-\mathbf{f}(\mathbf{x}, \beta_m))^T \cdot \mathbf{P}_{m,y}^{-1} \cdot (\mathbf{y}-\mathbf{f}(\mathbf{x}, \beta_m))/2}, \quad (15.4)$$

where the mean value of the output Gaussian is replaced by the function $\mathbf{f}(\mathbf{x}, \beta_m)$ with unknown parameters β_m . Again the off diagonal terms in the output covariance matrices $\mathbf{P}_{m,y}$ can be neglected if needed.

To understand this form, consider the conditional forecast of the expected \mathbf{y} given \mathbf{x} :

$$\begin{aligned} \langle \mathbf{y}|\mathbf{x} \rangle &= \int \mathbf{y} p(\mathbf{y}|\mathbf{x}) d\mathbf{y} \\ &= \int \mathbf{y} \frac{p(\mathbf{y}, \mathbf{x})}{p(\mathbf{x})} d\mathbf{y} \\ &= \frac{\sum_{m=1}^M \int \mathbf{y} p(\mathbf{y}|\mathbf{x}, c_m) d\mathbf{y} p(\mathbf{x}|c_m) p(c_m)}{\sum_{m=1}^M p(\mathbf{x}|c_m) p(c_m)} \\ &= \frac{\sum_{m=1}^M \mathbf{f}(\mathbf{x}, \beta_m) p(\mathbf{x}|c_m) p(c_m)}{\sum_{m=1}^M p(\mathbf{x}|c_m) p(c_m)}. \end{aligned} \quad (15.5)$$

We observe that the predicted \mathbf{y} is a superposition of all the local functionals, where the weight of each contribution depends on the posterior probability that an input point was generated by a particular cluster. The denominator ensures that the sum of the weights of all contributions equals unity.

Likewise we compute the conditional error in terms of the expected covariance of \mathbf{y} given \mathbf{x}

$$\begin{aligned} \langle \mathbf{P}_y|\mathbf{x} \rangle &= \int (\mathbf{y} - \langle \mathbf{y}|\mathbf{x} \rangle)(\mathbf{y} - \langle \mathbf{y}|\mathbf{x} \rangle)^T p(\mathbf{y}|\mathbf{x}) d\mathbf{y} \\ &= \int (\mathbf{y}\mathbf{y}^T - \langle \mathbf{y}|\mathbf{x} \rangle \langle \mathbf{y}|\mathbf{x} \rangle^T) p(\mathbf{y}|\mathbf{x}) d\mathbf{y} \\ &= \frac{\sum_{m=1}^M \int \mathbf{y}\mathbf{y}^T p(\mathbf{y}|\mathbf{x}, c_m) d\mathbf{y} p(\mathbf{x}|c_m) p(c_m)}{\sum_{m=1}^M p(\mathbf{x}|c_m) p(c_m)} - \langle \mathbf{y}|\mathbf{x} \rangle \langle \mathbf{y}|\mathbf{x} \rangle^T \\ &= \frac{\sum_{m=1}^M [\mathbf{P}_{m,y} + \mathbf{f}(\mathbf{x}, \beta_m) \mathbf{f}(\mathbf{x}, \beta_m)^T] p(\mathbf{x}|c_m) p(c_m)}{\sum_{m=1}^M p(\mathbf{x}|c_m) p(c_m)} - \langle \mathbf{y}|\mathbf{x} \rangle \langle \mathbf{y}|\mathbf{x} \rangle^T \end{aligned} \quad (15.6)$$

which equals the expected variance if only a single output dimension is considered,

$$\langle \sigma_y^2 | \mathbf{x} \rangle = \frac{\sum_{m=1}^M [\sigma_{m,y}^2 + f(\mathbf{x}, \beta_m)^2] p(\mathbf{x}|c_m) p(c_m)}{\sum_{m=1}^M p(\mathbf{x}|c_m) p(c_m)} - \langle y | \mathbf{x} \rangle^2 .$$

There are two parameters to be determined beforehand: the number of clusters M and the form of the local models f which together control the model resources and hence under-fitting versus over-fitting. We trade off the complexity of the local models against the complexity of the global architecture, which is nicely illustrated in the case of a local polynomial expansion (Equ.15.8): If we use locally constant models together with a large number of clusters, the predictive power is determined by the number of Gaussian kernels. If, alternatively, we use a high-order polynomial model and a single kernel, the model reduces to a global polynomial model.

The choice of local models depends on the application. In general f expresses prior beliefs about the nature of the data or insights in the mechanics of a system and thus functions as a regularizer of the model. Machine learning architectures and estimation algorithms typically depend on global regularizers that handle prior beliefs about what a good model is. This is problematic since global statements may not apply locally. For example, the maximum entropy principle is good at handling discontinuities, but it has no notion of local smoothness, whereas integrated curvature is good for enforcing local smoothness but rounds out discontinuities. In our approach the model is constrained only by the local architecture which may enforce local smoothness but at the same time allows for discontinuities where needed.

15.2.2 Model Estimation

Non-linear function fitting uses models with linear coefficients β_m and non-linear basis functions $f(\mathbf{x})$,

$$\mathbf{y}(\mathbf{x}) = \sum_{m=1}^M \beta_m \mathbf{f}_m(\mathbf{x}) , \quad (15.8)$$

for example, a polynomial expansion, or models that have the coefficients inside the nonlinearities,

$$\mathbf{y}(\mathbf{x}) = \sum_{m=1}^M \mathbf{f}_m(\mathbf{x}, \beta_m) , \quad (15.9)$$

for example, a neural network. In the case of a generalized linear model (Equ.15.8) only a single matrix pseudo-inverse is needed to find the set of

coefficients yielding the minimum mean-square error. However, the number of coefficients in Equation 15.8 is exponential in the dimension of \mathbf{x} . A model with non-linear coefficients (Equ.15.9) has more expressive power, which can reduce the number of coefficients needed for a given approximation error to linear in the dimension of \mathbf{x} [2]. Yet, the non-linear parameters of Equ.15.9 require an iterative search [8].

CWM uses simple local models, which satisfy (15.8), to create globally powerful models as described by (15.9) and hence combines the efficient estimation of the former with the benefits of the latter models. We fit the local model parameters by a matrix inversion of the local covariance matrix and find the remaining cluster parameters in charge of the global weighting, using a variant of the expectation-maximization (EM) algorithm [7]. EM is an iterative search that maximizes the model likelihood given a data set and initial conditions [16, 1]. We pick a set of starting values for the cluster parameters and enter the iterations with the expectation step.

In the **E-step** we assume the current cluster parameters are correct and evaluate the posterior probabilities that relate each cluster to each data point. These posteriors can be interpreted as the probability that a particular data was generated by a particular cluster or as the normalized responsibility of a cluster for a point:

$$\begin{aligned} p(c_m|\mathbf{y}, \mathbf{x}) &= \frac{p(\mathbf{y}, \mathbf{x}|c_m) p(c_m)}{p(\mathbf{y}, \mathbf{x})} \\ &= \frac{p(\mathbf{y}, \mathbf{x}|c_m) p(c_m)}{\sum_{l=1}^M p(\mathbf{y}, \mathbf{x}|c_l) p(c_l)}, \end{aligned} \quad (15.10)$$

where the sum over clusters in the denominator causes clusters to interact, fight over points and specialize in data they best explain.

In the **M-step** we assume the current data distribution is correct and find the cluster parameters that maximize the likelihood of the data. The new estimate for the unconditioned cluster probabilities is

$$\begin{aligned} p(c_m) &= \int p(c_m|\mathbf{y}, \mathbf{x}) p(\mathbf{y}, \mathbf{x}) d\mathbf{y} d\mathbf{x} \\ &\approx \frac{1}{N} \sum_{n=1}^N p(c_m|\mathbf{y}_n, \mathbf{x}_n), \end{aligned} \quad (15.11)$$

Here the idea is that an integral over a density can be approximated by an average over variables drawn from the density.

Next we compute the expected input mean of each cluster which is the

estimate of the new cluster means:

$$\begin{aligned}
 \mu_m &= \int \mathbf{x} p(\mathbf{x}|c_m) d\mathbf{x} & (15.12) \\
 &= \int \mathbf{x} p(\mathbf{y}, \mathbf{x}|c_m) d\mathbf{y} d\mathbf{x} \\
 &= \int \mathbf{x} \frac{p(c_m|\mathbf{y}, \mathbf{x})}{p(c_m)} p(\mathbf{y}, \mathbf{x}) d\mathbf{y} d\mathbf{x} \\
 &\approx \frac{1}{N p(c_m)} \sum_{n=1}^N \mathbf{x}_n p(c_m|\mathbf{y}_n, \mathbf{x}_n) \\
 &= \frac{\sum_{n=1}^N \mathbf{x}_n p(c_m|\mathbf{y}_n, \mathbf{x}_n)}{\sum_{n=1}^N p(c_m|\mathbf{y}_n, \mathbf{x}_n)}. & (15.13)
 \end{aligned}$$

The apparently formal introduction of \mathbf{y} into the density as a variable to be integrated over has the important result that cluster parameters are found with respect to the joint input-output space. Clusters get pulled based on both where there is data to be explained and how well their model explains the data. In a similar way we can define a cluster-weighted expectation of any function $\Theta(\mathbf{x})$,

$$\begin{aligned}
 \langle \theta(\mathbf{x}) \rangle_m &\equiv \int \theta(\mathbf{x}) p(\mathbf{x}|c_m) d\mathbf{x} & (15.14) \\
 &\approx \frac{1}{N} \sum_{n=1}^N \theta(\mathbf{x}_n) \frac{p(c_m|\mathbf{y}_n, \mathbf{x}_n)}{p(c_m)} \\
 &= \frac{\sum_{n=1}^N \theta(\mathbf{x}_n) p(c_m|\mathbf{y}_n, \mathbf{x}_n)}{\sum_{n=1}^N p(c_m|\mathbf{y}_n, \mathbf{x}_n)},
 \end{aligned}$$

which lets us update the cluster weighted covariance matrices:

$$[\mathbf{P}_m]_{ij} = \langle (x_i - \mu_i)(x_j - \mu_j) \rangle_m. \quad (15.15)$$

It also lets us compute the matrices needed to update the local models. The model parameters are found by taking the derivative of the *log* of the total likelihood function with respect to the parameters:

$$0 = \frac{\partial}{\partial \beta} \log \prod_{n=1}^N p(\mathbf{y}_n, \mathbf{x}_n). \quad (15.16)$$

Considering a single output dimension y and a single coefficient β_m , we

get:

$$\begin{aligned}
 0 &= \sum_{n=1}^N \frac{\partial}{\partial \beta_m} \log p(\mathbf{y}_n, \mathbf{x}_n) & (15.17) \\
 &= \sum_{n=1}^N \frac{1}{p(\mathbf{y}_n, \mathbf{x}_n)} p(\mathbf{y}_n, \mathbf{x}_n, c_m) \frac{y_n - f(\mathbf{x}_n, \beta_m)}{\sigma_{m,y}^2} \frac{\partial f(\mathbf{x}_n, \beta_m)}{\partial \beta_m} \\
 &= \frac{1}{N p(c_m)} \sum_{n=1}^N p(c_m | \mathbf{y}_n, \mathbf{x}_n) [y_n - f(\mathbf{x}_n, \beta_m)] \frac{\partial f(\mathbf{x}_n, \beta_m)}{\partial \beta_m} \\
 &= \left\langle [y - f(\mathbf{x}, \beta_m)] \frac{\partial f(\mathbf{x}, \beta_m)}{\partial \beta_m} \right\rangle_m .
 \end{aligned}$$

Plugging (15.8) into (15.17) we obtain an expression to update β_m ,

$$\begin{aligned}
 0 &= \langle [y - f(\mathbf{x}, \beta_m)] f_j(\mathbf{x}) \rangle_m & (15.18) \\
 &= \underbrace{\langle y f_j(\mathbf{x}) \rangle_m}_{a_{j,m}} - \sum_{i=1}^I \beta_{m,i} \underbrace{\langle f_j(\mathbf{x}) f_i(\mathbf{x}) \rangle_m}_{\mathbf{B}_{j,i,m}} , \\
 \Rightarrow \beta_m &= \mathbf{B}_m^{-1} \cdot \mathbf{a}_m ,
 \end{aligned}$$

where the matrix inverse should be done by a singular value decomposition to avoid numerical problems with singular covariance matrices.

Considering the full set of model parameters we get

$$\beta_m = \mathbf{B}_m^{-1} \cdot \mathbf{A}_m \quad (15.19)$$

with

$$\begin{aligned}
 [\mathbf{B}_m]_{ij} &= \langle f_i(\mathbf{x}, \beta_m) \cdot f_j(\mathbf{x}, \beta_m) \rangle_m \\
 [\mathbf{A}_m]_{ij} &= \langle y_i \cdot f_j(\mathbf{x}, \beta_m) \rangle_m .
 \end{aligned} \quad (15.20)$$

Finally the output covariance matrices associated with each model are estimated,

$$\begin{aligned}
 \mathbf{P}_{y,m} &= \langle [y - \langle y | \mathbf{x} \rangle]^2 \rangle_m & (15.21) \\
 &= \langle [y - \mathbf{f}(\mathbf{x}, \beta_m)] \cdot [y - \mathbf{f}(\mathbf{x}, \beta_m)]^T \rangle_m .
 \end{aligned}$$

Clusters should not be initialized arbitrarily because the algorithm is only guaranteed to terminate in a local likelihood maximum. Also, initializing clusters in places that are close to their final position saves time, since they don't have to walk their way through the data set. We use a method that performs well empirically: Choose $1/N$ as the initial cluster probabilities.

Randomly pick as many points from the training set as there are clusters and initialize the cluster input means, and the cluster output mean with these points. Set the remaining output coefficients to zero. Use the size of the data set in each space dimension as the initial cluster variances. It is also a good idea to normalize the training data to zero mean and unit variance since arbitrary data values may cause probabilities to become too small.

To summarize the model estimation process: (1) Pick some initial conditions; (2) evaluate the probability of the data $p(\mathbf{y}, \mathbf{x}|c_m)$; (3) from those find the posterior probability of the clusters $p(c_m|\mathbf{y}, \mathbf{x})$; (4) update the cluster weights $p(c_m)$, the cluster-weighted expectations for the input means μ_m^{new} and variances $\sigma_{m,d}^{2,new}$ or covariances \mathbf{P}_m^{new} , the maximum likelihood model parameters β_m^{new} , and finally the output variances $\sigma_{m,y}^{2,new}$; go back to (2) until the total data likelihood does not increase anymore [7].

15.2.3 Error Estimation and Characterization

From the probability density of the training data set (15.1) several error estimates and statistics can be derived, each of which provides useful insights and a self-consistency check on the model. The density itself indicates the model uncertainty in that we can't expect to obtain a valid model where the data density is low. The certainty of the model estimate is proportional to the data density in a subspace.

The conditional covariance (15.6) on the other hand indicates the prediction uncertainty given an input \mathbf{x} . It can be related to other characterizations of uncertainty, such as entropy and Lyapunov exponents. The differential entropy of a Gaussian process is $H = \log_2(2\pi e\sigma^2)/2$. Because only differences in a differential entropy matter, we ignore the additive and consider $H = \log_2(\sigma)$. The asymptotic rate of growth of the entropy with time is equal to the source entropy h , which in turn is equal to the sum of positive Lyapunov exponents times the time lag τ between samples, $h = \tau \sum \lambda^+$. Therefore, assuming that the prediction errors are roughly Gaussian, the asymptotic value of the log of the output width as the input dimension is increased provides a local estimate of the source entropy of the system. The sum of the negative exponents can similarly be found by analyzing the time series in reverse order (thereby exchanging positive and negative exponents).

Because clusters find the subspace that is occupied by data, we can use the cluster parameters to find the dimension of the data set even in a high-dimensional space. Intuitively, the number of significant eigenvalues of the local covariance matrices provides an estimate of the dimensionality of the data manifold. For example, we obtain three significant eigenvalues for the Lorenz attractor embedded in 6 dimensions (Fig.15.2). To quantify this further we use the eigenvalues of the local covariance matrices $E_m =$




FIGURE 15.1. The plot shows the Lorenz set, embedded in a three dimensional lag space. The dense dots show the embedded data. Below it are the cluster means and covariances, and the derived input density estimate; above it is the forecasting surface shaded by the conditional uncertainty, showing the maxima associated with the orbit re-injection.

$\{e_{1,m}, e_{2,m}, \dots, e_{3,m}\}$ to evaluate the radial correlation integral

$$\begin{aligned}
 C_m(r) &= \int_{-r}^r \dots \int_{-r}^r p(x_1, \dots, x_D | C_m) dx_1 \dots dx_D \quad (15.22) \\
 &= \operatorname{erf} \left(\frac{r}{\sqrt{2e_{1,m}^2}} \right) \dots \operatorname{erf} \left(\frac{r}{\sqrt{2e_{D,m}^2}} \right),
 \end{aligned}$$




FIGURE 15.2. Fitting the Lorenz set. Top: Data likelihood as a function of iterations. Middle: Mean square error as a function of iteration: Bottom: Sorted eigenvalues of the local covariance matrices.

which in turn lets us compute the cluster’s correlation dimension [8] as

$$\begin{aligned}
 \nu_m &= \frac{\partial \log C_m(r)}{\log r} & (15.23) \\
 &= \sum_{d=1}^D \frac{1}{\operatorname{erf}\left(\frac{r}{\sqrt{2e_{d,m}^2}}\right)} \sqrt{\frac{2}{\pi e_{d,m}^2}} e^{-r^2/2e_{d,m}^2} r.
 \end{aligned}$$

In the limit $r \rightarrow 0$, this dimension is equal to the dimension of the space, because locally the curvature of the clustered space cannot be seen. If it is evaluated at $r = 0.1\sigma_{\max}$, for the e_{\max} direction the contribution is still 0.997, but for a direction with variance $e_{\max}/100$ the contribution drops to 10^{-21} . The expected dimension of the whole data set is finally given by the expectation

$$\langle \nu \rangle = \sum_{m=1}^M \nu_m p(c_m). \tag{15.24}$$

Unlike a conventional $O(N^2)$ calculation of the dimension of a data set from all the inter-point pairs, the clusters find the significant places to evaluate the dimension, and the appropriate length scale at which to test the scaling.

15.3 Application: How to Build a Digital Strad

Mimic synthesis of musical instruments tries to infer models that behave and sound like the original instrument ideally to the extent that original and model become indistinguishable. Given this general goal there have been a variety of different modeling approaches. Global sampling, for example, has been particularly successful in commercial keyboard synthesizers. Each single note of a piano is recorded at many different volume levels and with varying duration and these sounds are replayed during synthesis. Since memory is cheap, only very little interpolation between samples is required and the sound quality is close to the original recordings. However, the method works only for instruments with low dimensional control space, namely, keyboard instruments. Since the model does not know about the instrument's internal state, but only reuses what it has seen before, there is no notion of control on part of the player.

Another successful synthesis technique is physical modeling [21]. It is based on first principles analysis of the acoustics of the instrument, which are implemented in numerical methods. This method provides a lot of flexibility; for example, it allows one to create new instruments that are derived from physical mechanisms but could not be implemented physically. However, the approach has also serious limitations. Current computers can barely run a full-scale model of the violin as can be shown in a simple calculation on a finite element approximation of a violin. Assuming ten body modes per body axis and ten finite element nodes per cycle, we get 10^4 nodes per violin plate and in the order of 10^5 nodes per instrument. If we multiply this by a CD quality sample rate of 40 kHz, we end up with roughly 10 giga instructions per second needed to run a model in real time.

As a further fundamental problem of physical models there is no systematic parameter search within a model structure and instrument family. Given a basic model of a violin there is no way to find the parameters that distinguish a Guanerius from a Stradivarius other than trying out combinations of parameters in a high-dimensional space.

The method we are presenting here lies conceptually between the sampling and the physical modeling approaches and hence is best described as a "physics-sampler". Although we infer our model from recorded data and even use stored samples, we create a model that has the flexibility of a physical model, since we synthesize the physics of the instrument, not the sound. At the same time we are doing computational compression on data, since the physical device is represented in an efficient description.

It was mentioned earlier that the mechanics of violin playing involve stochastic behavior. The stochastic aspects become clear when one considers player and instrument jointly. The violinist only partially controls the instrument. While the player has an idea of the spectral characteristics desired, he or she has no means to hear and control the phase of the produced signal. Naturally there is a causal relationship between the player action

and the spectral content of the sound, whereas the phase of the different partials is random and hence unpredictable.

Fortunately, since phase is not perceived as a discriminating feature in a typical playing situation, we may pick it arbitrarily as long as we avoid discontinuities in the signal components. The general lesson to learn is that we need to model the process, not an instantiation of a particular process. While we can predict deterministic aspects of the signal, stochastic behavior needs to be summarized in appropriate statistics such as the power spectrum.

The violin, like most musical instruments, is characterized by slowly varying boundary conditions that map into a fast audio signal. The non-linear interaction between bow and string causes the slow player motion to be turned into the famous Helmholtz motion which contains the frequency components of the final audio signal [6]. The slow and fast elements describe two different times scales which, if mixed, confuse each other. Instead, fast and slow dynamics and the corresponding state variables need to be treated differently. CWM provides the means to implement such distinction: The slowly varying boundary conditions are used to select the domain of operation (cluster) in the configuration space (Equ.15.3), while the fast dynamics are handled by the local models and the associated state variables (Equ.15.4).

The previous section introduced CWM as a machine learning framework that allows one to predict and characterize arbitrary input-output data. Given this inference tool we need to consider a second important aspect of data analysis and prediction, which is data representation. Although linear transforms such as Fourier or wavelet transforms do not change the information content of the data, they make a considerable difference in which domain we try to predict. CWM lets us embed a variety of specific local representations. In this section we discuss cluster-weighted spectral modeling and cluster-weighted sampling as examples of two local implementations of CWM. We also introduce ways of higher order factorization and show how the CWM structure can be included in a hidden-Markov model to the end of explicitly encoding timing in the model.

15.3.1 Cluster-Weighted Spectral Modeling

It is our goal to build an input-output model of a violin given a data set that contains physical input features measured on the bow and the finger-board along with synchronized audio data. In the training process the network learns the mapping between the physical input and the sound. After training, the network knows how to generate appropriate audio, given new input.

We decompose the audio training signals into spectral frames at a frame rate that equals the sampling rate of the slowly varying physical input. Each frame contains of the coefficients of a short-term Fourier transform

(STFT) applied to a fixed number of audio samples weighted by a Hamming window. The underlying assumption is that the player operates on the spectral composition of the sound and that these spectral characteristics do not change faster than the actual control. From those frames we retain only the harmonic partials of the violin signal. The amplitudes of the harmonic partials are taken to be the magnitude of the power spectrum in the frequency bin, while precise frequency estimates are obtained from the phase difference in closely spaced sample windows [13]. Given a total of P partials the output vector \mathbf{y} has $2P$ components.

The input vector \mathbf{x} consists of physical input data, such as bow velocity, pressure, finger position, and bow-bridge position. Driven by the belief that past input conditions the current state of the instrument the input vector is augmented with respect to past input data. Adding time lagged input samples to \mathbf{x} , we balance the need to include the past and the burden of a big input space. While the model scales linearly in the output dimension, it is very sensitive to large input spaces, since the required amount of training data exponentially increases the input data dimension. Also the model is more sensitive to over-fitting given a bigger input space.

In training we use the set of vector pairs $\{\mathbf{y}_n, \mathbf{x}_n\}_{n=1}^N$ to train a CWM input-output model using simple linear local models of the form $\mathbf{y} = \beta_m \cdot \mathbf{x}$. In the synthesis process the vector of spectral information \mathbf{y} is predicted from new input data \mathbf{x} . Given the spectral data we compute the time domain audio data by sinusoidal additive synthesis, where phase and amplitude of the partials are taken to be the predicted components, linearly interpolating between frames [19]. The final signal is obtained from summing the different components [18].

15.3.2 Cluster-Weighted Sampling

Global sampling has been a successful synthesis technique for instruments with low-dimensional control space, such as the piano [12]. However, the technique is less appropriate for instruments with continuous complex control, such as the violin. In the case of the violin the amount of data required to cover all possible playing situations is prohibitive, since control possibilities are essentially unlimited. To overcome this problem we parameterize the available sample material in an efficient way. CWM learns how to select the appropriate samples, but also to predict the parameters needed to reassemble the sound from the raw material.

Clusters now have multiple output models covering sample selection, amplitude prediction and pitch prediction. The first expert is a pointer into sample space. The cluster that most likely generated a control data takes over and its sequence of samples stands in for the particular playing situation. The cluster is replayed until another cluster becomes more likely and takes over with its own samples. We will come back to the issue of sequencing time domain sound samples later.




FIGURE 15.3. Cluster-weighted sampling: a) overlapping samples of the string signal. b) input-output model, from the bottom: bow position; bow speed; finger position; predicted out-of-samples amplitude (solid) and given sampled amplitudes (dashed); predicted out-of-samples pitch (solid) and given sampled pitch (dashed); the doubled dashed lines indicate overlapping sample windows: the old window is slowly faded out while the new window is faded in, in such a way that the total weight of data adds up to unity at any given moment.

The second output model is a pitch predictor. Given a control input that typically includes the left-hand finger position on the finger board a local linear model predicts the appropriate pitch at any moment in time. The samples selected for synthesis almost certainly won't match this desired pitch exactly. Therefore they are re-sampled with respect to the predicted target pitch. The resampling is done in real time according to

$$\hat{s}(t) = \sum_{n=-N}^{n=N} s(n \cdot T_s) h_s(t - n \cdot T_s) \quad (15.25)$$

with

$$h_s = (\min\{F_s/F'_s\}) \text{sinc}(\min\{F_s, F'_s\}t) , \quad (15.26)$$

where F_s is the stored sampling frequency and F'_s is the target sampling frequency [20]. Sample pitch and target should not differ too much, since big pitch shifts result in audible artifacts. However, resampling can easily compensate for effects such as vibrato. Since we cannot hope to record

any possible vibrato sequence and frequency, we choose to superpose the desired vibrato behavior on the sampled material.

The third output model predicts the sound volume at any moment in time using, once again, simple locally linear predictors. The selected samples are re-scaled with respect to the target volume. Strong modifications of the sample volume should be avoided so that the correct timber is not altered.

This approach requires a number of preprocessing steps that extract the high level properties from the audio data. We need both pitch and volume to label, parameterize and correct the audio data at any moment in time. These properties are easier to obtain than it may seem. Although pitch extraction is a problem that has not been solved in full generality, it turns out to be surprisingly simple in our approach. Since we are measuring physical input data, we have a rather good estimate of pitch to start with. Given a certain finger position, the possible pitch is within a very small frequency interval, which makes it practically impossible for a pitch tracker to get confused in the audio analysis.

An important detail is the sequencing of pieces of audio when there is looping within a sample interval or when a change of cluster occurs. We choose to match samples by minimizing the least square error between the old and the new samples. Additionally we fade out the old sound and fade in the new sound using a Hamming window overlap-add.

Because we re-sample the audio material anyway, we can increase the resolution of our fit, allowing for non-integer alignment of sounds without increasing the complexity of the synthesis algorithm. The success of the overlap-add depends on the length of the permissible fading interval and on the character of the sound. Figure 15.3 shows the overlap of two highly phase-coherent pieces of the string signal of a violin describing a Helmholtz motion. In that case the partials line up nicely with the fundamental and discontinuities are not a problem. However, the sound signal loses its regularity after the filtering by the bridge and the resonant body of the instrument, which makes it much harder to deal with.

15.3.3 Higher Order Factorization: Hierarchical Mixture Models and Hidden-Markov Models

We have demonstrated a flat network structure that is easily applied to many problems and sufficiently complex for most applications. However, there are cases where additional hierarchical structure is helpful if not crucial. Identical models may want to be reused in different areas in the configuration space or systems may have long-term temporal dependences. Jordan and Jacobs [10] introduce mixture models of arbitrary hierarchical depth. Similarly we can add higher level factorization describing global states of our system. For example, the top-level state of a violin model could




FIGURE 15.4. Hidden-Markov model, *from bottom*: cluster/model input space, two clusters per state; state probabilities; predicted out-of-samples amplitude (measured-dashed and predicted-line); predicted out-of-samples pitch (measured-dashed and predicted-line). Although measured and predicted data are visibly different, the reconstructed audio sounds very similar to the original audio data, since the spectral characteristics and the basic characteristics of the sound envelope are preserved.

distinguish global playing conditions such as pizzicato and arco playing or the use of a particular string. The probability density is then expanded as

$$p(\mathbf{y}, \mathbf{x}) = \sum_k \sum_m p(\mathbf{y}, \mathbf{x}, c_m, \text{Model}_k) . \quad (15.27)$$

In the previous sections we used time lags of the input signal to encode temporal structure and memory of the system. Another way of stating this dependence is to say that the current state depends on the past state and the current input. Hidden-Markov-Models (HMMs) have been developed to precisely implement this dependence in a probabilistic framework [17]. If we embed CWM in an HMM structure, we obtain an input-output synthesis

model with an explicit time dependence built into it.

HMMs are typically defined in terms of the number of distinct states q_1, q_2, \dots, q_N ; the state transition probability matrix $\mathbf{A} = \{a_{i,j}\}$, where $a_{i,j}$ denotes the probability that state i follows state j and the emission probability $b_j(k)$, which denotes the probability that the system generates observation k , given that it is in state j . We replace the discrete emission probabilities by a continuous probability density function of the form of $p(\mathbf{x}, \mathbf{y}|q_j)$, which means the cluster probabilities $p(c_m)$ become effectively time dependent, conditioned on past system states.

A cluster (or more than one) now represents a specific state q_j given a set of possible states $q_1 \dots q_N$. The likelihood of a sequence of input-output observations $(\mathbf{X}, \mathbf{Y}) = \{\mathbf{x}_1, \mathbf{y}_1, \mathbf{x}_2, \mathbf{y}_2, \dots, \mathbf{x}_T, \mathbf{y}_T\}$ is

$$p(\mathbf{X}, \mathbf{Y}) = \sum_Q p(\mathbf{X}, \mathbf{Y}|Q) \cdot p(Q) , \tag{15.28}$$

with

$$p(Q) = \pi_{q_1} a_{q_1 q_2} a_{q_1 q_2} \dots a_{q_{T-1} q_T} , \tag{15.29}$$

$$p(\mathbf{X}, \mathbf{Y}|Q) = b_{q_1}(\mathbf{x}_1, \mathbf{y}_1) \cdot b_{q_2}(\mathbf{x}_2, \mathbf{y}_2) \dots b_{q_T}(\mathbf{x}_T, \mathbf{y}_T) . \tag{15.30}$$

$b_{q_i}(\mathbf{x}, \mathbf{y})$ is the emission probability of a pair (\mathbf{x}, \mathbf{y}) given the state q_i . These probability densities may be simple clusters or themselves a sum over clusters,

$$b_{q_i}(\mathbf{x}, \mathbf{y}) = \sum_{m=1}^M p(\mathbf{y}|\mathbf{x}, c_m) \cdot p(\mathbf{x}|c_m) \cdot p(c_m) , \tag{15.31}$$

where the probability distributions are identical to (15.3) and (15.4).

The model estimation is more complicated but is based on the same probabilistic ideas as shown earlier. HMMs are typically trained in a forward-backward procedure which is a special implementation of EM and makes the estimation problem tractable. In synthesis the model is evaluated in a forward procedure since output has to be generated causally [17]. The output sequence at any moment in time is taken to be the expected value of \mathbf{y} given estimated past states and current observed input,

$$p(q_{j,t}) = \frac{\sum_{i=1}^N p(q_{i,t-1}) \cdot a_{j,i} \cdot b(\mathbf{x}_t|q_j)}{\sum_{j=1}^N \sum_{i=1}^N p(q_{i,t-1}) \cdot a_{j,i} \cdot b(\mathbf{x}_t|q_j)} \tag{15.32}$$

$$\langle \mathbf{y}_t | \mathbf{x}_t, q_{t-1} \rangle = \sum_{j=1}^N \mathbf{f}(\mathbf{x}_t, \beta_j) \cdot p(q_{j,t}) .$$

A particular sequence of states now reflects a sequence of input gestures and internal states of the violin. Figure15.4 illustrates a state sequence for simple détaché bowing. We can follow a note from the attack, to the sustained part, to the next bow change.

15.4. Summary

The valuable insights that are possible into signals from complex systems have not penetrated into routine data analysis and engineering practice because of algorithms with limited applicability or reliability. The cluster-weighted modeling framework that we have presented cannot solve all problems, but it does handle nonlinearity and stochasticity in a transparent fashion that provides a clear connection to past practice in a domain (through the choice of the local models), with just a single hyper-parameter (the number of clusters). A natural extension exists for problems that require internal states in the model, without needing to incur the architectural uncertainty of more general graphical probabilistic networks.

One of the most valuable consequences of this probabilistic setting is the range of statistics that can be derived from the underlying model. Rather than imposing a cost function for a learning algorithm at the outset, prediction questions can be answered directly from the density estimate. This is possible with reasonable amounts of data because the estimate is constrained by the local models. Further, the many possible kinds of characterization of the data are done more reliably in a context that can also make falsifiable predictions about the data, including internal consistency checks such as predicting the model's own errors. The resulting models are efficient in storage and computation because the model resources are allocated only where there is data to describe, and the out-of-sample generalization is limited to the reasonable behavior of the local models. These features point to the possibility of broadly applicable "physics sampling," building phenomenological models of driven systems in the space of effective internal degrees of freedom, thereby enabling new applications that figuratively and literally sound great.

Acknowledgments

We are grateful for support from the MIT Media Lab Things That Think Consortium.

References

- [1] Shunichi Amari. Information Geometry of the EM and em Algorithms for Neural Networks. *Neural Networks*, 8(9):1379–1408, 1995.
- [2] Andrew R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39:930–945, 1993.
- [3] W.L. Buntine. A guide to the literature on learning probabilistic networks from data. *IEEE Transactions on Knowledge and Data Engineering*, 1996.

- [4] Martin Casdagli. A dynamical systems approach to modeling input-output systems. In M. Casdagli and S. Eubank, editors, *Nonlinear Modeling and Forecasting*, Santa Fe Institute Studies in the Sciences of Complexity, pages 265–281, Redwood City, 1992. Addison-Wesley.
- [5] W.S. Cleveland and S.J. Devlin. Regression analysis by local fitting. *J. A. Statist. Assoc.*, 83:596–610, 1988.
- [6] Lothar Cremer. *The Physics of the Violin*. MIT Press, Cambridge, Massachusetts, 1984.
- [7] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum Likelihood From Incomplete Data via the EM Algorithm. *J. R. Statist. Soc. B*, 39:1–38, 1977.
- [8] Neil Gershenfeld. *The Nature of Mathematical Modeling*. Cambridge University Press, New York, 1999.
- [9] D. Heckerman and M. Wellman. *Bayesian Networks*. Communications of the Association Machinery. 1995.
- [10] M.I. Jordan and R.A. Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6:181–214, 1994.
- [11] Michael Jordan, editor. *Learning in Graphical Models*. MIT Press, Cambridge, Massachusetts, 1998.
- [12] Dana C. Massie. Wavetable sampling synthesis. In Mark Kahrs and Karlheinz Brandenburg, editors, *Applications of Digital Signal Processing to Audio and Acoustics*, pages 311–341. Kluwer Academic Publishers, 1998.
- [13] R.J. McAulay and T.F. Quatieri. Speech analysis/synthesis based on a sinusoidal representation. *IEEE Transactions on Acoustics, Speech and Signal Processing*, ASSP-34 No.4:744–754, 1986.
- [14] M.E. McIntyre and J. Woodhouse. On the fundamentals of bowed-string dynamics. *Acustica*, 43(2):93–108, 1979.
- [15] Radford M. Neal. *Bayesian Learning for Neural Networks*. Springer, New York, 1996.
- [16] Radford M. Neal and Geoffrey E. Hinton. A new view of the em algorithm that justifies incremental and other variants, 1993.
- [17] Lawrence R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77:257–286, 1989.
- [18] B. Schoner, C. Cooper, C. Douglas, and N. Gershenfeld. Data-driven modeling of acoustical instruments. *Journal for New Music Research*, 28(2):81–89, 1999.
- [19] Xavier Serra and Julius O. Smith. Spectral modeling synthesis: A sound analysis/synthesis system based on a deterministic plus stochastic decomposition. *Computer Music Journal*, 14(4):12–24, 1990.
- [20] J. Smith and P. Gosset. A flexible sampling-rate conversion method. In *Acoustics, Speech, and Signal Processing, San Diego*, volume 2, 1984.
- [21] Julius O. Smith. Physical modeling using digital waveguides. *Computer Music Journal*, 6(4), 1992.
- [22] Floris Takens. Detecting strange attractors in turbulence. In D.A. Rand and L.S. Young, editors, *Dynamical Systems and Turbulence*, volume 898 of *Lecture Notes in Mathematics*, pages 366–381, New York, 1981. Springer-Verlag.

Chapter 16

Data Compression, Dynamics, and Stationarity

Matthew B. Kennel¹
Alistair I. Mees

ABSTRACT One of the main themes of this book is the considerable progress that has been made in modeling data from nonlinear systems that may be affected by noise. In this chapter, we describe a modeling method based on an idealization that gives fast algorithms with known properties based on rigorous results from data-compression theory. The idealization is that the system outputs symbols from a finite alphabet, rather than outputting a real number; we also make a reasonable assumption which is the discrete analogue of the standard embedding theorem. The models that result can be used to simulate and to estimate many of the usual dynamically interesting quantities such as topological entropy. They are also well-suited for a specific new application: testing the stationarity of time-series of discrete symbols, whether two data streams appear to originate from the same underlying unknown dynamical system.

16.1 Introduction

A system with complicated temporal evolution may nevertheless be able to be modeled in a relatively simple way. When the model is constructed directly from data y_1, \dots, y_T , the process is called *reconstruction* in the dynamical systems literature. Building a model may help in answering questions such as whether there is a degree of predictability, whether the system is stationary, and so on. Indeed, questions like stationarity cannot really be separated from the existence of a model. Ideally, reconstruction should settle these questions by selecting the most appropriate model from a wide range that includes elements of all the above classes: If, for example, the best model is non-stationary, then it is reasonable to declare the original system to be non-stationary, as long as we remember that we can never find *the* best model, only the best among those classes we have chosen to consider.

The choice of classes of model will be guided by other knowledge about

¹ Author for correspondence.

the system, but there do exist “universal” classes which can in principle model any system. The modeling method described in this chapter is universal within the assumptions made. Likewise, more traditional dynamical reconstruction works with relatively general model classes, such as discontinuous local linear functions [1, 7, 28], continuous piecewise-linear models [19, 20], neural nets [32], radial basis functions [3, 21], and so on. Theoretical asymptotic guarantees of universality are only useful in practice if they give good quality models for finite amounts of data, and if they also give insights as well as high quality models.

The embedding theorem [29], discussed elsewhere in this volume, applies to real-valued measurements from smooth dynamical systems; it allows system states to be represented by vectors of delayed output signals. The embedding theorem as originally stated does not allow for noise, and its implication for modeling is that if there is finite dimensional dynamics, the measured data can be written

$$y_t = f(y_{t-1}, y_{t-2}, \dots, y_{t-k}) \quad (16.1)$$

for some unknown function f and integer k , the so-called *embedding dimension*. The value of k is determined by methods such as false nearest neighbors [14] and the function f is determined by one of the above approximation methods. Realistic modeling of dynamical systems requires that f be allowed to be nonlinear, and the emphasis on this is the distinguishing characteristic of dynamical systems reconstruction.

All realistic methods allow for noise as well as for nonlinearity in the dynamics; generalizations of the embedding theorem to noisy systems now exist [27]. Usually, the model in (16.1) is modified to

$$y_t = f(y_{t-1}, y_{t-2}, \dots, y_{t-k}, \varepsilon_t), \quad (16.2)$$

where ε_t is the unknown noise (and fitting errors!) at time t . The question of how much variation to ascribe to noise and how much to nonlinearity is best answered using information-theoretic methods [21, 12]. Model (16.2) is not as general as it needs to be, but only rarely is it generalized because of the difficulty of combining nonlinear function fitting and explicit state error estimation.

A considerable simplification is to assume the observed data consist only of a finite set of symbols. In return for this simplification, we gain powerful reconstruction methods which can have good qualities even with short data series. The reduction to symbols is realistic in that observed data are invariably digitized, but we shall allow for, and typically use, substantially coarser discretization than is inherent in the measuring process. Symbolic methods have been used in the study of dynamical systems from the earliest days, but in contrast to most of the theory which has flowed from Kolmogorov and other mathematical dynamicists, we make no attempt to craft a special discretization to match known dynamics; we are not trying to make a Markov partition [17], even approximately.

For our purposes, symbolization is merely a reduction—possibly a considerable reduction—in precision of the original data. Typically, we divide the one-dimensional histogram of the observed data between a small number of bins, either by equal probability mass or by equal width, and then code each datum by the bin number, producing a symbol stream, $s_1, s_2, s_3, \dots, s_N$, each symbol from some alphabet A re-expressed as integers $s \in \{1, 2, \dots, |A|\}$. There will typically be a trade-off between the alphabet size and the complexity of the model: the relationship between dynamics and shifts [17] suggests that a smaller alphabet may be acceptable if we are prepared to deal with longer histories (that is, longer multi-symbol words) when model-building.

The distribution of multi-symbol words provides information about time-dependent structure and correlation, just as, with continuous nonlinear data, time-delay embedding provides a vector space revealing dynamical information. The techniques introduced by Fraser [8, 9] to measure mutual information and related concepts are familiar to applied nonlinear dynamists, but our aim is to make explicit the modeling step, which is implicit in Fraser's methods. A significant benefit is that, at least in principle, we can deliver guaranteed convergence and optimality, something that is rare in nonlinear reconstruction.

The central payoff of symbolization is that we may exploit powerful contemporary methods taken from the field of data compression. It is apparent that useful data compression involves modeling: The better one's model, the better one may compress the data. In data compression there exists a class of compression algorithms known as "universal", which have been proven to be able to model any source with a non-pathological probability law. The metric of quality is, of course, a model good enough to compress data down to the Shannon entropy of the generating source². Universal compression methods "achieve entropy" in that their asymptotic code length per input symbol converges to the true entropy rate. We will exploit the modeling phase of such algorithms, rather than being interested in the literal output of compressed bits.

There is more to time-series analysis than good prediction and modeling. Our application of data compression methods solves, in a fundamental way, a common problem in deriving valid statistical tests for complex dynamical data: accounting for arbitrary correlation and dependence between the data. As a result, it allows calculation of dynamical invariants as well as application of other statistical tests. As an illustration, in this contribution we construct a direct test for dynamical stationarity.

²The Shannon entropy is an average and is only a fundamental minimum in the asymptotic sense; however, if the data stream is more compressible than implied by the entropy of the source, then it is atypical of the source and any model we make from it will not be a good model of the source.

16.2 Symbolic Prediction and Coding

From observed data one may estimate either absolute or conditional probabilities. Statisticians call the first task “density estimation” and the second, “regression”. Many information theoretical quantities have equivalent formulations using either absolute or conditional probabilities; however, there is a practical difference. Given finite observed data, especially from chaotic dynamics, it is usually easier to estimate conditional rather than absolute probabilities. For example, the invariant measure of a strange attractor is a complex geometric object with a difficult probability distribution, while it often arises from a simple and smooth dynamical evolution law which is easier to model parametrically. Additionally, modern powerful Bayesian numerical methods, such as Markov Chain Monte Carlo, rely on iterative estimation of marginal distributions to discover absolute distributions.

Assume, then, that we have an observed symbol string s_1, \dots, s_N , which was output from a finite-memory source, so that

$$P(s_t | s_{t-1}, \dots, s_1) = F(s_{t-1}, \dots, s_{t-k}). \quad (16.3)$$

That is, the probability of the next symbol conditional on the entire past depends only on the previous k symbols, where again we have to somehow estimate k (which may depend on the “context”, i.e., on the actual symbols observed), and we also have to estimate³ the conditional probability F .

Notice the parallels between this model and the standard embedding models: Our context size k is like a “local” embedding dimension [13] and estimation of the conditional probability is equivalent to reconstructing the dynamics in more conventional models. Recasting the procedure in a more statistical way, our intention is to model serial correlations in the observed data by conditioning on past observations, choosing sufficient observations to ensure that, as far as can be detected from the data, successive observations are independent under the conditioning. That is, we split the problem into one of detecting the conditional *structure*, and another of estimating *distributions* of I.I.D. random variables. It is convenient to think of the structure determination as the deterministic part of the modeling process and the distribution estimation as the stochastic part.

An obvious model is a simple finite-order Markov chain, estimating the histogram of future symbols s_t given every possible combination of the k past symbols. There are $|A|^k$ possible bins, each corresponding to a specific past history, and containing an estimate of the next symbol. Although not infrequently used by dynamicists doing symbolic time-series analysis, this is a very poor modeling technique. There is a global free parameter

³If we restrict the choice of approximating functions \hat{F} for F , then we may increase the uncertainty, which will show up in the model as transition probabilities which are more nearly uniform than they should be. We are, as in most dynamical modeling, confounding modeling error and stochasticity.

k which somehow must be selected, and there is an exponential explosion in the number of local parameters which must be estimated. Larger k lets one model higher-dimensional dynamical laws, but because of the exponential explosion in the number of bins, the amount of data available to estimate the next symbol can be very small, producing a bad estimate. Furthermore, there are no guarantees that such an estimate even converges to the true answer given more and more data: the canonical “over-fitting” problem. Such a model is not remotely competitive in the data compression community.

As a replacement, we build a so-called *tree machine* which acts like a stochastic dynamical system. In effect, this technique automatically discards or under-weights distant past histories when they have little discriminatory benefit, but keeps them when they improve performance.

What is the proper measure of performance? There are good reasons [16] to define “better” to mean “able to compress the data more”. The ultimate bound on performance is, of course, that one may not, averaging over realizations of the data source, compress data at a rate better than the Shannon entropy of the source. That is, if P were the *true* underlying probability law governing the source generating the data, the best compression rate achievable would be, assuming ergodicity and stationarity,

$$h = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N -\log_2 P(s_t | s_{t-1} \dots s_1) \quad (16.4)$$

bits per symbol. Any other probability law $\hat{P} \neq P$ produces a larger code length when substituted in in (16.4), except for exponentially unlikely atypical strings. Thus, minimizing the code length

$$L = \sum_k -\log_2 \hat{P}(s_k | s_{k-1} \dots s_1) \quad (16.5)$$

is a good universal metric of performance for any modeling scheme which learns \hat{P} from observed data.

This code length is not just an impractical lower bound. Given a probability estimator \hat{P} and subsequently the actual symbol $s(t+1)$, a contemporary coding scheme called *arithmetic coding* [4] produces a literal compressed bit string of length less than $\lceil L \rceil + 2$. There is an accompanying decompressor that can reconstruct the same probability estimator and, given the compressed bit string, the actual symbol $s(t+1)$. The coding and decoding stages are not relevant for our current needs, but we mention them here to show the relationship between on-line modeling of symbolic dynamics and data compression.

For the decoding process to work, the encoder (and hence the estimated code length) has to be causal: It cannot use information about the present or future symbols when it encodes the present symbol. This is the

information-theoretic version of the ubiquitous over-fitting problem: the statistical equivalent is that if a model has too many parameters, its error in “predicting” in-sample data will be smaller than when predicting previously unseen data. The more free parameters there are in the model, the more likely this problem is. A considerable statistical literature is devoted to schemes to regularize or validate models to prevent over-fitting or to add penalty terms for free parameters.

The information theoretic version of the problem suggests a solution which is particularly easy to apply in the symbolic case. The idea is to include the cost (in bits) of specifying the values of the free parameters, including any necessary to specify the model’s structure, as well as the data set itself, denoting the sum as the “description length” (DL). Assuming the broad model class is known at both ends of a transmission link, it is possible to reconstruct the data set at the receiver using the number of bits specified by the DL. Rissanen’s minimum description length (MDL) principle [24] tells us to select the specific model out of the class of models considered which gives the minimum description length on the dataset considered. This criterion is a practical analogue of the fundamental Kolmogorov complexity of a dataset: the minimum size of a program on a universal computing machine capable of reconstructing the data, a metric which is, unfortunately, non-computable [16].

The MDL principle can be applied either in a batch mode or sequentially. The former is more familiar in the dynamical systems literature, having been used to prune regression models with radial basis [12, 21] and polynomial [2] functions. The description length is the sum of the regression errors plus a term which accounts for the number and size of non-zero free parameters. Although the principle is legitimate, the calculations necessary to fairly account for the free parameters can be tricky and conceptually obscure, especially in the continuous domain, requiring a number of approximations of varying quality. The sequential coding method, more common in traditional data compression, avoids these difficult problems. Here, we update the estimating model as appropriate after each observation has been processed. That is, the estimate for $\hat{P}(s_t)$ is generated using only knowledge of the previously occurring symbols, and the code length L in (16.5) is thus a true description length. The costs for the parameters are automatically included in L because the estimator is not as good in the beginning of the dataset, when few symbols have been observed, as toward the end.

16.3 Context Trees

A *context tree* is a model of dynamics which is a close symbolic-system analogue of the embedding theorem. It defines the conditional probabili-

ties of the next symbol, given a number of preceding symbols. The number required may vary, just as, in the English language, the next-letter probabilities are almost completely defined if the current letter is “q”, whereas it is much more informative to condition on “th” than on “t”. In general, we should choose sufficient past observations to ensure that, as far as can be detected from the data, successive observations are independent under the conditioning.

One way to look at the context tree construction is that it splits the problem into one of detecting the conditional structure, and another of estimating distributions of i.i.d. random variables.

First, suppose we already know the conditional structure; that is, we are given a set of *contexts* C_i , $i = 1, \dots, |C|$, such that $P(y|C_i)$ is independent of $P(y|C_j)$ for $i \neq j$. Being past sequences of variables, the contexts may be conveniently represented as the leaves of a tree. Figure 16.1 shows a




FIGURE 16.1. A simple context tree which generates the periodic sequence 011011011 Each leaf is a *context*; it represents the most recent symbols seen, read backwards in time from the present and downwards from the root of the tree. Probabilities at the leaves are conditional probabilities for the next output symbol, given that context.

possible context tree for the alphabet $\{0, 1\}$, with conditioning contexts 0, 01 and 11. In the tree, the most recent symbol is nearest the root so, for example, the leaf corresponding to context 01 is found by following the arc labeled 1 at the top level and the arc labeled 0 one level deeper. The context tree is a machine that outputs independent realizations given any context: in Figure 16.1 there is probability 1 that the next symbol will be a 1 and probability 0 that the next symbol will be a 0, regardless of previous symbols before the two represented in the context.

16.3.1 Zero-Memory Estimation

If the context tree structure is given then we can estimate the future-symbol conditional probabilities for any input string by recording at each

leaf the number c_i of times the i th symbol has been seen already at each leaf, then estimating the multinomial distribution of symbols there. One widely-used estimator in the binary case is the *Krichevsky-Trofimov* (KT) estimator [34], which has good asymptotic properties.⁴ It is

$$P(0) = \frac{c_0 + 1/2}{c_0 + c_1 + 1} \quad (16.6)$$

where c_0 and c_1 are the counts of observations of 0 and 1; that is, it is equivalent to dividing a “ballast” of one observation equally between the binomial bins before the start of observations. If we have prior knowledge or prejudices about the system, we may want to use a different amount of ballast; more if we believe we have a good estimate of the distribution, less if we do not have such an estimate but nevertheless believe the system to be nearly deterministic. The binary KT estimator is in fact the posterior for a Dirichlet(1/2, 1/2) prior.

We need ballast (a zero ballast corresponds to the maximum likelihood estimator) because this probability estimate is used as input to a coder, or at least to estimate the code length that one would achieve. Unseen future symbols would lead to an estimated probability of zero, but if that symbol did actually occur (and this is likely in sequential coding) it would imply an infinite code length. In circumstances when an estimate is not used to compute a code length, the ballast can be zero.

In accordance with our earlier discussion about estimation of description lengths, we estimate the distribution predictively; that is, the probability we ascribe to any symbol is the probability ascribed by the KT estimator based only on what has been seen in the past. In the binary case, the probability ascribed to an input sequence containing c_0 zeros and c_1 ones is therefore

$$\hat{P}(c_0, c_1) = \frac{\frac{1}{2} \times \frac{3}{2} \times \cdots \times (c_0 - \frac{1}{2}) \times \frac{1}{2} \times \frac{3}{2} \times \cdots \times (c_1 - \frac{1}{2})}{1 \times 2 \times \cdots \times (c_0 + c_1)}. \quad (16.7)$$

In the general case, each node accumulates $|A|$ integers which record the occurrences c_i of every symbol $i \in 1 \dots |A|$ that occurred immediately after its particular context C , and estimate the conditional probability for the k th symbol using the Krichevsky-Trofimov [26] estimator, to give

$$\hat{P}(k) = (c_k + \beta) / \sum_j (c_j + \beta). \quad (16.8)$$

Here β is the ballast factor, which was 1/2 in (16.6). In this chapter we use a fixed ballast of $1/|A|$ in each bin; we have found this to work well in practice over a wide range of problems.

⁴The non-asymptotic properties of the KT estimator may not be so good, especially for larger alphabets. Other estimators may be more suitable for particular problems.

16.3.2 *The Structure of the Tree Machine*

Unless there is some special knowledge available, we do not know the tree structure in advance and we will have to induce it from the data; this is, in effect, the deterministic part of our modeling procedure. It is possible, and common, to specify a maximum memory, and hence a maximum tree depth. This can be avoided by recording the entire past for every symbol. This means the maximal depth of the tree is equal to the length T of the input, but with efficient implementation which avoids explicit construction of unbranched chains of links, particularly at tails [33], the storage required is linear in T .

Each leaf now represents just one observation, and the KT estimator is trivial there. We also find KT estimators for the internal nodes by scanning backward from the leaves and combining the counts of the children of every node: If we were to use the KT estimator at a particular node this would be equivalent to pruning its children, since we would be claiming that the context at that node is sufficient for conditioning in prediction of future symbols. The problem is to decide where to prune the tree. The critical issue is balancing between the more detailed dynamical reconstruction possible in deeper contexts and the increased quantity of observations at shallower contexts which gives greater robustness against statistical fluctuation.

One must choose from or combine the estimates at different depths to make the best estimate. We discuss three alternatives: selecting a node for each new observation at each time step; weighting among all matching nodes; and batch pruning the tree after all observations have been seen.

The question central to all these methods is “When is the model at a node better than the best model obtainable from its children?” Better performance here means better compression ability. Compression can be defined in terms of probability distributions: In fact, if $P_s(x)$ is the probability ascribed to data sequence x as seen⁵ at node s by (16.7) then the code length required to encode x there is

$$-\log_2 P_s(x)$$

bits. So all of the alternatives need to recognize the data compression abilities of each context; for example, we could define a rule that prunes off any nodes that jointly compress worse than their parents in the context tree.

In each of the following implementations, we insert the symbol into the tree by adding branches and nodes as needed to store its entire history, with any required optimizations. One updates the histograms at all nodes whose counts have changed (i.e., for every node which matches the current history) and recalculates corresponding description lengths. Usually one or more code-lengths will be stored at each node. We will not comment further

⁵The sequence “seen” at a node is the subsequence of the original data for which the contexts match the node’s context.

on the insertion and update process; any reference to “inserting the new symbol” will assume that all of this is being done.

On line Node Selection

First we describe a method of selecting contexts as the tree is built; that is, we are defining optimal coding nodes predictively, in accordance with our guidelines for calculating description length. Our state-selecting algorithm is a version descended from Rissanen’s original work introducing predictive context trees [23].

Given the existing tree at time t (before the insertion of the new symbol), the algorithm selects a particular node for encoding. In addition to the accumulators c_k used in (16.8), each node retains an internal code length L_e , the cost of coding using that node’s own local probability estimator all symbols which have occurred after histories matching that node’s context. This is a measure of how good that node is at coding the future when the past matches its context, or, since $P_e = 2^{-L_e}$, the total probability of observing all futures after matching the current node.

In this algorithm we insert the current symbol before any other computations, but keep the updated counts separate until the symbol’s code-length has been computed.

Given the current symbol, we traverse the nodes which match the current history from the root node down. If the tree is as in Figure 16.2, and the




FIGURE 16.2. Example of a small context tree for a 3 symbol alphabet. Internal nodes (nodes with deeper children) are the root node, A, C, and AC, and terminal nodes, AA,BA,CA,B,BC,CC. Descendents of AC continue off the figure. Each node accumulates counts of future symbols and two internal code lengths.

recently seen symbols are ABBBC (C being the most recent), the nodes in the tree which match the current history are the root node (matches everything), C and BC. The current node n is the “encoding node” for the

next symbol (in the example, coming after the 'C') if it is a terminal node (i.e., if it has no deeper children whatsoever as one has hit the beginning of the data), or barring that, if

$$L_e(n) \leq \sum_c L_e(c) \quad (16.9)$$

with c summing over all extant children of n . Otherwise, one descends one level deeper to the matching child and repeats. The notion is that we wish to find nodes at the level where, until now, the current counts better predicted the future (measured via L_e) versus descending to a deeper level.

The third step is to increment L_e at each node n which matches the history by the code length obtained with that node's estimator evaluated on the actual subsequent symbol: $L_e \leftarrow L_e - \log_2 \hat{P}(s_{t+1})$. The total code length is incremented by the estimator at the chosen encoding node. Finally, the appropriate count is incremented, $c_k \leftarrow c_k + 1, s_{t+1}$, for every matching node. For the comparison to be fair, $\sum_k c_k(n) = \sum_c \sum_k c_k(c)$ must be maintained as an invariant of the tree.

Weighted Contexts

Willems and coworkers [34] introduced a universal coding scheme called *context-tree weighting*, which, instead of choosing between a node and its children, weights them according to a recursive formula. This method has excellent theoretical properties, with upper bounds on the coding redundancy (code length above Shannon's limit) for individual strings, not just in probability.

Here the code-length of the string up to time t always appears at the root node of the tree constructed up to that time. For any empty node the code-length is zero; otherwise, for any node n it is $L_w(n) = -\log_2 P_w(n)$ where $P_w(n)$ is the weighted probability, computed as follows. Let $P_e(n) = 2^{-L_e(n)}$ be the KT probability at the node (that is, the probability computed assuming zero additional memory beyond that implied by the node's context). Define

$$P_w(n) = \begin{cases} P_e(n) & \text{if } n \text{ is a leaf;} \\ \frac{1}{2} \left(P_e(n) + \prod_{i=1}^{|A|} P_{i,w}(x) \right) & \text{otherwise.} \end{cases} \quad (16.10)$$

Thus $P_{i,w}(n)$ is the probability ascribed to child i of s by the weighted estimator.

We now insert the current symbol into the tree, updating L_e and L_w at each node. The code-length for the symbol currently observed is the increment in $L_w(\text{root})$ on adding the symbol.

Here no single node has been used to provide a probability estimator to an arithmetic coder which would encode s_{t+1} ; instead, all contexts which match the current history will contribute to some degree, weighted by their

past performance. In effect, the implicit source model of this method is a weighted blend of all possible tree topologies. Although it is not obvious from this description, the Willems method is in fact implementable as literal data compression with an arithmetic coder.

Batch Pruning

This blending of trees, though theoretically elegant, presents some practical difficulties for some of our desired applications. We would like to be able to assign every history to a particular node⁶ which best models its future evolution. We use the quantities accumulated by the weighted tree construction to extract a single tree and we use its leaves as states. This simplifies our statistical tests and allows a convenient conversion of a tree model to an equivalent Markov chain.

To construct the single tree model, we build the weighted tree for all of the data, in exactly the same way as the Willems method. Then we examine the final code lengths by a depth-first scan of the tree, declaring a node n to be “terminal” by pruning off all children when $L_e \leq L_w$.

Given additional data from the same source, we would make a probability estimate by following the history as far as possible down the tree until a leaf has been reached, and using that node’s estimator for the future distribution. We may also similarly re-run the original data back into this fixed tree thus finding distinct “encoding nodes” (the terminal matching contexts) for each datum, as the first algorithm provided. This is important for our stationarity test.

Constructing a batch model for the data set via an intermediary of a weighted, sequential coding algorithm may seem to be a needless circumlocution. We found it necessary, however, because direct batch estimators of the description length and application of an MDL principle did not provide stable or empirically satisfactory trees. The point is that the L_e and L_w quantities were accumulated fairly by causal sequential estimation and thus implicitly include costs for parameters. We have post facto selected a model from the weighted tree, something that can certainly be done by the receiver at the end of the thought-experiment of transmitting the data using the Willems compressor, without any further communication from the transmitter.

⁶Such a node plays the role of a state in normal dynamical modeling. Physicists and engineers tend to prefer such models, though Bayesian statisticians appear to be happy with weighted models having no identifiable states.

16.4 Using Tree Machine Models

16.4.1 Compression

Since the purpose of the original algorithm is to compress, we should check that it performs in that respect. In the first 5 rows of Table 16.1 we show

	Bits (uncomp)	Bits (comp)
Fair coin tosses	10,000	10,008
Base 5 random	232,193	232,224
Period 3 binary	57	12.9
Period 4 binary	92	17.8
Above 2 concatenated	149	31.1
Lorenz 84, alphabet 10	132,877	57,234
Shakespeare play, 128 symbol ASCII	1,082,952	383,558
Same, lower case, no punctuation, 27 symbols	1,010,824	284,739

TABLE 16.1. Compression performance of algorithm on various simple inputs: random and simple deterministic cases. Lorenz84 is a 3-dimensional chaotic dynamical system; the data came from integrating the differential equations and discretizing the x coordinate into 10 equal-width bins. See the text for discussion of the Shakespeare play.

results of applying the Willems method to inputs which are random, deterministic with short periodicity, and a concatenation of two deterministic series with different periods. In each case the expected results are obtained: we fail to compress the random data, with the overhead being approximately that expected from the known redundancy of the weighted tree method; we compress the simple deterministic data greatly (and the machines, which are not shown, are easily seen to be the simplest possible); and the code length for the concatenated deterministic sequences is slightly longer than the sum of the code lengths when the two are compressed separately, giving a weak indication of non-stationarity of this sequence. The 6th row compresses very crudely discretized data from a chaotic dynamical system. The last two rows are discussed in Section 16.4.2.

16.4.2 Simulation

To produce a realization of the context tree machine, allow it to start in any context and emit a symbol according to the estimated probability distribution there. This, together with previous output symbols possibly as far back as the initial context, defines a new context, and we iterate the process. If we believe that the tree machine was built from a good sample of the process, it makes sense to estimate the conditional probabilities *for simulation purposes only* via the maximum likelihood estimator; that is, the

emission probability for the i th symbol is taken to be $n_i / \sum n_j$ rather than the KT estimate. This ensures that, for example, a deterministic system is not randomly noised-up.

We already saw a deterministic example in Figure 16.1; it is clear that the tree there emits the sequence claimed. All of the examples in Table 16.1

a	dear	her	nature	son
ace	deed	here	nay	soon
acknowledge	demand	hereafter	never	speak
act	denial	high	news	speedily
actiseriously	devil	him	no	stake
admiringly	diana	himself	noble	stay
afraid	die	his	nobleman	steal
after	digents	hither	noise	stepone
again	direct	holy	nose	steward
against	dischart	home	not	stire
aim	disclown	honest	now	strong

TABLE 16.2. A test of the tree machine on some English text (Shakespeare's *As You Like It*). The text was preprocessed to replace all punctuation by spaces and convert all upper case letters to lower case; multiple spaces were reduced to single spaces and the processed text was used as input data to a context tree building computer program, which was then used to generate simulation output. The resulting text was split on spaces and some of the "words" produced are shown above. In spite of a couple of nonsense words and a couple of plausible but incorrect English words, the machine has captured a remarkable amount of the structure of English words. In other tests without the input preprocessing (not shown) it also captured a great deal of the punctuation and layout of the original, and a little of the phrase and sentence structure.

were tested in this way and produced good results. The compression performance on a sample of English text was also shown in Table 16.1. In Table 16.2 we show some of the simulation output from the tree machine built from the same data, which has arguably performed very well.

16.4.3 *Converting Tree Machines to Markov Chains*

It is tempting to assume that because a tree machine is based on conditional probabilities, it is equivalent to a Markov chain with states corresponding to the leaves of the tree. This is not so; for example, the tree machine represented by solid lines in Figure 16.3 is not a Markov process because it requires additional memory outside the tree. If at leaf A a 1 is emitted, the context is now 01, corresponding to internal node E. But node E is not a leaf and so is not part of the machine as such: We need to know what came




FIGURE 16.3. The solid lines show a tree machine which is not a Markov chain: that is, the contexts (leaves of the tree) cannot be used as states because they are incomplete. By extending the tree as shown by the dashed lines, we do obtain a Markov chain. This can be done in general.

before the 0 so that we can choose between nodes G and H to determine the distribution for the next output symbol.

By adding children C and D to node A as shown by the dashed lines, we do obtain a Markov chain. In general, we need to add children until the tree has the property that every subtree is also a subtree at the root. The probability distributions at added child nodes such as C and D are defined to be the same as at their parent A.

Clearly, a simulation output from such a Markov chain will be identical to a simulation produced by the original tree machine. More usefully, we can easily compute the equilibrium probability distribution of the Markov chain and hence obtain an invariant probability distribution for the contexts. For a system which really is discretized, this allows immediate calculation of any of the standard dynamical invariants, which are averages over the equilibrium distribution. For a system assumed to have continuous states, which are merely approximated by symbols, there is further work to be done in proving convergence of the equilibrium probabilities to one of the invariant measures, most likely the SBR measure [10]. Our experience is that the equilibrium distribution is indeed a good approximation to the SBR measure. Applications will be described elsewhere.

16.4.4 Topological Entropy

Topological entropy is a particularly important invariant in dynamical systems [36]. It gives the rate of generation of new information (or equivalently, the uncertainty per iterate) for the dynamical system. For a finite state machine, it can be calculated from the topological transition matrix T which is the same as the Markov chain transition matrix with all nonzero elements replaced by 1. The topological entropy in bits per iteration is then $\log_2 \lambda$ where λ is the maximal eigenvalue of T .

For example, the tree machines in Figures 16.1 and 16.3 have topological transition matrices

$$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 \end{pmatrix}$$

where in each case the states are the leaves of the (possibly augmented) tree read from left to right. The corresponding topological entropies are 0 and 0.91 bits per iteration. The 0 for the first case reflects the fact that the machine is deterministic, so there is no uncertainty at all in the prediction. Again, we defer realistic applications to a later paper.

16.5 Application: Simulating a Chaotic Laser

We have already shown the results of simulations on toy examples, such as repeated short periodic sequences of symbols, and sequences produced by pseudo-random number generators. We now present one somewhat realistic simulation, output from a laboratory laser experiment where the electric field amplitude fluctuates in a complex manner. This complex dynamics is not easy to model by most methods. The most successful models are those produced by nonlinear reconstruction methods [1] but although they tend to predict well over the short term, they do not simulate well; that is, they get short-term predictions right but do not model the iterated dynamics in a fully satisfactory manner. We show here that a tree machine estimator does better on the dynamics and produces a good estimate of topological entropy, though the simple version considered in this chapter does not yet give an extremely high quality emulation of the dynamics. The best current continuous reconstruction methods [22] do somewhat better but are enormously more computationally intensive.

The original data were digitized to 8 bits. By experimenting with discarding different numbers of low-order bits, we determined that the best overall compression of the data resulted when we truncated the data to 5 bits accuracy, corresponding to an alphabet of size 32. We used 10,000

data as input to the the tree machine algorithm, which compressed the 50,000 bit input to approximately 17,900 bits using a Willems weighted context tree in around one second on a 266MHz Pentium-based computer. By construction, the tree machine has rather accurately captured all the simple statistics like mean and variance; it also has an RMS fitting error of around 2 percent which is less than the truncation error (one part in thirty two). Its one-step prediction accuracy, which is of the same order, is therefore near the limit of what is possible given the truncated data.

The machine was used as a simulator and gave the results in Figure 16.4,




FIGURE 16.4. Experimental data from an irregularly fluctuating laser, and output of a tree machine model made from 10,000 data, truncated to 5 bits precision. The model output is a “free run”: at every instant, the output depends only on the model’s current context without reference to the real data. The two graphs are simply typical segments, and do not refer to the same time or initial conditions.

which shows a segment of the experimental data and a segment of the simulation . The simulation performance is very good but further examination (with much longer simulation runs) reveals that it is not as good as the very best recent dynamical modelers [22] can produce; there are well-understood reasons for this and further work may be able to improve this. In the meantime, we have an extremely rapid and entirely automatic modeler which

allows calculation of dynamical invariants, can predict accurately, and has useful simulation power.

16.6 Application: Stationarity Testing

Stationarity, the notion that the dynamical law describing the system does not change over long time-scales, is a prerequisite for the vast majority of nonlinear data analysis techniques and is an assumption we have implicitly made so far in this chapter. Only recently have hypothesis tests suitable for realistic chaotic and nonlinear dynamical data been proposed [15, 25, 35]. The symbolic approach of this chapter allows us to justify the statistical assumptions in a deeply principled way.

A first attempt at a stationarity test might be to apply the classical χ^2 -test to observed counts of distinct multi-symbol words of a fixed length observed in, say, the front and back halves of the data. Unfortunately, the assumption underlying this test—that each datum is randomly and *independently* drawn from some distribution—is not true in realistic dynamical data. Short time correlations in physical data strongly couple symbols near in time; thus naive application of such tests fails miserably. Indeed, arbitrary dynamical dependence makes it difficult to construct a proper statistical null test for any hypothesis which allows chaotic or general nonlinear data in the null class, and few examples of this sort exist.

A better stationarity test is to construct a context tree from the entire data set and, instead of using the model to predict or emit a compressed binary stream, to examine the statistics observed in the model to discern stationarity. Answering the question “do two data sets appear to arise from the same underlying dynamical system” translates to combining hypothesis tests performed at each encoding node regarding whether the distribution of future symbols actually encoded—whether from the first set or the second—could have come from a single underlying probability distribution and if any apparent difference is statistically significant. At encoding contexts, we may use standard tests because these events ought to be nearly independent; after all, a context tree model assumes that once the appropriate matching context has been found, the best model of the dynamics is that the next symbol is randomly drawn according to some distribution, which it estimates for that node. If there were more usefully extractable dependence in the data, then the context tree modeler would have gone to deeper nodes to distinguish them until the future looked random again.

We modify the context tree construction by counting the two data sets separately. Every node records the frequency with which symbol k was encoded there, $e_{k;1}$ in the first set and $e_{k;2}$ in the second. (Note that $e_k \neq c_k$, the latter accumulating frequencies whenever a context was excited.)

Assuming independence, the statistic

$$\chi^2 = \sum_{k=1}^{|A|} \frac{(R^{1/2}e_{k;1} - R^{-1/2}e_{k;2})^2}{e_{k;1} + e_{k;2}} \tag{16.11}$$

with $R = \sum e_{k;2} / \sum e_{k;1}$ follows the standard χ^2 distribution with $|A| - 1$ degrees of freedom under the null hypothesis that both empirical probability distributions came from the same underlying distribution⁷. Given the value of χ^2 and the degrees of freedom, standard numerical algorithms provide a likelihood L asymptotically uniform $L \in (0, 1)$ under the null. Small values of L reject the null at the given significance level, e.g. $L < 0.01$.

It is known that the analytic approximation used for the asymptotic distribution of the χ^2 statistic becomes increasingly inaccurate as the number of observations decreases. Thus for $\sum e_k < 75$ (a somewhat arbitrary cut-off) we switch over to a combinatorial test for differences in proportions, called *Fisher’s exact test*. The calculations for this test are easy only in the 2×2 case. We coalesce bins by keeping the observation for the most frequent symbol (bin m which achieves $\max(e_{m;1} + e_{m;2})$) and merging the others into $e_{o;1}, e_{o;2}$, resulting in four quantities conventionally expressed in a “contingency table”, with cumulative row and column sums:

$e_{m;1}$	$e_{o;1}$	n_1
$e_{m;2}$	$e_{o;2}$	n_2
n_m	n_o	N

Under the null that the difference in proportions between m and o counts is independent of being in set 1 and 2, the probability for seeing any particular table with the given marginal sums is:

$$p_T = n_m!n_o!n_1!n_2! / (e_{m;1}!e_{m;2}!e_{o;1}!e_{o;2}!N!).$$

One directly enumerates all tables with the given observed marginals (only a one-dimensional sum for a 2×2 table) and sums p_T for every table with a difference in proportions at least as great as that observed⁸, resulting in a likelihood L for accepting the null hypothesis at this node.

We combine these M likelihoods, each measuring some aspect of the same null hypothesis, into a single overall test. Under the null, the quantity

$$X^2 = \sum_{k=1}^M -2 \ln L_k \tag{16.12}$$

⁷The χ^2 analytics degrade for small bin counts. For those bins, we merge any symbols whose expected count—for either set one or two—is less than five, reducing the degrees of freedom appropriately. If, after merging, there remain fewer than two symbols passing this criterion, then this node is wholly excluded.

⁸In this discrete case, the sum will encounter tables exactly as likely as the observed one (such as the observed table itself); the summed p_T for these tables is weighted by a uniform random deviate $r \in [0, 1)$.

is χ^2 distributed with $2M$ degrees of freedom, from which we compute our final \mathcal{L} , again uniform in $(0, 1)$ under the null. Especially small values of \mathcal{L} imply a small likelihood that this level of difference would have been observed had the two symbol datasets been generated by the same underlying dynamical process. This completes our desired test procedure. In the following numerical examples, we use the state-selecting algorithm descended from Rissanen’s method. We have not yet used the batch tree extraction method for this application yet.

We first test the accuracy of the statistic under the null. We produced an ensemble of 1000 time series from the x coordinate of the “Lorenz 84” attractor: a tiny geophysical model with attractor dimension $d \approx 2.5$ [18]. This system is higher dimensional and more complex (see Fig. 16.5) than




FIGURE 16.5. Sample orbit of Lorenz dynamical system in reconstructed state space.

the traditional Lorenz dataset, and is thus a somewhat more stringent test. Figure 16.6 shows the distribution of \mathcal{L} comparing the first and second halves of each set, demonstrating that \mathcal{L} is close to uniform $\in (0, 1)$. This is a stringent requirement and shows the success of our independence assumption, as it is difficult to get a high-quality null distribution with complicated arbitrarily correlated chaotic data in the null class. With this number of data, the test is also quite powerful.

We demonstrate discrimination power with a set of pressure data from an experimental model of a “fluidized bed reactor” [6]. This experimental system consists of a vertical cylindrical tube of granular particles excited from below by an externally input gaseous flow. In some regimes (“slugging”), the particles exhibit a combination of collective low-dimensional bulk dynamics and small-scale high-dimensional turbulence of the individual particles [6]. The observed variable was an azimuthally averaged pressure difference between two vertically separated taps. Figure 16.7 shows portions of time-delay embedding of orbits sections of the dataset taken




FIGURE 16.6. Quantile-quantile plot of \mathcal{L} under the null hypothesis. The observed values of \mathcal{L} are sorted and plotted vs their normalized index $(i+1)/1001$. Asymptotically the curve should approach the diagonal under the null. Bars are \pm two standard deviations for 100 samples of 1000 uniform deviates $\in [0, 1]$ processed similarly.




FIGURE 16.7. Phase space plots of the differential pressure signal from a fluidized bed reactor. Three are from the same parameters; one is different.

at the same experimental parameters, and one when the flow was boosted by 5%. The change in the attractor is rather subtle and difficult to reliably diagnose by eye. Figure 16.8 shows \mathcal{L} on a data set whose flow was




FIGURE 16.8. Non-stationary fluidized bed results with air flow altered at the 50% mark. Plotted statistic \mathcal{L} as a function of hypothesized breakpoint in time series and symbolic alphabet precision. Results for $|A| > 2$ numerically underflowed to $\mathcal{L} = 0$ toward the center and are not plotted. Null hypothesis emphatically rejected on account of the very small values of \mathcal{L} .

increased at the midpoint. As the alphabet size increased and the hypothesized breakpoint approached the true value of 50%, the strength of the rejection increased, $\mathcal{L} \rightarrow 0$. Even the binary alphabet case showed a significant rejection of the null. On data taken in stationary conditions \mathcal{L} fluctuates randomly in $(0, 1)$, as expected. We performed the same statistical test, with qualitatively identical results, on an experimental system whose flow rate was ramped slowly by the same amount. Even with this sort of data, the greatest rejection tends to occur in the middle of the dataset because the statistical discrimination power is greatest when there are equal numbers in the two sets considered in the test. We saw qualitatively identical results on a dataset whose flow was adiabatically boosted by the same degree during the run. This is not surprising, even though the test assumes a discrete breakpoint in order to lump the symbols into one class or another, as even given a smooth change, front and back sets will have different characteristics, and the ability to detect this (strongest rejection) will peak with approximately equal quantity of data in each set. Using a break-point test on smoothly changing dynamics might result in a small loss of statistical power compared to an ideal test, but in our experience with the proposed method on experimental datasets of at least one thousand points, adequate power to detect physically significant non-stationarity is rarely a concern with this method in our experience.

The Southern Oscillation Index, the normalized pressure difference be-

tween Tahiti and Darwin, is a proxy for the El Niño Southern Oscillation, as ocean temperature influences atmospheric dynamics. The period from mid-1990 to 1995 exhibited an anomalously sustained period of El Niño-like conditions (Fig. 16.9), perhaps indicative of global climate change. One




FIGURE 16.9. Three month moving average of the Southern Oscillation Index, the normalized pressure difference between Tahiti and Darwin, Australia. Strongly negative values correspond to El Niño events. Is the extended negative period from mid-1990 through 1995 especially anomalous?

statistical analysis [31] found such an anomaly quite unlikely assuming stationarity, but another group [11] used a different analysis and found it significantly more likely to be a chance fluctuation. Both papers used traditional linear forecasting models, with the difference centered around an auto-correlation-based correction for serial correlation to arbitrarily reduce the degrees of freedom. We applied our algorithm to the three-month moving average SOI (binary symbolized) testing the 5.4-year period in question against the rest of the series (starting from 1900), with a resulting $\mathcal{L} \approx 0.01$, meaning that one would expect to see a region this anomalous by chance every 540 years. The result is closer to those of [11] than [31] but we certainly do not want to take any particular position regarding climate; rather, we wish to point out an application for our method where correcting for serial correlation automatically is useful.

We point out that the proposed method is not exclusively limited to testing or finding a single break point—all that is needed is a sensible *a priori* hypothesized division of the dataset into discrete multiple classes. For instance, one might want to test for the presence of cyclo-stationarity, that the dynamics are externally modulated at some slow frequency Ω . In this case, one could choose elements of set 1 and set 2 depending on whether $\sin(\Omega t + \theta)$ is positive or negative, given fixed Ω and θ . Here the hypothesis is that the dynamics are significantly different when trying to

predict symbols on one part of the cycle compared to the other. Testing against three or more classes would require an upgrade of the χ^2 or Fisher test of proportions procedures; there exist conventional methods in the statistical literature.

Recent work has successfully used a distance in the symbolic space to fit unknown parameters of a physically motivated continuous model to observed data, including substantial observational and dynamic noise all in one framework. Tang et al [30] first proposed minimizing over free parameters the difference between an observed distribution of symbol words and that produced by discretizing some proposed model's output. Daw et al [5] successfully used this technique to fit experimental internal combustion engine measurements to a low-dimensional dynamical model. The optimization target was a Euclidean norm in [30] and a chi-squared distance in [5]. Because of serial correlation, a true hypothesis test confirming the apparent compatibility of observed data to a well-fitting model was not possible in those works. We feel our current method could provide a less ad hoc optimization goal, e.g. maximizing average \mathcal{L} or minimizing the code length of the physical model's output, encoded using the symbolic model learned from the observed data.

16.7 Conclusions

We have shown how to use a simple and rapid algorithm to model the dynamics of symbolic systems. The context tree weighting algorithm, and the proof that it is asymptotically optimal and has good finite-time properties, is well-known to workers in information technology but less so to workers in dynamics. The state selection algorithm is known to have some small flaws by comparison, but is still a high-quality symbolic modeler. These modeling methods are fast, deterministic, and mathematically sound, requiring no iterative fitting or ad-hoc convergence tests, as is common in the "neural network" regression community.

Specific contributions of this chapter, besides pointing out the usefulness to dynamics of the information-theoretic work, are that we have demonstrated a new method for estimating invariant measure and topological entropy, two of the most fundamental quantities in nonlinear dynamics, as well as a test for dynamical stationarity.

Acknowledgments

Both authors thank the Australian Research Council for supporting this research and the Isaac Newton Institute for an outstanding working environment. AIM thanks the Department of Systems Engineering and Engineering Management at The Chinese University of Hong Kong for hospitality.

References

- [1] H. D. I. Abarbanel. *Analysis of Observed Chaotic Data*. Institute for Nonlinear Science. Springer, New York, 1996.
- [2] R. Brown, V. In, and E.R. Tracy. Parameter uncertainties in models of equivariant dynamical systems. *Physica D*, 102(3-4):208–226, 1997.
- [3] M. Casdagli. Nonlinear prediction of chaotic time series. *Physica D*, 35(3):335–356, 1989.
- [4] T. Cover and J. Thomas. *Elements of Information Theory*. Wiley Interscience, New York, 1991.
- [5] C S Daw, M B Kennel, C E A Finney, and F T Connolly. Observing and modeling nonlinear dynamics in an internal combustion engine. *Phys. Rev. E*, 57(3-A):2811–2819, 1998.
- [6] C.S. Daw, Finney C.E.A., Vasudevan M., N.A. van Goor, K. Nguyen, D.D. Bruns, E.J. Kostelich, C. Grebogi, Ott E., and J.A. Yorke. Self-organization and chaos in a fluidized bed. *Phys. Rev. Lett.*, (75):2308–2311, 1995.
- [7] J. D. Farmer and J. J. Sidorowich. Predicting chaotic time series. *Phys. Rev. Letters*, 59(8):845–848, 1987.
- [8] A. Fraser. Information and entropy in strange attractors. *IEEE Trans. Information Theory*, 35:245–262, 1989.
- [9] A. M. Fraser and H. L. Swinney. Independent coordinates for strange attractors from mutual information. *Physical Review A*, 33(2):1134–1140, 1986.
- [10] G. Froyland, K. Judd, A. I. Mees, K. Murao, and D. Watson. Constructing invariant measures from data. *International Journal of Bifurcation and Chaos*, 5(4):1181–1192, 1995.
- [11] D.E. Harrison and N.K. Larkin. Darwin sea level pressure, 1876–1996: evidence for climate change? *Geophys. Res. Lett.*, 24:1779–1782, 1997.
- [12] K. Judd and A. I. Mees. On selecting models for nonlinear time series. *Physica D*, 82:426–444, 1995.
- [13] K. Judd and A. I. Mees. Embedding as a modeling problem. *Physica D*, 120:273–286, 1998.
- [14] M. B. Kennel, R. Brown, and H. D. I. Abarbanel. Determining embedding dimension for phase-space reconstruction using a geometrical construction. *Physical Review A*, 45(6):3403–3411, 1992.
- [15] M.B. Kennel. Statistical test for dynamical nonstationarity in observed time-series data. *Phys. Rev. E*, 78:316, 1997.
- [16] M. Li and P.M.B. Vitanyi. *An Introduction to Kolmogorov Complexity and its Applications*. Springer, New York, 2nd edition, 1997.
- [17] D. Lind and B. Marcus. *Symbolic Dynamics and Coding*. Cambridge University Press, 1995.
- [18] E.N. Lorenz. *Tellus*, 36A:98–110, 1984. The model is $dx/dt = -y^2 - z^2 - a(x - F)$, $dy/dt = xy - bxz - y + 1$, $dz/dt = bxy + xz - z$, $a = 1/4$, $b = 4$, $F = 8$. Each set was 5000 points long sampled every $\delta t = 0.08$.
- [19] A. I. Mees. Modelling complex systems. In T. Vincent, A. I. Mees, and L. S. Jennings, editors, *Dynamics of Complex Interconnected Biological Systems*, volume 6 of *Mathematical Modeling*, pages 104–124. Birkhauser, Boston, 1990.

- [20] A. I. Mees. Dynamical systems and tessellations: Detecting determinism in data. *International Journal of Bifurcation and Chaos*, 1(4):777–794, 1991.
- [21] A. I. Mees. Parsimonious dynamical reconstruction. *International Journal of Bifurcation and Chaos*, 3(3):669–675, 1993.
- [22] B. Pilgram, K. Judd, and A. I. Mees. Modelling the dynamics of nonlinear time series using canonical variate analysis. *Physica D*, in press, 2000.
- [23] J. Rissanen. Universal coding, information, prediction and estimation. *IEEE Trans. Inf. Theory*, IT-30(4):629–636, 1984.
- [24] J. Rissanen. *Stochastic Complexity in Statistical Inquiry*, volume 15 of *Series in Computer Science*. World Scientific, Singapore, 1989.
- [25] T. Schreiber. Detecting and analyzing nonstationarity in a time series using nonlinear cross predictions. *Phys. Rev. Lett.*, 78:843, 1997.
- [26] Y.M. Shtarkov, T.J. Tjalkens, and F.M.J. Willems. Multialphabet weighting universal coding of context tree sources. *Problems of Information Transmission*, 33(1):17–28, 1997.
- [27] J. Stark, D.S. Broomhead, M.E. Davies, and J. Huke. Takens embedding theorems for forced and stochastic systems. *Nonlinear Analysis*, 30:5303–5314, 1997.
- [28] G. Sugihara and R. M. May. Nonlinear forecasting as a way of distinguishing chaos from measurement error in time series. *Nature*, 344:734–741, 1990.
- [29] F. Takens. Detecting strange attractors in turbulence. In D. A. Rand and L. S. Young, editors, *Dynamical Systems and Turbulence*, volume 898 of *Lecture Notes in Mathematics*, pages 365–381. Springer, Berlin, 1981.
- [30] X.Z. Tang, E.R. Tracy, Boozer A.D., A. deBrauw, and R. Brown. Symbol sequence statistics in noisy chaotic signal reconstruction. *Phys. Rev. E*, 51:3871, 1995.
- [31] K.E. Trenberth and T.J. Hoar. The 1990–1995 el nino-southern oscillation event: longest on record. *Geophys. Res. Lett.*, 23:57–60, 1996.
- [32] A. S. Weigend, B. A. Huberman, and D. E. Rumelhart. Predicting the future: A connectionist approach. *International Journal of Neural Systems*, Vol 1(No 3):193–209, 1990.
- [33] F.M.J. Willems. The context tree weighting method: Extensions. *IEEE Trans IT*, 44(2):792–798, 1998.
- [34] F.M.J. Willems, Y.M. Shtarkov, and T.J. Tjalkens. The context tree weighting method: Basic properties. *IEEE Trans IT*, 41(2):653–664, 1995.
- [35] A. Witt, J. Kurths, and A. Pikovsky. Testing stationarity in time series. *Phys. Rev. E*, 58:1800, 1998.
- [36] L.-S. Young. Entropy, Lyapunov exponents, and Hausdorff dimension in differentiable dynamical systems. *IEEE Transactions on Circuits and Systems*, CAS-30(8):599–607, 1983.

Chapter 17

Analyzing Nonlinear Dynamical Systems with Nonparametric Regression

Henning U. Voss

ABSTRACT The analysis of dynamical systems data often can be considerably simplified using some knowledge of the system's structure rather than performing a general phase space reconstruction. For the common case when the evolution equations are given by a sum of functions of measurements, the statistical problem of model estimation is reduced from a multidimensional density estimation problem to several two-dimensional problems, connected by an in general nonlinear relationship. To recover this relationship, we use the statistical approach of nonparametric nonlinear regression analysis. This allows (a) for the analysis of systems with high-dimensional dynamics, like spatially extended and time-delayed feedback systems, and (b) for the further investigation of the resulting models. To illustrate these points, we review the application of nonparametric regression analysis to two physical experiments and numerical examples of nonlinear dynamics.

17.1 Introduction

Many methods for the analysis of data from nonlinear dynamical systems can be traced back mainly to the proper estimation of a multidimensional probability density function, e.g., in phase space reconstruction and nonlinear forecasting. The drawback of most of these methods is that already for small dimensions probability density functions are hard to estimate, i.e., the results have a large uncertainty. On the contrary, regression analysis is based on two-dimensional probability density estimations only, even in multivariate problems. The price one has to pay is that one is somewhat more restricted in the choice of a model, because one cannot rely on reconstruction theorems, and one needs to have some prior knowledge of the system's structure.

However, as it will be shown, for a broad class of nonlinear dynamical systems this poses no real restriction to an analysis. Using nonparametric regression, it turns out that the models often can be estimated from very limited amounts of data which allows for the analysis of spatially extended systems and systems with a time delayed feedback which may possess a

high-dimensional state space.

Rather than performing a least squares calculation in coefficient space in parametric regression analysis, in the nonparametric approach one performs an optimization in function space, and the results are non-parametrically given functions. These can be further investigated for the search of possible explanations of the observed dynamics.

In this chapter, these points will be elaborated on several experimental and numerical examples. For model selection, we use two distinct approaches: The statistical concept of maximal correlation and the elimination of redundant functions appearing in the results.

After an introduction to nonparametric nonlinear regression analysis and the notion of maximal correlation in the second section these concepts are applied to model time-delayed feedback systems in the third section. In the fourth section the analysis of a passive laser resonator by fitting a map to a time series of experimental measurements is illustrated. The fifth section concerns the analysis of spatially extended systems. It will be applied to experimental data from a fluid convection experiment in the sixth section, and, finally, the results are discussed.

17.2 Nonlinear Regression Analysis and Maximal Correlation

In dynamical systems modeling a common task is to find a model

$$\mathbf{x}_{t+1} = \Phi(\mathbf{x}_t) \quad (17.1)$$

based on a time series y_t ($t = 1, \dots, N$), where $x_{t+1} \in \mathbb{R}$, $\mathbf{x}_t \in \mathbb{R}^K$, for example, an embedding vector $(x_t, x_{t-\tau}, \dots, x_{t-(K-1)\tau})$, and $\Phi: \mathbb{R}^K \rightarrow \mathbb{R}$ a function to be estimated from the time series. To let (17.1) be the model with the most optimal prediction performance, $\Phi(\mathbf{x}_t)$ has to be estimated from data as the conditional expectation value of x_{t+1} given \mathbf{x}_t , $E[x_{t+1}|\mathbf{x}_t]$. Therefore,

$$\Phi(\mathbf{x}) = E[x|\mathbf{x}] = \int p(x|\mathbf{x}) x dx, \quad (17.2)$$

where $p(\cdot)$ is a conditional probability density function (PDF) and the time index has been dropped; due to the relation $p(x|\mathbf{x}) = p(x, \mathbf{x})/p(\mathbf{x})$, the estimation of $\Phi(\mathbf{x})$ becomes equivalent to the estimation of a $K + 1$ -dimensional PDF. Already for small K this is generally a problematic task given only a finite amount of data [31], known sometimes as the ‘‘curse of dimensionality.’’

For many problems, however, it is not necessary to solve problem (17.2) in its most general form; the inclusion of prior knowledge about $p(x, \mathbf{x})$ can help to simplify the problem considerably. In this contribution it will be

especially assumed that the model can be written as a linear combination of transformations $\Phi_i(x_i)$, i.e.,

$$\Phi_0(x_0) = \sum_{i=1}^K \Phi_i(x_i), \tag{17.3}$$

with $\Phi_i : \mathbb{R} \rightarrow \mathbb{R}$. As will be shown in the examples, this ansatz applies for a broad class of problems, including spatiotemporal dynamics.

To estimate model (17.3) from data we use nonlinear nonparametric regression analysis. In this case, the functions Φ_i can be estimated based on the estimation of only two-dimensional PDFs.

Let us start with $K = 1$, where the two functions Φ_0 and Φ_1 for the model $\Phi_0(x_0) = \Phi_1(x_1)$ are searched for. To estimate these, one has to minimize the expression $E[(\Phi_0 - \Phi_1)^2]$. To exclude the trivial solution, the constraint of a finite variance of at least one of the functions is added, say, $D^2[\Phi_0] = 1$, where $D[\cdot]$ is the standard deviation. The solution of this estimation problem for Φ_0 and Φ_1 becomes equivalent to the maximization of the correlation coefficient between Φ_0 and Φ_1 [3]; the linear correlation coefficient R is defined by

$$R(x_0, x_1) = (E[x_0 x_1] - E[x_0] E[x_1]) / D[x_0] D[x_1],$$

the quantity

$$\Psi(x_0, x_1) = \sup_{\Phi_0, \Phi_1} |R(\Phi_0(x_0), \Phi_1(x_1))|$$

is called *maximal correlation* [25] between the two random variables x_0 and x_1 , and the solutions Φ_0 and Φ_1 for which the supremum is attained are called *optimal transformations*. By definition, Ψ is restricted to the interval $[0, 1]$. The maximal correlation has the important property that it is a statistical measure that captures *any* dependence between the variables x_0 and x_1 . For the case of complete dependence, the maximal correlation attains unity.

Hence, the main idea is to transform x_0 and x_1 by suitable, generally nonlinear, transformations to a linear relationship between the new random variables $\Phi_0(x_0)$ and $\Phi_1(x_1)$. In case of a linear dependence between x_0 and x_1 , the optimal transformations are linear functions, too. It has been shown that optimal transformations do exist quite generally [3].

Here, a generalization of the notion of maximal correlation for multivariate regression problems will be used, i.e.,

$$\Psi(x_0, \dots, x_K) = \sup_{\Phi_0, \dots, \Phi_K} \left| R \left(\Phi_0(x_0), \sum_{i=1}^K \Phi_i(x_i) \right) \right|. \tag{17.4}$$

The problem of calculating the maximal correlation (17.4) is equivalent to estimating optimal transformations from data; the optimal transformations

can numerically be estimated in a nonparametric way from data by the *alternating conditional expectation algorithm* (ACE), invented by Breiman and Friedman [3]¹.

In the ACE-algorithm, which is an iterative algorithm, for $K = 1$ at each iteration step the quantities $E[\Phi_0(x_0)|x_1] = \int p(x_0|x_1) \Phi_0(x_0) dx_0$ and $E[\Phi_1(x_1)|x_0] = \int p(x_1|x_0) \Phi_1(x_1) dx_1$ are non-parametrically estimated, thus involving only two-dimensional PDFs. This also holds for the multivariate case with $K \geq 2$. Therefore, “in general, the problem of nonparametric regression is much easier than the problem of nonparametric density estimation” [31].

There are several possibilities to estimate conditional expectations from finite data sets. Here we use an efficient algorithm, in which the data are rank-ordered before the optimal transformations are estimated. This makes the result more insensitive to the data distribution which is often rather inhomogeneous. By definition, the maximal correlation is not affected by these monotonous transformations. The conditional expectation values are then estimated with local boxcar smoothing.

In the following chapters this concept is applied to four different modeling problems, chosen to emphasize the different possibilities of application. Especially, the maximal correlation will be used for model selection, and we put stress on the applicability to real-world measurements.

17.3 High-Dimensional Dynamics (I): Analysis of Delayed-Feedback Systems

Time-delay induced dynamics, as described by univariate delay differential equations (DDEs), play an important role in modeling natural phenomena. Such models are used in many different scientific disciplines (see [13, 15, 9] and references therein.)

Dynamical systems given by a DDE have an infinite-dimensional state space, and attractors of the solutions can also be high-dimensional [11, 10]. By evaluating a method for the identification of such systems, given only a single and finite solution, we demonstrate that these properties do not constitute severe limitations for a nonparametric regression analysis. (Another approach to this problem was first given by Büchner et al. [4].) Besides the estimation of the DDE, the method allows us to judge whether the data are in agreement with a description by a nonlinear feedback model. The DDE is estimated in a nonparametric way, such that the method covers a wide class of models.

¹For a brief description of the algorithm see also [16], [33] or [34]. A C-program is available from the author or by <http://www.agnld.uni-potsdam.de/~hv/ace.c>.

We consider the reconstruction of DDEs of the form

$$h(\dot{x}(t)) = f(x(t)) + g(x(t - \tau)), \quad (17.5)$$

where the functions f, g, h and the delay time τ are unknown. The functions are assumed to be continuous; usually, h is the identity. (For the estimation of models with more than one delay time, see [33].)

In spite of that the dynamics of equation (17.5) has to be described in an infinite-dimensional phase space, the evolution of the triple $(\dot{x}(t), x(t), x(t - \tau))$ is restricted to a three-dimensional manifold. This is also reflected in each finite realization of the DDE, y_t ($t = 1, \dots, N$), with an arbitrary initial condition: The evolution of the triple $(\Delta y_t, y_t, y_{t-\tau})$, where $\Delta y_t = (y_{t+1} - y_{t-1})/2$ is an estimate of the time derivative, is also restricted approximately to the same three-dimensional manifold via

$$\tilde{h}(\Delta y_t) = \tilde{f}(y_t) + \tilde{g}(y_{t-\tau}). \quad (17.6)$$

The functions \tilde{f}, \tilde{g} and \tilde{h} are then approximations of f, g and h obtained from a finite realization of the process. Equation (17.6) expresses the inverse problem to equation (17.5), i.e., finding a DDE from a finite solution.

The triple $(\Delta y_t, y_t, y_{t-\tau})$ can be seen as a three-dimensional embedding vector where the first component comes from a differential-embedding and the second and third components come from a delay-embedding. Therefore, for the reconstruction of a univariate DDE (supposing, the measurement function is one-to-one), such a three-dimensional embedding always suffices, no matter how large the attractor dimension actually is.

To estimate the functions in equation (17.6), we use nonparametric regression analysis: If one solves equation (17.4) for $x_0 = \Delta y_t$, $x_1 = y_t$ and $x_2 = y_{t-\tau}$, the estimates for Φ_0 , Φ_1 and Φ_2 can serve as estimates for the functions \tilde{f}, \tilde{g} and \tilde{h} in equation (17.6). As a quantitative criterion for how well the estimated functions fit equation (17.6), the maximal correlation is used:

$$\Psi(\tau) = |R(h(\Delta y_t), f(y_t) + g(y_{t-\tau}))|. \quad (17.7)$$

Now, if the maximum of $\Psi(\tau)$ is close to one, its location points to the estimate for τ . If none of the $\Psi(\tau)$ appears to be close to one, the time series is not likely to be the solution of a DDE of the form (17.5). What “close to one” means exactly depends of course on subjective judgment, taking all circumstances into account — like the length of the time series, the sampling rate, the quality of measurement and the used ACE implementation.

As an example, we consider chaotic data produced by numerical integration of the passive optical resonator system of Ikeda and Matsumoto [19],

$$\dot{x}(t) = -x(t) + \mu \sin(x(t - \tau) - x_0). \quad (17.8)$$

The parameters are chosen to be $\mu = 20$, $\tau = 2$ and $x_0 = \pi/3$. For $\mu \gg 1$, as in this case, the data are high-dimensionally chaotic [19]. The data set




FIGURE 17.1. The Ikeda model (17.8): (a) the data sample, (b) the ACF, (c) the regression plot as explained in the text, and (d) the delay estimation.

consists of 2000 values (Fig. 17.1a) which were sampled with a time step of 0.02. The small delay time leads to a mixing of the local correlations, corresponding to the damping term on the right hand side of equation (17.8) and the nonlocal, delay-induced correlations, corresponding to the sinusoidal feedback term. This can be seen in the autocorrelation function (ACF), which is dominated by the damping term but not by the delayed-feedback term (Fig. 17.1b).

Performing a nonparametric regression analysis on the variables Δy_t , y_t and $y_{t-\tau}$ using the ACE-algorithm to solve equation (17.6), one obtains a clear peak in the maximal correlation $\Psi(\tau)$ at the correct delay of $\tau = 2.00$ (Fig. 17.1d). Figures 17.2a–c show the estimates for the three terms of equation (17.5) retrieving clearly the identity as the left-hand side of the Ikeda model, the linear damping term and the nonlinearity.

It can be seen that the original, very complicated relationship between the three variables Δy_t , y_t and $y_{t-\tau}$ (Fig. 17.1c) has been changed to a linear relation between the transformed variables $h(\Delta y_t)$, $f(y_t)$ and $g(y_{t-\tau})$ (Fig. 17.2d). The remaining scatter is mainly due to an inaccurate estimation of the time derivatives Δy_t from the data.

To summarize, both the delay time and the DDE have been estimated with high accuracy from a short time series; the system has been completely reconstructed. Since the optimal transformations are estimated in

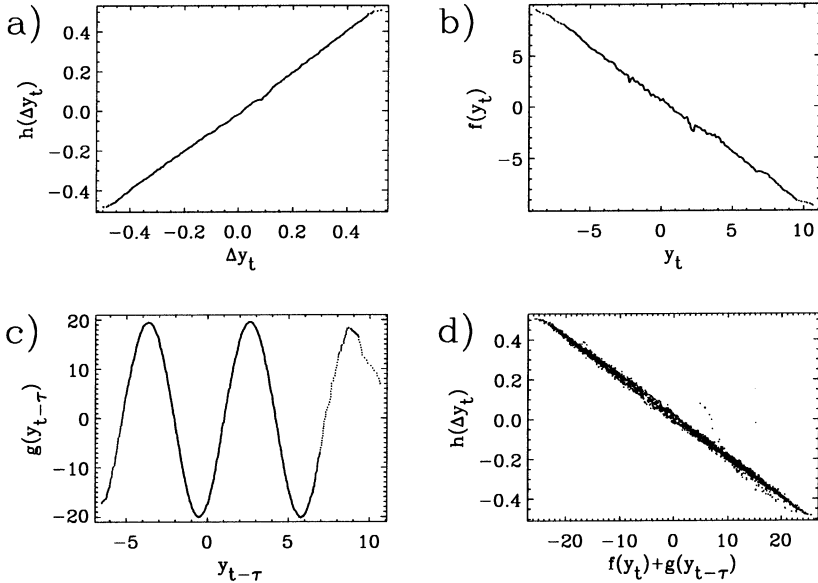


FIGURE 17.2. The estimated optimal transformations for the l.h.s (a), the damping term (b) and the nonlinear feedback term (c) for $\tau = 2.00$. In (d) a regression plot of the transformed variables shows that the complicated statistical dependence in the time series (Fig. 17.1c) has almost completely been transformed to a linear dependence.

a nonparametric way, even the very strong sinusoidal nonlinearity could be approximated well. This would not have been possible by an ansatz in which the feedback function is given by a parameterized Taylor expansion.

The function estimates should be reliable provided that the data cover enough space of the manifold described by the DDE. This is not a restriction of the method itself but a general limitation due to the inverse nature of the problem. In the case of high-dimensional chaotic data this requirement should always be fulfilled, if a sufficient amount of data is given. For the behavior for noisy and very short time series, see [33] and [34].

17.4 Interpretation of Models (I): Analysis of a Fiber Ring Resonator System

The explicitly given model (17.3) allows for an investigation of itself because the functions Φ_0, \dots, Φ_K are known as the result of the regression analysis, whereas the model (17.1) in its most general form has to be considered to be a “black-box” model. The main reasons are that the embedding of the time series usually leads to a complicated nonlinear transformation of the

phase space, and multidimensional functions are difficult to visualize.

However, in this section it is shown on experimental data of a system from nonlinear optics that using nonparametric regression analysis, it may still be possible to give the model a physical interpretation.

The experiment consists of a passive optical fiber ring resonator containing 10 m of single mode optical fiber and a beam splitter to couple light in and out. This resonator is driven by a stream of pulses, about 1 ps wide and spaced out by $\Delta t = 12$ ns, from a color center laser which in turn is synchronously pumped from a Nd:YAG laser. As light pulses travel around the ring cavity, they experience a phase shift $\Theta(t) \propto n(t)$. The effective refractive index of the fiber $n(t)$ is given as $n(t) = n_0 + n_2 I(t)$. The first term n_0 denotes the familiar classical refractive index, and the term $n_2 I(t)$ gives rise to the intensity-dependent phase known as the optical Kerr effect, with $I(t)$ the temporal intensity profile of the pulses, and n_2 the Kerr coefficient. The round trip time is made synchronous with the laser pulse repetition time Δt so that after each round trip the pulse circulating in the cavity can interfere with the next incoming pulse. The combination of the stationary and the nonlinear cavity phase leads to complex modifications of the pulse structure on interference and can lead to chaotic dynamics for some parameter regimes. For more details of the experiment, see [29].

We analyze a time series of 6000 points of the light intensity. Return maps of the data are shown in Figs. 17.3a,b. Fitting models

$$I_{t+1} = \Phi_0^{-1} \left(\sum_{i=1}^K \Phi_i(I_{t+1-i}) \right) \quad (17.9)$$

with increasing order $K = 1 \dots 10$ yields, as expected, a monotonic increase of the maximal correlation that saturates at about $K = 4$ (Fig. 17.4). Since the data are rather noisy, the maximal correlation remains somewhat smaller than unity. Truncating the model (17.9) at $K = 4$, one yields the optimal transformations as depicted in Figure 17.5.

Next, the model given by the functions in Figure 17.5 is numerically iterated in a free run to compare its dynamics with the measurements. As a result, the over-all shape of the attractor is well reproduced by the model data (Figs. 17.3c,d). However, one also observes a fine-structure of the attractor that is not visible in the original measurements. Since model (17.9) does not include noise, the measurement noise in the data cannot be reproduced by this approach, and the modeling result is expected to contain a lower amount of noise. The fine-structure of the attractor is also revealed after application of a simple nonlinear noise reduction scheme [26] to the data based on local averaging in embedding space (Figs. 17.3e,f).

Now we come to the interpretation of the estimated model functions, the optimal transformations. It has been shown [29] that for certain parameter regimes (which we consider here), the dynamics can be well approximated by the Ikeda equation [18]. The Ikeda equation is given as a map for the

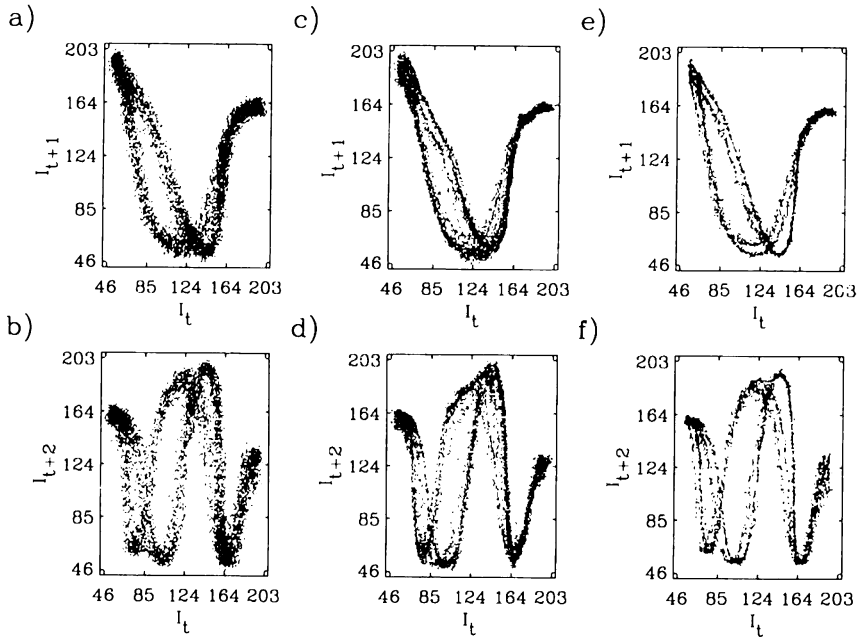


FIGURE 17.3. The first and second return maps of the measured laser light intensities (a,b), the model data for a model order of $K = 4$ in Eq. (17.9) (c,d), and the noise-reduced measured laser light intensities (e,f). Since for the nonlinear noise reduction the embedding space must have an odd dimension, here a five-dimensional embedding has been used. (From [36].)




FIGURE 17.4. The maximal correlation Ψ for model (17.9) with order K from one to ten. (From [36].)

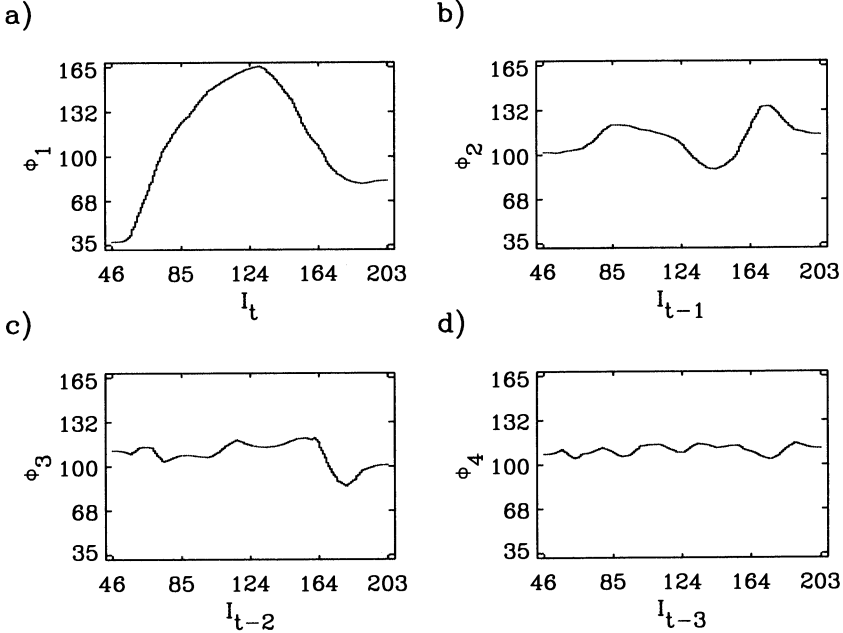


FIGURE 17.5. The optimal transformations $\Phi_1(I_t)$ to $\Phi_4(I_{t-3})$. (From [36].)

complex field amplitude E_t in the cavity after the t -th round trip,

$$E_{t+1} = aE_0 + be^{i(\Theta_0+c|E_t|^2)} E_t , \tag{17.10}$$

where a, b and c are real coefficients. The measured data are the intensities $I_t = |E_t|^2$. Therefore, we transform the complex Ikeda equation (17.10) to a model for the intensities. That means, instead of the Ikeda equation for the complex field amplitudes,

$$E_{t+1} = f(E_t) , \quad f : \mathbb{C} \rightarrow \mathbb{C} , \tag{17.11}$$

here we will use a map for the real intensities,

$$I_{t+1} = g_1(I_t) + g_2(I_{t-1}) + \dots + g_K(I_{t-K+1}) , \tag{17.12}$$

where $g_1, \dots, g_K : \mathbb{R} \rightarrow \mathbb{R}$. It can be shown [34] that this map is

$$I_{t+1} = a^2 I_0 + b^2 I_t + 2a^2 b I_0 \cos(\Theta_0 + cI_t) + T_{t-1,t-2,\dots} , \tag{17.13}$$

where $T_{t-1,t-2,\dots}$ denotes further lagged terms. Truncating this infinite series at a finite order K , one can compare the estimated optimal transformations in Figure 17.5 with the functions appearing in this model for the intensities:




FIGURE 17.6. The zero-lag estimated optimal transformation $\Phi_1(I_t)$ (bold line) and a numerical fit of Eq. (17.14). The static phase is estimated to be $\Theta_0 = -0.95$. (From [36].)

The zero-lag function $\Phi_1(I_t)$, that can be seen as an estimate of the zero-lag function in equation (17.13), resembles a superposition of a linear and a cosine-function, as expected from equation (17.13). The remaining function estimates for the right hand side of equation (17.13) show a doubling of the number of maxima with increasing time lag, as is also expected from a further expansion of equation (17.13). Furthermore, the first function Φ_0 , corresponding to the estimate for the left-hand side of equation (17.13), is approximately the identity (not shown).

As already noted, the zero-lag function estimate looks like a superposition of a linear and a cosine term (Fig. 17.5a). Therefore, next we try to retrieve the not directly measurable static phase Θ_0 of the Ikeda model from a parameterized fit to the zero-lag optimal transformation. That is, a function

$$f(I_t) = a_0 + a_1 I_t + a_2 \cos(a_3 + a_4 I_t) \quad (17.14)$$

is adjusted to the optimal transformation $\Phi_1(I_t)$ by numerically estimating the coefficients a_0 — a_4 . The coefficient a_3 then yields the stationary phase Θ_0 . A comparison of the fitted function (17.14) to the zero-lag optimal transformation $\Phi_1(I_t)$ (Fig. 17.6) indicates that the optimal transformation can indeed be parameterized by a function of the form (17.14), at least up to minor imperfections. In [36] the analysis has been performed also for non-stationary data with a temporally linearly increasing phase, allowing for tracking the phase in the course of time.

The inverse problem to yield equations of motion from measured data is usually ill-posed [14] in the sense that the solution cannot be expected to be unique or stable against small perturbations in the data. Considering the problem of uniqueness, here a pragmatic approach has been used; the model has been validated by comparing the numerical output of the model with the data. The problem of stability has been addressed by the application to a data set with a slowly changing bifurcation parameter [36], observing that the result follows that change smoothly, although the attractor changes qualitatively at certain values of the bifurcation parameter.

To sum up our results, we have shown on experimental data that in contrast to modeling this dynamical system with equation (17.1), it has been advantageous to use a more restricted model of the form (17.3). In the considered example, it was possible to give the estimated dynamical model a physical interpretation, thereby verifying the validity of the Ikeda equation.

17.5 High-Dimensional Dynamics (II): Analysis of Spatially Extended Systems

The unstable dynamics observed in spatially extended systems attracted huge experimental and theoretical research activity in the last decades (see [5, 24] and references therein). Due to the technical development of high-precision measurement techniques, it is now possible to resolve measurements in space and time sufficiently to perform a quasi-continuous analysis of experimental data. Before this will be shown in the next section for a specific example, we illustrate how in the analysis of numerical data non-parametric regression analysis can be used for model selection. Whereas in the previous two sections the maximal correlation was used for this purpose, here we will stress the elimination of redundant terms.

We analyze data $v(\vec{x}, t)$ from the Swift-Hohenberg equation [30]:

$$\begin{aligned}\partial_t u &= [r - (\nabla^2 + k^2)^2] u - u^3 \\ &= (r - k^4)u - u^3 - 2k^2(\partial_{xx} + \partial_{yy})u \\ &\quad - (\partial_{xxxx} + \partial_{yyyy} + 2\partial_{xxyy})u.\end{aligned}\tag{17.15}$$

The parameters are $r = 0.1$ and $k = 1$. The global dynamics of the model can be derived from a potential, such that the asymptotic time dependence is trivial [24]. Therefore, a transient state is analyzed to have a sufficient variation in the time derivative. The field size is 100×100 points, i.e., the data set $v(\vec{x}, t)$ contains 3×10^4 values. The data for the central time point are shown in Fig. 17.7. Details about the numerical integration can be found in [32].

For spatially extended data the “natural” embedding is a differential embedding if the system is described by a partial differential equation. The differential operators can be estimated by symmetric differencing schemes, e.g., $\partial_t v(\vec{x}, t) \approx [v(\vec{x}, t + \Delta t) - v(\vec{x}, t - \Delta t)]/2\Delta t$, or by estimation in the frequency domain. In the former case, which is used in this section, to estimate the time derivatives of first order in each spatial data point, one needs at least three consecutive “pictures” of data.

Suppose one does not know anything about the origin of the data. Then, to identify the unknown system, one would use as many independent input variables for the nonparametric regression analysis as possible. Here, an

ansatz of type (17.3) with all non-mixed terms (like $\partial_x v$) up to fourth order in the spatial derivatives is used. Additionally, the three product terms $v\partial_x v$, $v\partial_y v$ and $\partial_x v\partial_y v$ are included:

$$\begin{aligned} \Phi_0(\partial_t v) = & \Phi_1(v) + \Phi_2(\partial_x v) + \Phi_3(\partial_y v) + \Phi_4(\partial_{xx} v) \\ & + \Phi_5(\partial_{xy} v) + \Phi_6(\partial_{yy} v) + \Phi_7(\partial_{xxx} v) + \cdots \\ & + \Phi_{10}(\partial_{yyy} v) + \Phi_{11}(\partial_{xxx} v) + \cdots + \Phi_{15}(\partial_{yyy} v) \\ & + \Phi_{16}(v\partial_x v) + \Phi_{17}(v\partial_y v) + \Phi_{18}(\partial_x v\partial_y v), \end{aligned} \quad (17.16)$$

where the twelve terms $\Phi_2(\partial_x v)$, $\Phi_3(\partial_y v)$, $\Phi_5(\partial_{xy} v)$, $\Phi_7(\partial_{xxx} v)$, $\Phi_8(\partial_{xxy} v)$, $\Phi_9(\partial_{xyy} v)$, $\Phi_{10}(\partial_{yyy} v)$, $\Phi_{12}(\partial_{xxy} v)$, $\Phi_{14}(\partial_{yyy} v)$, $\Phi_{16}(v\partial_x v)$, $\Phi_{17}(v\partial_y v)$ and $\Phi_{18}(\partial_x v\partial_y v)$ are statistically independent from the others; these should vanish as a result of the nonparametric regression analysis, though the corresponding arguments, of course, do not vanish. Note that all functional dependencies on single arguments (like v^α , $\alpha \in \mathbb{R}$) will be captured by the optimal transformations, making the ansatz (17.16) already a very general one. Comparing equation (17.15) with equation (17.16), one expects in particular the following for the solution of equation (17.3): Up to an arbitrary common factor, Φ_0 should be the identity, Φ_1 should be a polynomial of third order, and for $i = 4, 6, 11, 13, 15$ the Φ_i should be linear functions in their respective arguments. All other estimates should vanish. Furthermore, one expects that for the slopes of the linear functions, $\alpha_4 = \alpha_6 = \alpha_{13} = -2$ and $\alpha_{11} = \alpha_{15} = -1$.

Performing the nonparametric regression analysis, one finds a maximal correlation of 0.9993 and optimal transformations as shown in Figure 17.8. All functions approximate the expected shape and the terms that were expected to vanish are indeed very small compared to the others. This can be quantified using the relative variances D_i^2 of the estimated functions, $D_i^2 = D^2[\Phi_i]/D^2[\Phi_0]$ ($i = 1, \dots, K$). For the six non-redundant terms one yields a relative variance of 0.06 to 0.50 and for the others 8×10^{-6} to 3.8×10^{-5} , a difference of three orders of magnitude. Comparison of the slope of the linear functions yields the possibility of estimating parameters; we obtain $\alpha_4 = \alpha_6 = -1.9$, $\alpha_{13} = -2.0$ and $\alpha_{11} = \alpha_{15} = -1.0$, in good coincidence with the expected values in equation (17.15).

Finally, it is important to note that the estimation of the derivatives sensitively depends on the data quality, which has been more closely investigated in [32]. This can make an application to real-world data problematic, especially if higher-order derivatives, like in the example, have to be incorporated. Nevertheless, in the next section this method is successfully applied to experimental measurements.




FIGURE 17.7. The data sample $v(x, y, t_0)$ for the central time point t_0 encoded in gray values (small values dark).

17.6 Interpretation of Models (II): Analysis of a Convection Experiment

In this last example nonparametric regression analysis is applied to extract dynamical equations that describe an experiment on traveling-wave convection in a binary fluid. A fundamental problem in the study of spatially-extended dynamical systems is the quantitative comparison of experimental data with models based on partial differential equations. In the study of nonlinear pattern-forming systems, theoretical models usually take the form of amplitude equations, and comparisons between data and models can only be performed qualitatively. In the case of traveling-wave convection in binary fluids, the complex Ginzburg-Landau equation (CGLE) model has been derived directly from the Navier-Stokes equations (see references in [22]), and a quantitative comparison of this model with data is warranted.

We analyze data from an experiment on convection in an ethanol/water mixture in a long, narrow, annular container heated from below. The system can be considered approximately one-dimensional, with periodic boundary conditions. The convection pattern is visualized by a shadowgraph system which illuminates a circular array of photodiodes, whose signals are digitized to provide data for analysis. The bifurcation parameter ε is defined as the fractional distance above the temperature difference ΔT applied across the fluid layer at onset of convection. We analyse data obtained at seven different values of the bifurcation parameter $\varepsilon' = \varepsilon\tau_0^{-1} \times 10^3$ (scaled by the characteristic time τ_0 defined later): $\varepsilon' = 1.77, 4.22, 6.38, 9.32, 12.07, 14.03$ and 16.28 .

The first dynamical state observed above onset ($\varepsilon = 0$) consists of pairs of weakly-nonlinear wave packets which propagate around the system in opposite directions, referred to as “left” and “right” [21] (see Fig. 17.9). The left- and right-going complex wave amplitudes $A_L(x, t)$ and $A_R(x, t)$ are extracted from the shadowgraph data using complex demodulation [23].

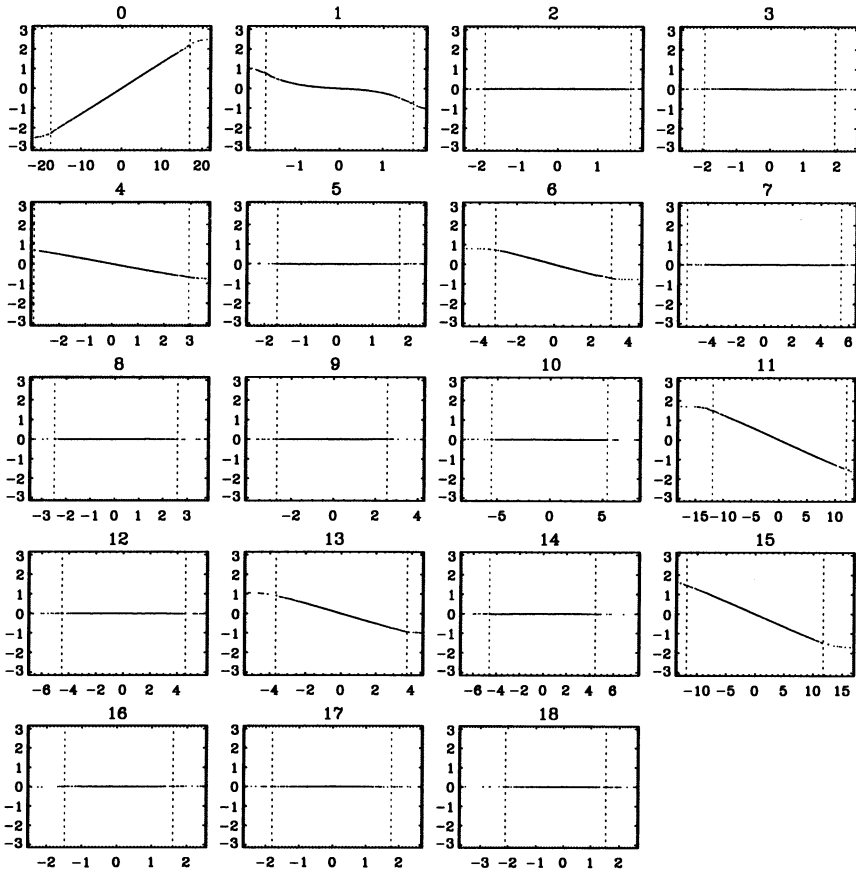


FIGURE 17.8. All estimates for the optimal transformations. As in the previous figures, the abscissae are the random variables estimated from the data, and the ordinates are the optimal transformation. The dotted lines mark the intervals on the abscissae in which 98% of the data points are located. Due to a very non-homogeneous distribution of the data, the optimal transformations outside the marked interval cannot be estimated reliably.




FIGURE 17.9. Space-time plots of the sum of the left- and right-going real amplitudes, given as gray values (small values dark) for a) $\varepsilon' = 1.77$ and b) $\varepsilon' = 12.07$. The gray-scale contrast in (a) has been magnified by a factor 2.5 with respect to that in (b) to compensate for weaker amplitudes. Since the amplitudes for $\varepsilon' = 12.07$ are about 2.5 times larger than the amplitudes for $\varepsilon' = 1.77$, in (b) the waves are more pronounced and the signal-to-noise ratio is higher. (From [35].)

In contrast to a reconstruction by embedding, this technique is based on a separation of different temporal and spatial scales.

The actual data fields we analyze are the real amplitudes and phases $a_L(x, t)$, $a_R(x, t)$, $\phi_L(x, t)$, and $\phi_R(x, t)$, defined by

$$A_{L,R}(x, t) = a_{L,R}(x, t)e^{i\phi_{L,R}(x,t)}. \quad (17.17)$$

These fields are sampled on a spacetime mesh of 180 spatial points by 760 time steps. For each value of ε' , the left- and right-wave fields are analyzed separately.

For the first dynamical states above onset, theory proposes as a quantitative description by amplitude equations two coupled CGLEs for left- and right-going traveling waves [21]. For ease of notation, we consider only right-going waves coupled to left-going ones and suppress the “R”-subscript. The equations for left-going waves follow symmetrically with a change of the sign of the velocity s from negative to positive values. The complex amplitude $A(x, t)$ is described by the CGLE

$$\begin{aligned} \tau_0 (\partial_t + s \partial_x) A &= \varepsilon(1 + ic_0) A + \xi_0^2(1 + ic_1) \partial_{xx} A & (17.18) \\ &+ g(1 + ic_2) |A|^2 A + h(1 + ic_3) |A_L|^2 A. \end{aligned}$$

The involved parameters are explained elsewhere [22]; τ_0 is a characteristic time which can be determined experimentally. Inserting Eq. (17.17), two

equations for real amplitude and phase are obtained:

$$\begin{aligned} \partial_t a &= -s \partial_x a + \varepsilon \tau_0^{-1} a \\ &\quad + \xi_0^2 \tau_0^{-1} \{ \partial_{xx} a - a(\partial_x \phi)^2 - c_1(2\partial_x a \partial_x \phi + a \partial_{xx} \phi) \} \\ &\quad + g \tau_0^{-1} a^3 + h \tau_0^{-1} a_L^2 a, \end{aligned} \quad (17.19)$$

$$\begin{aligned} a \partial_t \phi &= -s a \partial_x \phi + \varepsilon c_0 \tau_0^{-1} a \\ &\quad + \xi_0^2 \tau_0^{-1} \{ 2\partial_x a \partial_x \phi + a \partial_{xx} \phi + c_1(\partial_{xx} a - a(\partial_x \phi)^2) \} \\ &\quad + g c_2 \tau_0^{-1} a^3 + h c_3 \tau_0^{-1} a_L^2 a. \end{aligned} \quad (17.20)$$

First, the derivative fields $\partial_t a$, $\partial_x a$, $\partial_x \phi$, etc., that appear in Equations (17.19) and (17.20) are estimated from the experimental measurements, using spectral estimators. Then, these variables are taken as inputs for the nonparametric regression analysis.

Since one can apply the method to Equations (17.19) and (17.20) independently, we use two sets (**I** and **II**) of variables v_i ($i = 0, \dots, 5$) as input for the regression analysis. Set **I**, corresponding to equation (17.19): $v_0 = \partial_t a$, $v_1 = \partial_x a$, $v_2 = a$, $v_3 = \partial_{xx} a - a(\partial_x \phi)^2$, $v_4 = 2\partial_x a \partial_x \phi + a \partial_{xx} \phi$, $v_5 = a_L^2 a$, and set **II**, corresponding to equation (17.20): $v_0 = a \partial_t \phi$, $v_1 = a \partial_x \phi$, $v_2 = a$, $v_3 = 2\partial_x a \partial_x \phi + a \partial_{xx} \phi$, $v_4 = \partial_{xx} a - a(\partial_x \phi)^2$ and $v_5 = a_L^2 a$.

Numerical studies on several dynamical model equations [32] revealed that the CGLE could be estimated with high accuracy from noise-free data, leading to a maximal correlation of almost unity.

As a first question, we want to check that the spatiotemporal evolution of the system can be described by the coupled CGLEs (17.19,17.20). In this case, one expects the following optimal transformations: The function Φ_0 should be the identity, Φ_2 should be a third-order polynomial in a , and all the other functions should be linear, with slopes corresponding to the coefficients in Equations (17.19,17.20). As a check, we will compare our results with experimentally obtained coefficients from the same experiment as presented in [22]. There it was also shown that most of the experimental values agree reasonably well with the ones calculated from first principles. These experimental values are represented as smooth curves in Figure 17.10. Since the polynomials $\varepsilon \tau_0^{-1} a + g \tau_0^{-1} a^3$ and $\varepsilon c_0 \tau_0^{-1} a + g c_2 \tau_0^{-1} a^3$ have large uncertainties, curves representing their extremal values are shown in the upper and lower panels for Φ_2 , respectively. The distribution of the amplitudes, phases, and derivatives are rather inhomogeneous with heavy tails. Therefore, in Fig. 17.10 the range on the abscissa that is covered by 96 percent of the data values is marked by vertical dotted lines. Since the optimal transformations are harder to estimate for very sparse data, each 2 percent of the transformed data values at the edges are considered as outliers.

For the seven data sets that were analyzed, one obtains the following results (see Fig. 17.10):

For large bifurcation parameters ($\varepsilon' \geq 12.07$), the expected functions coincide very well with the coefficients found in [22]. In particular: Set **I**

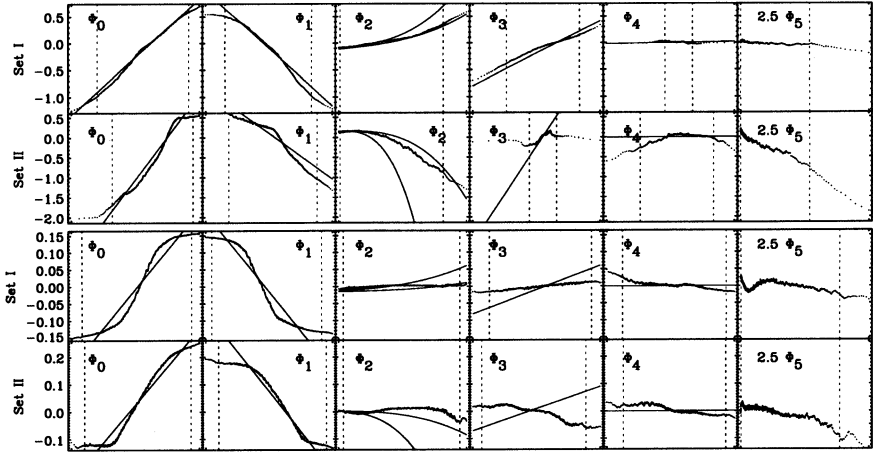


FIGURE 17.10. Estimated optimal transformations for the set of terms I and II, both for $\varepsilon' = 12.07$ (upper two rows) and for $\varepsilon' = 1.77$ (lower two rows). The ordinates are the optimal transformations multiplied by 1000. They are the same for all plots in one row, except in the frames for Φ_5 where they have been magnified by 2.5. The abscissae are given by the terms v_0 to v_5 , respectively, and are not labelled for clarity. Additionally, smooth curves indicate the theoretically expected functions, and vertical dotted lines mark the ranges on the abscissae where 96 percent of the data values are located, as explained in the text. The results for $\varepsilon' = 14.03$ and 16.28 resemble the results for $\varepsilon' = 12.07$ and are therefore not shown; similarly, the results for $\varepsilon' = 4.22, 6.38$ and 9.32 resemble the results for $\varepsilon' = 1.77$. (From [35].)

(top row of Fig. 17.10): The estimate for the left-hand side, Φ_0 , turns out to be approximately the identity; the estimate for Φ_1 is an approximately linear function in $\partial_x a$ with a slope in good agreement with the wave velocity s measured in [22]; and the estimate for Φ_2 can be described by a cubic polynomial in a ; the estimates for Φ_3 and Φ_4 are approximately linear, also with correct slopes. The estimate for the coupling term, Φ_5 , appears to be approximately linear in $a_1^2 a$ with a clearly negative coupling coefficient.

Set II (second row of Fig. 17.10) yields similar results, but obviously the estimates for Φ_3 and Φ_4 are worse. The reason is that due to the law of error propagation, the error of the phase ϕ is on average two orders of magnitude larger than the error of the amplitude, as explained in [35]. This leads to worse estimates for the derivatives and, consequently, to worse results for the optimal transformations of the set of terms II. (Note especially that Φ_3 depends on second-order derivatives of the phase.) This is also revealed in a smaller maximal correlation (17.4): $\Psi(\text{I}) = 0.985$ vs. $\Psi(\text{II}) = 0.945$.

For smaller bifurcation parameters ($\varepsilon' \leq 9.32$), the results are much worse (bottom rows of Fig. 17.10) and give only a rough impression of the true equations. Here, owing to smaller complex amplitudes, the signal-to-

noise ratio, given by the ratio of the standard deviations of signal and noise, changes by a factor of 2.7 as ε' is changed from 1.77 to 16.28. This leads again to bad estimates of the derivatives for small bifurcation parameters.

Finally, for the three data sets with highest signal-to-noise ratio, one clearly finds that the coupling term is approximately of the form $h(1 + ic_3)|A_L|^2 A$. This allows one to fit linear functions to the optimal transformations Φ_5 to yield estimates for the coupling coefficients: $h\tau_0^{-1} = -13.7 \pm 4.7$ and $hc_3\tau_0^{-1} = -142 \pm 15$. The negativity of the coupling coefficients is in agreement with the experimentally observed slight decrease in phase and group velocities measured during the interaction of the oppositely propagating wave packets [21].

To summarize, it has been shown from an analysis of the dynamics of a state of traveling-wave convection that it is possible to extract the governing amplitude equations from high-quality experimental spatiotemporal data, using nonparametric regression analysis. Limitations are mainly due to inaccurate estimates of spatial and temporal derivatives.

17.7 Conclusions

The practical analysis of complex nonlinear systems often requires solutions that are well-suited for certain problems. Especially for high-dimensional dynamics one has to find alternatives to common reconstruction methods (although considerable progress can be expected in this direction, as follows from recent results [27]). Often, even a well-performing phase-space model is not satisfying since it cannot be interpreted with respect to understanding the physics behind it.

In this chapter we have reviewed a method that stands in-between generally valid results, like reconstruction theorems, and specific problem solutions. On the one hand, nonparametric regression analysis is a very general approach of nonlinear statistics (which is still under development); on the other hand one does not have strong theorems that ensure the validity of the results with respect to their uniqueness.

Especially when the point is reached where the result of a nonparametric regression analysis is given a (physical) meaning, as we have done repeatedly in this chapter, one has to proceed with care; "As a practical matter, regression models, linear or nonlinear, should be considered exploratory devices, to be replaced as soon as possible by theory based models supported by consideration of the mechanism underlying the process" [31]. In our case, the ill-posedness of the inverse problem of determining a differential equation from data is, of course, of crucial importance. Nevertheless, we have given some examples of a useful application of nonparametric regression analysis, and the results were consistent with theory. Sometimes, like in the analysis of the fluid convection experiment, we clearly came close to

the limit of what is possible with today's measurement accuracy.

The applicability of nonparametric regression analysis has been elaborated on four examples chosen to highlight different aspects of this method as applied to nonlinear dynamical systems. There are other applications; the ACE algorithm in particular has been applied to analyze nonlinear diffusion processes [8, 20]. Monographs giving an overview of methods (but not directly considering nonlinear dynamical systems) are [16] and [31]. Properties of the maximal correlation, as compared to other measures of dependence, can be found in [1, 2, 6, 7, 12, 17] and [25]. The asymptotic distribution for a vanishing maximal correlation, which can be used for testing the hypothesis " $\Psi \neq 0$," has been derived by Sethuraman [28].

Acknowledgments

The author would like to thank M. Abel, M.J. Bünner, P. Kolodner, J. Kurths, F. Mitschke and A. Schwache for useful discussions and providing experimental data. Parts of this chapter have been adopted from joint publications with them. He acknowledges the hospitality of the Georgia Institute of Technology and the University of Surrey, and financial support from the Max-Planck-Gesellschaft.

References

- [1] C.B. Bell. Mutual information and maximal correlation as measures of dependence. *Ann. Math. Stat.*, **33**:587–595, 1962.
- [2] R.C. Bradley. Equivalent measures of dependence. *J. Multivariate Analysis*, **13**:167–176, 1983.
- [3] L. Breiman and J.H. Friedman. Estimating optimal transformations for multiple regression and correlation. *J. Am. Stat. Assoc.*, **80**:580–598, 1985.
- [4] M.J. Bünner et al. Tool to recover scalar time-delay systems from experimental time series. *Phys. Rev. E*, **54**:R3082–3085, 1996.
- [5] M.C. Cross and P.C. Hohenberg. Pattern formation outside of equilibrium. *Rev. Mod. Phys.*, **65**:851, 1993.
- [6] P. Csaki and J. Fischer. Contributions to the problem of maximal correlation. *Publ. Math. Inst. Hung. Acad. Sci.*, **5**:325–337, 1960.
- [7] P. Csaki and J. Fischer. On the general notion of maximal correlation. *Magyar Tud. Akad. Mat. Kutato Int. Kozl.*, **8**:27–51, 1963.
- [8] R.D. De Veaux and J.M. Steele. ACE guided-transformation method for estimation of the coefficient of soil-water diffusivity. *Technometrics*, **31**:91–98, 1989.
- [9] O. Diekmann et al. *Delay Equations*. Springer, New York, 1995.
- [10] B. Dorizzi et al. Statistics and dimension of chaos in differential delay systems. *Phys. Rev. A*, **35**:328–339, 1987.

- [11] J.D. Farmer. Chaotic attractors of an infinite dimensional dynamical system. *Physica D*, **4**:366–393, 1982.
- [12] H. Gebelein. Das statistische Problem der Korrelation als Variations- und Eigenwertproblem und sein Zusammenhang mit der Ausgleichsrechnung. *Zeitschrift für angew. Math. und Mech.*, **21**:364–379, 1941.
- [13] K. Gopalsami. *Stability and Oscillations in Delay Differential Equations of Population Dynamics*. Kluwer, Dordrecht, 1992.
- [14] C.W. Groetsch. *Inverse Problems in the Mathematical Sciences*. Vieweg, Braunschweig, 1993.
- [15] J.K. Hale and S.M. Verduyn Lunel. *Introduction to Functional Differential Equations*. Springer, New York, 1993.
- [16] W. Härdle. *Applied Nonparametric Regression*. Cambridge Univ. Press, Cambridge, 1990.
- [17] H.O. Hirschfeld. A connection between correlation and contingency. *Proc. of the Camb. Phil. Soc.*, **31**:520–524, 1935.
- [18] K. Ikeda, H. Daido, and O. Akimoto. Optical turbulence: Chaotic behavior of transmitted light from a ring cavity. *Phys. Rev. Lett.*, **45**:709–712, 1980.
- [19] K. Ikeda and K. Matsumoto. High-dimensional chaotic behavior in systems with time-delayed feedback. *Physica D*, **29**:223–235, 1987.
- [20] H.G. Kim and J.C. Lee. Critical heat flux correlation through the alternating conditional expectation algorithm. *Transactions of the American Nuclear Society*, **74**:161–163, 1996.
- [21] P. Kolodner. Counterpropagating quasilinear wave packets in binary-fluid convection. *Phys. Rev. Lett.*, **69**:2519, 1992.
- [22] P. Kolodner, S. Slimani, N. Aubry, and R. Lima. Characterization of dispersive chaos and related states of binary-fluid convection. *Physica D*, **85**:165–224, 1995.
- [23] P. Kolodner and H. Williams. Complex demodulation techniques for experiments on travelling-wave convection. In F.H. Busse and L. Kramer, editors, *Proc. of the NATO Advanced Research Workshop on Nonlinear Evolution of Spatio-Temporal Structures in Dissipative Continuous Systems, Streitberg, Germany*, page 73, Plenum, New York, 1990.
- [24] P. Manneville. *Dissipative Structures and Weak Turbulence*. Academic Press, New York, 1990.
- [25] A. Rényi. *Probability Theory*. Akadémiai Kiadó, Budapest, 1970.
- [26] T. Schreiber. Extremely simple nonlinear noise reduction method. *Phys. Rev. E*, **47**:2401–2404, 1993.
- [27] C.G. Schroer, T. Sauer, E. Ott, and J.A. Yorke. Predicting chaos most of the time from embeddings with self-intersections. *Phys. Rev. Lett.*, **80**:1410–1413, 1998.
- [28] J. Sethuraman. The asymptotic distribution of the Rényi maximal correlation. *Commun. Stat., Theory Methods*, **19**:4291–4298, 1990.
- [29] G. Steinmeyer et al. Dynamical pulse shaping in a nonlinear resonator. *Phys. Rev. A*, **52**:830, 1995.
- [30] J. Swift and P. Hohenberg. Hydrodynamic fluctuations at the convective instability. *Phys. Rev. A*, **15**:319, 1977.
- [31] J.R. Thompson and R.A. Tapia. *Nonparametric Function Estimation, Modeling, and Simulation*. SIAM, Philadelphia, 1990.

- [32] H. Voss, M.J. Bünner, and M. Abel. Identification of continuous, spatiotemporal systems. *Phys. Rev. E*, **57**:2820–2823, 1998.
- [33] H. Voss and J. Kurths. Reconstruction of nonlinear time delay models from data by the use of optimal transformations. *Phys. Lett. A*, **234**:336–344, 1997.
- [34] H. Voss and J. Kurths. Reconstruction of nonlinear time delay models from optical data. *Chaos, Solitons & Fractals*, **10**:805–809, 1999.
- [35] H.U. Voss, P. Kolodner, M. Abel, and J. Kurths. Amplitude equations from spatiotemporal binary-fluid convection data. *Phys. Rev. Lett.*, **83**:3422–3425, 1999.
- [36] H.U. Voss, A. Schwache, J. Kurths, and F. Mitschke. Equations of motion from chaotic data: A driven optical fiber ring resonator. *Phys. Lett. A*, **256**:47–54, 1999.

Chapter 18

Optimization of Embedding Parameters for Prediction of Seizure Onset with Mutual Information

Alfonso M. Albano¹
Christopher J. Cellucci
Richard N. Harner
Paul E. Rapp

ABSTRACT Normal neurological function is characterized by a high volume of information transferred between different parts of the central nervous system. An imperfect, but nonetheless useful, assessment of this spatially distributed transfer process can be obtained by examining multichannel EEG records that are measured with an array of scalp electrodes. A nonlinear quantitative measure of the transfer can be obtained by calculating the pairwise mutual information of each electrode pair. The average mutual information of two time series is the amount of information of one that can be predicted by measuring the other. As suggested in our previous work, a reduction in information transfer, as estimated by this metric, can occur prior to clinically discernible seizure onset. In the case of the focal seizure examined here, this reduction first occurred in the area of the brain that was subsequently shown to contain the epileptogenic focus. Thus, the calculation contains two clinically valuable elements: a prediction of seizure onset and a preliminary localization of the epileptogenic focus. We predict that in the case of seizures that are generalized at onset, the initiation of a seizure will be preceded by a near-simultaneous reduction in cross-channel average mutual information for several electrode pairs.

18.1 Introduction

Mutual information has been used to study EEG data for at least two decades ([3, 6, 7, 8, 9, 10] and references quoted therein). Although it may have been the first nonlinear measure used for this purpose, some of the early works made use of linearizing assumptions in the course of their cal-

¹ Author for correspondence.

culations, which kept them from detecting nonlinear correlations in the data. In their 1987 work, Mars and Lopes da Silva [9] showed clearly how mutual information, calculated without linearizing assumptions, could detect correlations that were not detected by the correlation coefficient which was then, and now still is, widely used to seek correlations in multivariate data. Using techniques and terminology introduced earlier by Vastano and Swinney [11] in another context, Xu, *et al.* [12] used mutual information to assess the “information transfer” among various scalp EEG sites. Information transfer between two EEG channels in this context is the average mutual information of a data segment from one channel and a segment from the other channel after some time delay

In a recent work on multichannel EEG data that includes a well characterized epileptic seizure [2], we found that even when using epochs with as few as 1000 data points, information transfer calculated using embedded data yielded details of spatio-temporal structures not obtainable when unembedded data were used. For this data set, these spatio-temporal structures predicted the sites where the seizure will occur several seconds before expert visual detection of seizure onset in the raw EEG.

In this contribution we investigate some higher-dimensional embeddings of the same dataset we studied earlier and further explore the effects of varying embedding dimension, embedding lag, and the time delay between epochs being compared. To determine embedding parameters, we use the method of “Global False Nearest Neighbors” [1] to find a value for the embedding dimension, Fraser’s criterion using mutual information for determining the embedding lag [5], and a variant of this criterion to determine the time delay during which to assess information transfer [2]. We find that provided one does not stray too much from the parameters so obtained, and provided one does some embedding, there is considerable robustness of the results relative to variations in the actual values of parameters used. Since mutual information calculations in high embedding spaces is costly in terms of computational resources and time, this flexibility makes it possible to choose those embedding parameters that are computationally more economical.

18.2 Mutual Information

Let \mathbf{X} and \mathbf{Y} be two sets of measurements. Let $p_X(x)$ be the probability density that measurement of \mathbf{X} yields x ; $p_Y(y)$ the probability that measurement of \mathbf{Y} yields y ; and $p_{XY}(x, y)$ the joint probability that measurements of \mathbf{X} and \mathbf{Y} yield x and y . The *average mutual information* of \mathbf{X} and \mathbf{Y} is defined by [see, e.g., [1, 9, 11]].

$$I_{XY} = \sum_{x,y} p_{XY}(x, y) \log \left[\frac{p_{XY}(x, y)}{p_X(x)p_Y(y)} \right]. \quad (18.1)$$

On average, I_{XY} gives the amount of information gained about x from a measurement of y . The quantities x and y could be either scalars – the values of the signals from different time series, or different segments of the same time series – or vectors obtained from a time-delay embedding of these data. Earlier work compared scalar time series (see, e.g., [9]). Working with data computed from a reaction-diffusion equation, Vastano and Swinney [11] used embedded data, obtaining values of the embedding dimension and embedding lag using mutual information-based criteria due to Fraser and Swinney [5]. Vastano and Swinney [11] argue that if x is a time series measured at location X of a spatially extended system and y is measured at location Y after some time delay, t , then $I_{XY}(t)$ is a measure of the the information transferred from X to Y in time t , or of the communication between X and Y in that time.

18.2.1 Embedding

A scalar time series, $\{x_1, x_2, \dots, x_N\}$ may be embedded in a d -dimensional space by forming d -dimensional vectors,

$$(x_i, x_{i+L}, \dots, x_{i+(d-1)L}); \quad i = 1, 2, \dots, N - (d-1)L.$$

There exists a large variety of techniques for determining appropriate values of the embedding parameters – *i.e.*, the “embedding dimension,” d , and the “embedding lag,” L . These techniques are typically applicable to single-channel, or univariate data, usually presumed to be stationary (see, e.g., [1]). There is growing awareness that an optimal embedding strategy may depend on both the time series and the applied measure. That is, the embedding criterion that is optimal when studying fluid flow data may not be optimal in the analysis of a time series from an EEG. Similarly, a procedure for selecting embedding parameters when the correlation dimension is to be estimated may not be as successful when calculating Lyapunov exponents or entropy [4]. The calculation of mutual information using nonstationary data from different locations in a spatially extended system introduces some further complications. Embedding parameters appropriate for a time series measured at one location may not be appropriate for one measured at another location. Parameters that may be optimal for one segment of a time series may not at all be appropriate for another segment if the time series is sufficiently nonstationary. If the same embedding parameters are to be used to study an entire multichannel data set, it may become necessary to use, as a compromise, those values that work well enough for all pairs, but are not necessarily optimal for any pair.




FIGURE 18.1. Model Data Set 1. Channel 1 (top panel) consists of points 51-2050 of an EEG time series; Channels 2 and 3 (lower panels) are identical and consist of points 1-2000 of the same time series.




FIGURE 18.2. Model Data Set 2. The three channels are segments of the same EEG time series. Channel 1 here (top panel) is the same as Channels 2 and 3 of Model Data Set 1.

18.2.2 Computationally Constructed Multivariate Data Sets

To illustrate some of these issues, we consider two “three-channel model data sets.” In the first (Model Data Set 1, Figure 18.1), Channel 1 consists of points 51-2050 of an EEG time series sampled at 200 Hz and Channels 2 and 3 are points 1-2000 of the same time series. That is, Channels 2 and 3 of this multichannel data set are identical not because they communicate with each other, but because both receive signals (“information”) from Channel 1 after some time delay. In the second (Model Data Set 2, Figure 18.2), the three channels come from widely separated segments of the same time series. Inspection of Figure 18.2 shows that Channels 1 and 2 are similar although not identical, but both are vastly different from Channel 3.²




FIGURE 18.3. Model Data Set 1. Average mutual information of the first 1,000 points of Channel 1 and a time-delayed epoch of the same length as a function of time delay.

One way of determining a value of the embedding lag is to calculate the average mutual information, $I_{XY}(T)$ for which X is an epoch of a time series starting at time, t , say, and Y is an epoch of the same time series, of the same length as X , but starting at time, $t + T$. The value of the time delay, T , where the first minimum of $I_{XY}(T)$ occurs is taken as the embedding lag [5]. Figure 18.3 shows average mutual information *vs.* time delay for the first 1,000 points of channel 1 of Set 1. It has a first minimum

²Both “model data sets” are derived from Channel 1 of Fig. 18.8. Set 1 uses the first 2050 points (10.25 s) of the time series. Channel 1 of Set 2 consists of the first 2000 points, Channel 2 of points 1001-3000 (5-15 s), and Channel 3 of points 5001-7000 (25-35 s).




FIGURE 18.4. Model Data Set 2. Similar to Figure 18.5, showing results for Channel 1 (+), Channel 3 (o), and all channel pairs (x).

in the neighborhood of fifteen time units. Since the other channels of this set are essentially the same as Channel 1, the results for the other channels are identical. This situation is not at all true for Set 2 (Figure 18.4). The graph for Channel 1 is the same as that in Figure 18.3; that for Channel 3 does not reach a minimum until about thirty time units. Yet another result emerges if the average mutual information of all channel pairs is computed. Rather surprisingly, the graph now has a first minimum at approximately six time units.

To determine a value for the embedding dimension, we use the method of “Global False Nearest Neighbors” [1]. This method is based on the premise that if too small an embedding dimension is used, there would be points in the embedding space that are close to each other because the reconstructed attractor has not been sufficiently unfolded. These artificially close points - the false neighbors - would be separated in an embedding space of appropriately high dimension. Figure 18.5 shows, for Channel 1 of both data sets, the percentage of nearest neighbors that are false as a function of embedding dimension. It suggests an embedding dimension of around 5.

Figure 18.6 shows the average mutual information between each of the channels of model data set 1 and channels of the same data set delayed by fifty time units. These calculations are done in embedding dimensions 1, 3, 5, and 7. In each of the graphs, the value of AMI in the region bounded by “Reference Channel” values i and $i + 1$ and “Delayed Channel” values between j and $j + 1$ is the average mutual information of data from channel i and delayed data from channel j . All graphs indicate that, indeed, Channels 2 and 3, delayed by fifty time units, are similar to Channel 1. This is not




FIGURE 18.5. Model Data Set 1. Percentage of nearest neighbors that are “false” *vs.* embedding dimension.

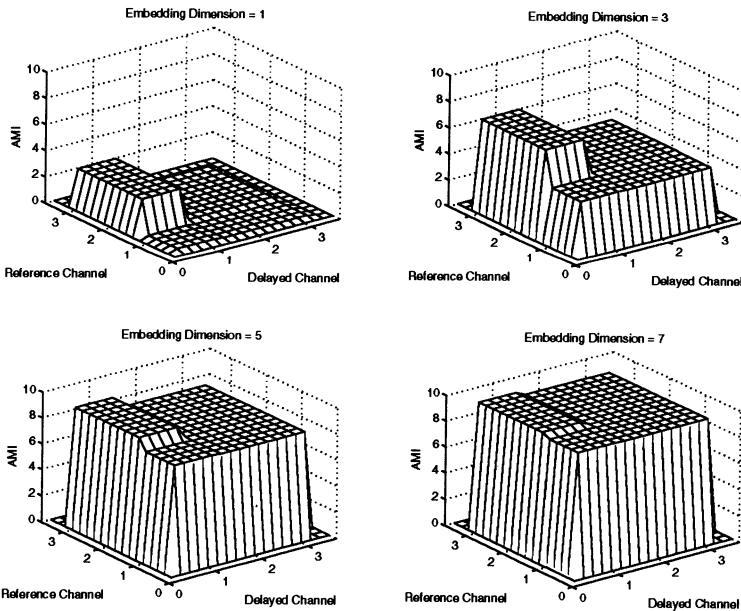


FIGURE 18.6. Model Data Set 1. Average mutual information of all “reference channel” – “delayed channel” pairs. The delay (50) is equal to the time delay that was introduced in the construction of the time series.

so for any other pairs. The results confirm the communication between Channel 1 and Channels 2 and 3 and is not affected by the identity of simultaneous epochs of the latter two channels. The maximum value of average mutual information increases with embedding dimension, but so does the minimum, so that the contrast between values for similar data and those for dissimilar data also deteriorates as embedding dimension increases. We investigate some consequences of this.

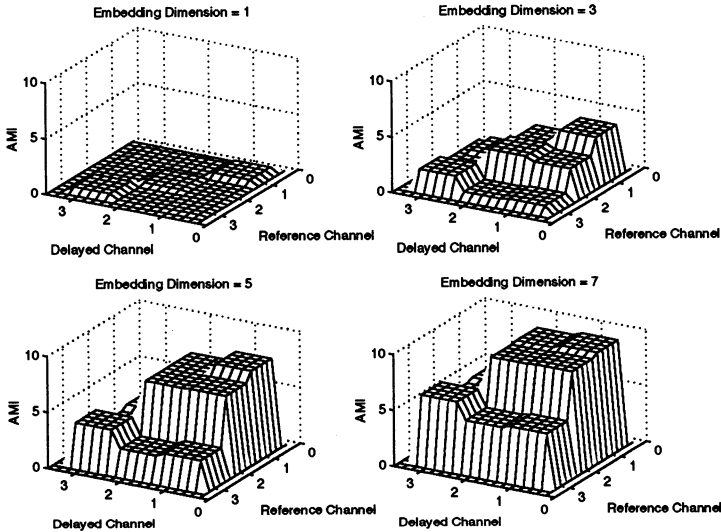


FIGURE 18.7. Model Data Set 2. Average mutual information of all “reference channel” – “delayed channel” pairs using lag = 6, delay = 8, and embedding dimensions 1, 3, 5, and 7.

Figure 18.7 shows the average mutual information for all channel pairs of Set 2 using a lag of 6, a time delay between channels of 8 and embedding dimensions of 1, 3, 5, and 7. The results for embedding dimension 1 indicate the similarities of each channel with delayed versions of itself. The results, however, do not indicate the similarity between Channels 1 and 2 and the differences between these and Channel 3. In other words, use of embedding dimension 1 (or the raw data) does not provide sufficient contrast to distinguish the similarity between Channels 1 and 2 and the differences between these and Channel 3. Results for dimensions 3 – 7 all show such a contrast. Embedding dimension 3 shows distinctions between the average mutual information of Channel 2 *vs.* Channel 2 (delayed), and Channel 1 *vs.* Channel 2 (delayed) or Channel 2 *vs.* Channel 1 (delayed). Embedding dimensions 5 and 7 do not show such distinctions as clearly. These results indicate some robustness relative to changes in embedding dimension, provided one does not stray too far from values suggested by techniques such as Global False Nearest Neighbors. Indeed, in the case of Figure 18.7, using an embedding

dimension in which approximately 0.05% of the nearest neighbors are false gives better results than higher dimensions in which the number of nearest neighbors is closer to zero. Since computational costs increase greatly with embedding dimension, and since there are well-known problems with scattering relatively few points in large embedding spaces, this might also justify using relatively small embedding spaces (but greater than one!).

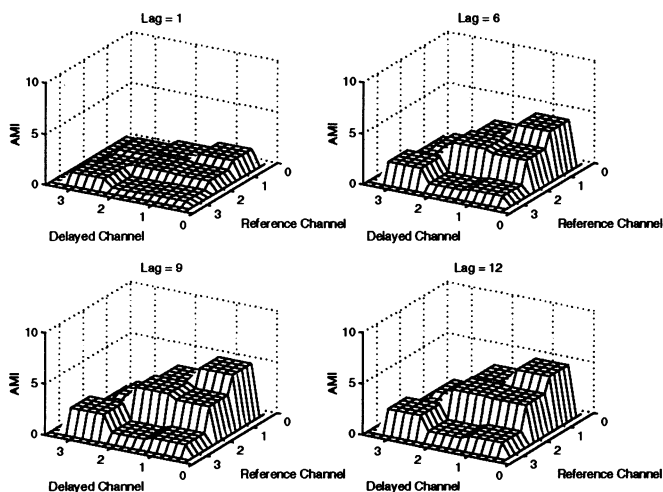


FIGURE 18.8. Model Data Set 2. Average mutual information of all channel pairs in embedding dimension 3, lag = 7 and delays = 0, 8, 16, 20.

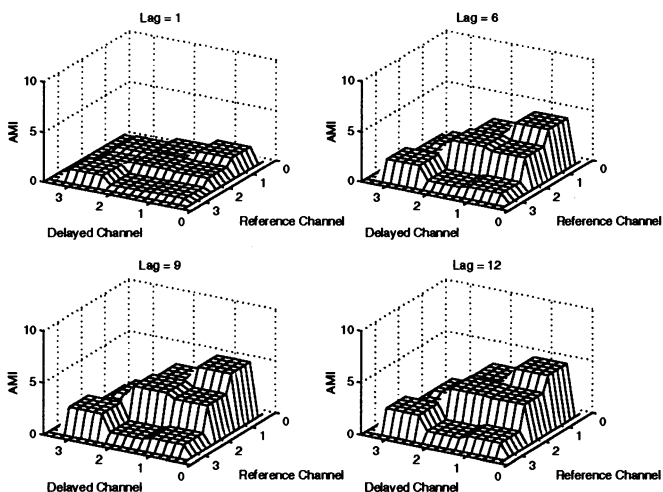


FIGURE 18.9. Model Data Set 2. Average mutual information of all channel pairs in embedding dimension 3, delay = 8, lag = 1, 6, 9, 12.

There is also some robustness against changes in the time delay between channels. Figure 18.8 shows the average mutual information calculated in embedding dimension 3, with lag 7, and delays of 0, 8, 16, and 20. All values of delay show the expected differences and similarities among all channel pairs. Figure 18.9 shows calculations similar to this but for different values of the lag. They are all done in dimension 3, with a delay of 8 and lags of 1, 6, 9, and 12. All show the expected behavior.

The results summarized in Figures 18.7–18.9 indicate that values of the embedding parameters and of the time delay can be obtained using the now familiar techniques—Global False Nearest Neighbors to determine embedding dimension, and the average mutual information of unembedded data to get the embedding lag and time delay. More importantly, they show that calculations of information transfer, or average mutual information of time-delayed epochs, is not unduly sensitive to the values of these parameters. We therefore have some flexibility in choosing parameter values to take into account limitations imposed by the amount of data available or by considerations of computational economy.

18.3 Multichannel EEG

18.3.1 Data and Analysis




FIGURE 18.10. Nineteen-channel EEG time series. The sequence of channel numbers (1-19) corresponds to the sequence of electrode sites given in the text.

Figure 18.10 shows our multichannel data set. The data were recorded from nineteen scalp sites, conventionally called Fp1, F3, F7, C3, T3, T5, P3,

O1, Fp2, F4, F8, C4, T4, T6, P4, O2, Fz, Cz and Pz. Measurements were made with a sampling rate of 200 Hz and digitized at 10 bits. There were 10,400 recorded values per channel. The seizure lasts from approximately $t = 12$ s to $t = 52$ s, with a focus near sites T3 and F7 (channels 3 and 5), corresponding to the left temporal lobe. Figure 18.11 shows the signal from site T3, which is the electrode closest to the focus. The seizure




FIGURE 18.11. Time series for site T3 (Channel 5).

is preceded by interictal spikes, also at T3 and F7. The large amplitude signals extending from approximately $t = 25$ s to $t = 42$ s are eyeblink artifacts, and appear in most of the channels, maximal in Channels 1 and 9 (Fp1 and Fp2) which are nearest the eyes.

Our analysis of the model data sets in the previous section suggests that there is considerable leeway in the choice of embedding parameters and of the time delay between epochs whose average mutual information is being calculated. In this spirit, we will use calculated values of these parameters as starting points and investigate the effects of varying these values. A Global False Nearest Neighbor calculation such as that shown in Figure 18.5 suggests an embedding dimension of approximately 5. Figures 18.12 and 18.13 are similar to Figures 18.3 and 18.4 and show the average mutual information of all channel pairs for embedding dimensions 1 (or no embedding at all) and 3, respectively. In the latter case, the lag was six time units or 30 ms. Figure 18.12 suggests a lag of seven units (35 ms) and a time delay of twelve units (60 ms). Figure 18.13 as well as a similar calculation in embedding dimension 5 (not shown) suggests a lag of four units (20 ms) and a time delay of six units (30 ms) between epochs. To the extent that these figures have some physiological basis unlike those obtained




FIGURE 18.12. Average mutual information for all channel pairs *vs.* time delay. Unembedded data using the first 1000 data points.




FIGURE 18.13. Average mutual information for all channel pairs *vs.* time delay. Embedding dimension 3, lag = 6, using the first 1000 data points.

from Figure 18.4, they suggest a characteristic “information transfer” time on the order of a few tens on milliseconds, which is not unreasonable.




FIGURE 18.14. Average mutual information for all channel pairs in nine 1000-point (5.0-s) epochs. Time sequence goes from left to right, top to bottom. The seizure starts in the top right panel. Unembedded data, delay = 12.




FIGURE 18.15. Similar to Figure 18.14 but in embedding dimension 3, lag = 4, delay = 6. In the first two panels, which precede seizure onset, Channels 3 and 5, which are nearest the eventual seizure focus, have significantly lower values of average mutual information. In the third panel (top right) when the seizure occurs, the overall value of average mutual information is reduced.




FIGURE 18.16. Similar to Figure 18.14, but in embedding dimension 5, lag = 20, delay = 20.




FIGURE 18.17. Similar to Figure 18.14 but in embedding dimension 7, lag = 10, delay = 0.

In Figures 18.14 to 18.17, we show the results of a number of calculations, some using parameters values at or close to those found earlier, others using very different values. In each of these figures, the top left panel shows the average mutual information for all channel pairs for points 1 - 1000 (0-5 s). The succeeding panels, from left to right and top to bottom are for succeeding, non-overlapping 1000-point segments of nineteen-channel data. That is, the second panel covers points 1001-2000 (5-10 s), the third covers

points 2001-3000 (10-15 s), *etc.* In each of these figures, the seizure starts in the time interval covered by the third panel (top right). For each panel, with the bottom left corner as $(0,0)$, the cell at column x , row y , is the average mutual information of Channel x and Channel y (delayed).

Figure 18.14 uses unembedded data (embedding dimension = 1). As suggested by Figure 18.12, the time delay is set to 12 time units (lag is meaningless when the data are not embedded). There are some changes in patterns of the average mutual information, but there are no dramatic indications of an impending or actual seizure until the seventh panel which starts at $t= 25.0$ s at which time the seizure is well underway and the EEG signals from practically all channels are dominated by eyeblink artifacts. Figures 18.15 - 18.17, on the other hand, clearly indicate that Channels 3 and 5, which are nearest the eventual seizure focus, are behaving differently from the rest. The low values of average mutual information between them and all other channels suggest that they are “not communicating” and this indication appears some 10 s before seizure onset. Figure 18.15 uses an embedding dimension of 3, and a lag of 4 and delay of 6 as suggested by Figure 18.13. In the interval that includes seizure onset (top right), the average mutual information of all channel pairs drops, recovering at the next epoch at which time values for Channels 1 and 9 drop, corresponding to the predominance of eyeblink artifacts in these channels. Low levels of average mutual information characterize these two channels until the end of the time series. Figure 18.16 (embedding dimension = 5, lag = 20, delay = 20) and Figure 18.17 (embedding dimension = 7, lag 10, delay = 0) use values of embedding and delay parameters that are quite different from those calculated using the techniques discussed earlier. Nevertheless, they both mark Channels 3 and 5 from the beginning of the time series, show the drop in average mutual information at seizure onset, and mark the eyeblink-dominated channels until the end. Indeed, we have obtained similar results in embedding dimension 3 for lags from 5 to 200 (lower values of lag gave results similar to those found for unembedded data), and in embedding dimension 5 for delays from 0 to 100. In Figure 18.18, we summarize the time course of the “information transfer” to each channel by showing the average mutual information for each “delayed channel” averaged overall “reference channels.” That is, for each of the panels in Figures 18.14- 18.17, we take the average of each row and then show how this average changes from one panel to the next, resulting in an “information transfer” *vs.* time plot. Showing similar data for each “reference channel” by averaging each column would give similar results because of the approximate symmetry of the panels in Figures 18.14- 18.17 about their main diagonals. Investigation of the actual values of the average mutual information, however, shows that the panels are not strictly symmetric, which must be so if these values really mean information transfer.

The top left panel of Figure 18.18 shows the results for unembedded data (Figure 18.14); top right: embedding dimension 3 (Figure 18.15), bottom left: embedding dimension 5 (Figure 18.16) and bottom right: embedding dimension 7 (Figure 18.17). As before, for all but the unembedded data,

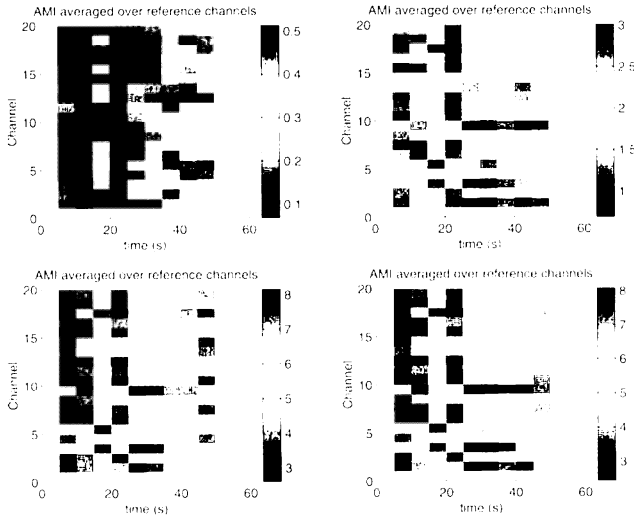


FIGURE 18.18. Average mutual information averaged over all reference channels vs. time. This shows the time dependence of the “information received” by each channel from all the rest, and how this changes in time. From left to right and top to bottom, these panels are for embedding dimensions 1, 3, 5, and 7. These were obtained by averaging the rows of Figures 18.14-18.17.

Channels 3 and 5 which contain interictal spikes are characterized by low information transfers. The third epoch (10-15 s) in which the raw EEG shows minimal visual evidence of seizure onset in a single channel is characterized by low information transfer in all channels. Channels 1 and 9 similarly show low information transfers during those epochs when they are dominated by eyeblinks. Calculations using embedded data show localized events such as spikes and eyeblinks as well as more global events characterizing the seizure. This figure also has the advantage of showing more clearly how the patterns evolve in time.

18.4 Concluding Remarks

The average mutual information of non-simultaneous data segments measured from different locations in a spatially extended system promises to be a useful tool for probing the inter-relationships among different regions of the system and for studying the time evolution of these relationships. It could be helpful in the analysis of multichannel EEG’s. If, for instance, it can provide a way of quantifying features of the EEG which clinicians use for diagnosis, it could become a useful diagnostic tool. We confirm a result we have reported elsewhere ([2]); that the average mutual informa-

tion calculated using embedded data is able to mark the sites where an epileptic seizure will occur several seconds before seizure onset. To perform the calculation, embedding parameters may be obtained using some of the now familiar techniques ([1]). However, we find that provided one does not stray much from the parameters so obtained, and provided one does some embedding, there is considerable robustness of the results relative to variations in the actual values of parameters used.

Acknowledgments

This work was supported in part by NIH Grant NS32983.

References

- [1] H. D. I. Abarbanel. *Analysis of Observed Chaotic Data*. Springer, New York, 1995.
- [2] A. M. Albano, C. Bedonie, C. J. Cellucci, V. Miller, J. Ree, A. Torruela, R. N. Harner, and P. E. Rapp. Spatiotemporal EEG information transfer in an episode of epilepsy. In J. Pradhan, R. Sreenivasan, and P. E. Rapp, editors, *Nonlinear Dynamics and Brain Functioning*, New York, In press. Nova Science.
- [3] E. Callaway and P. R. Harris. Coupling between cortical potentials from different areas. *Science*, 183:873, 1974.
- [4] C. J. Cellucci. *On Nonlinear Time Series Analysis*. PhD thesis, Bryn Mawr College, 1998.
- [5] A. M. Fraser and H. L. Swinney. Independent coordinates for strange attractors from mutual information. *Physical Review A*, 33:1134, 1986.
- [6] W. Gersch. Non-stationary multichannel time series analysis. In A. Gevins and A. Rémond, editors, *Methods of Analysis of Brain Electrical and Magnetic Signals, EEG Handbook*, volume 1, page 261, Amsterdam, 1987. Elsevier.
- [7] W. Gersch and B. R. Tharp. Spectral regression – amount of information analysis of seizures in humans. In P. Kellaway and I. Petersen, editors, *Quantitative Analytic Studies in Epilepsy*, page 509, New York, 1976. Raven Press.
- [8] F. Lopes da Silva, J. P. Pijn, and P. Boeijinga. Interdependence of EEG signals: linear vs. nonlinear associations and the significance of time delays and phase shifts. *Brain Topography*, 2:9, 1989.
- [9] N. J. I. Mars and F. H. Lopes da Silva. EEG analysis methods based on information theory. In A. Gevins and A. Rémond, editors, *Methods of Analysis of Brain Electrical and Magnetic Signals, EEG Handbook*, volume 1, page 297, Amsterdam, 1987. Elsevier.
- [10] N. J. I. Mars, P. M. Thompson, and R. J. Wilkus. The spread of epileptic seizure activity in humans. *Epilepsia*, 26:85, 1985.
- [11] J. A. Vastano and H. L. Swinney. Information transport in spatiotemporal systems. *Physical Review Letters*, 60:1773, 1988.
- [12] J. Xu, Z. R. Liu, R. Liu, and Q. F. Yang. Information transmission in human cerebral cortex. *Physica D*, 106:363, 1997.

Chapter 19

Detection of a Nonlinear Oscillator Underlying Experimental Time Series: The Sunspot Cycle

Milan Paluš

ABSTRACT After a brief review of a nonlinearity test based on information theoretic functionals (redundancies) and surrogate data technique, we discuss problems of this and similar tests for nonlinearity. In particular, we stress that a formal rejection of a linear stochastic null hypothesis does not automatically mean evidence for nonlinear dynamical origin of studied data. In an example we show how a variable variance could be mistaken for nonlinearity in a series of surface air pressures. Therefore we find a detection of nonlinearity in a series of sunspot numbers insufficient for an identification of a mechanism underlying the sunspot cycle. As a solution in this case we propose to test for a property of nonlinear oscillators – mutual dependence between their instantaneous amplitude and frequency. This behavior is detected in yearly and monthly records of the sunspot numbers using histogram-adjusted isospectral surrogate data and Barnes model as ARMA surrogates. The instantaneous amplitudes and frequencies are obtained by means of the analytic signal approach using the discrete Hilbert transform. In several tests the amplitude-frequency correlation has been found significant on levels ranging from $p < 0.03$ to $p < 0.07$, which supports the hypothesis of a driven nonlinear oscillator as a mechanism underlying the sunspot cycle.

19.1 Introduction

Let $\{y(t)\}$ be a time series, i.e., a series of measurements done on a system in consecutive instants of time $t = 1, 2, \dots$. Can we identify a mechanism underlying temporal evolution of such a system?

The time series $\{y(t)\}$ can be considered as a realization of a stationary linear stochastic process $\{Y(t)\}$. Without loss of generality we can set its mean to zero. Then the linear stochastic process $Y(t)$ can be written as:

$$Y(t) = Y(0) + \sum_{i=1}^{\infty} a(i)Y(t-i) + \sum_{i=0}^{\infty} b(i)N(t-i), \quad (19.1)$$

where $b(0) = 1$, $\sum_{i=1}^{\infty} |a(i)| < \infty$, $\sum_{i=0}^{\infty} |b(i)| < \infty$, and $\{N(t)\}$ is an independent, identically distributed (I.I.D.), normally distributed process with zero mean and finite (constant) variance. (For more details see [18].)

Alternatively, the time series $\{y(t)\}$ can be considered as a projected trajectory of a dynamical system, evolving in some measurable d -dimensional state space. To be more specific, let X_t denote a state vector in R^d . Then the measurements $y(t)$ are obtained as $y(t) = g(X_t)$, where $g(\cdot)$ is a projection (measurement function), and temporal evolution of X_t may be described by a discrete-time dynamical system (a difference equation):

$$X_t = F(X_{t-1}), \quad (19.2)$$

with $X_0 \in R^d$ and for $t \geq 1$.

Due to ubiquity of noise, it is more realistic to replace the preceding states by random variables and the dynamics by a Markovian model such as

$$X_t = F(X_{t-1}, e_t), \quad (19.3)$$

where $t \in Z_+$, $F : R^{2d} \rightarrow R^d$, $\{e_t\}$ is a sequence of independent and identically distributed d -dimensional random vectors and e_t is independent of X_s , $0 \leq s < t$. We call $\{e_t\}$ the *dynamic noise*. Following [25], we refer to Equation (19.2) as the *skeleton* of model (19.3), considering $F(X) = F(X, 0)$. For convenience, it is frequently assumed that the dynamic noise is additive so that Equation (19.3) reduces to the *model with additive noise*

$$X_t = F(X_{t-1}) + e_t. \quad (19.4)$$

Detection of nonlinearity in experimental time series, i.e., identification from experimental data of underlying mechanism such as the model (19.4) with a nonlinear function F is usually based on rejection of a linear null hypothesis by a statistical test. Typically, the considered null hypothesis is a linear Gaussian process such as (19.1) or a Gaussian process passed through a static nonlinear transformation, or a similar simple alternative. Rejection of such a null is frequently interpreted as a detection of a deterministic nonlinear relation (19.2 or 19.4) in data under study. This is, however, only one of possible alternatives. Other alternatives will be discussed in Sec. 19.3, in the following section we will briefly review a test for nonlinearity proposed in [12]. Then we will present an example of a spurious nonlinearity detection due to a variable variance in a case of air surface pressure data (Sec. 19.4). Nonlinearity in a series of sunspot numbers is tested in Sec. 19.5. Section 19.6 introduces the amplitude-frequency correlation in nonlinear oscillators. Detection of such a behavior in the sunspot cycles is presented in Sec. 19.7, and results are discussed in Sec. 19.8.

19.2 A Test for Nonlinearity Based on Redundancies and Surrogate Data

Consider n discrete random variables X_1, \dots, X_n with sets of values Ξ_1, \dots, Ξ_n , respectively. The probability distribution for an individual X_i is $p(x_i) = \Pr\{X_i = x_i\}$, $x_i \in \Xi_i$. We denote the probability distribution function by $p(x_i)$, rather than $p_{X_i}(x_i)$, for convenience. Analogously, the joint distribution for the n variables X_1, \dots, X_n is $p(x_1, \dots, x_n)$. The redundancy $R(X_1; \dots; X_n)$, in the case of two variables also known as mutual information $I(X_1; X_2)$, quantifies average amount of common information, contained in the n variables X_1, \dots, X_n :

$$R(X_1; \dots; X_n) = \sum_{x_1 \in \Xi_1} \dots \sum_{x_n \in \Xi_n} p(x_1, \dots, x_n) \log \frac{p(x_1, \dots, x_n)}{p(x_1) \dots p(x_n)}. \quad (19.5)$$

When the discrete variables X_1, \dots, X_n are obtained from continuous variables on a continuous probability space, then the redundancies depend on a partition ξ chosen to discretize the space. Various strategies have been proposed to define an optimal partition for estimating redundancies of continuous variables (see [12, 27] and references therein). Here we use the “marginal equiquantization” method described in detail in [11, 12].

Now, let the n variables X_1, \dots, X_n have zero means, unit variances and correlation matrix \mathbf{C} . Then, we define the *linear redundancy* $L(X_1; \dots; X_n)$ of X_1, X_2, \dots, X_n as

$$L(X_1; \dots; X_n) = -\frac{1}{2} \sum_{i=1}^n \log(\sigma_i), \quad (19.6)$$

where σ_i are the eigenvalues of the $n \times n$ correlation matrix \mathbf{C} .

If X_1, \dots, X_n have an n -dimensional Gaussian distribution, then $L(X_1; \dots; X_n)$ and $R(X_1; \dots; X_n)$ are theoretically equivalent.

In practical applications one deals with a time series $\{y(t)\}$, considered as a realization of a stochastic process $\{Y(t)\}$, which is stationary and ergodic. Then, due to ergodicity, all the subsequent information-theoretic functionals are estimated using time averages instead of ensemble averages, and the variables X_i are substituted as

$$X_i = y(t + (i - 1)\tau). \quad (19.7)$$

Due to stationarity the redundancies

$$R^n(\tau) \equiv R(y(t); y(t + \tau); \dots; y(t + (n - 1)\tau)) \quad (19.8)$$

and

$$L^n(\tau) \equiv L(y(t); y(t + \tau); \dots; y(t + (n - 1)\tau)) \quad (19.9)$$

are functions of n and τ , independent of t .

The surrogate-data based nonlinearity tests [23, 12] consist of computing a *nonlinear* statistic from data under study and from an ensemble of realizations of a linear stochastic process, which mimics “linear properties” of the studied data. If the computed statistic for the original data is significantly different from the values obtained for the surrogate set, one can infer that the data were not generated by a linear process; otherwise the null hypothesis, that a linear model fully explains the data, cannot be rejected and the data can be analyzed and characterized and predictions can be obtained by using well-developed linear methods. For the purpose of such a test the surrogate data must preserve the spectrum and consequently, the autocorrelation function of the series under study. An isospectral linear stochastic process to a series can be constructed by computing the Fourier transform (FT) of the series, keeping unchanged the magnitudes of the Fourier coefficients but randomizing their phases and computing the inverse FT into the time domain. Different realizations of the process are obtained using different sets of the random phases.

The FT surrogates tend to have a Gaussian distribution which is not always the case with the tested data. To avoid a possible influence of different histograms of the data and of the surrogates, the histogram adjusted FT (HAFT) surrogates are constructed. (In [23] the term “amplitude-adjusted” – AAFT surrogates is used.) In this case, the raw data undergo a nonlinear transformation which leads to a Gaussian distribution of the transformed data (Gaussianization — see [12] and references within). The gaussianized data are used to generate the FT surrogates as described earlier, and the obtained surrogate data are transformed to have the same histogram as the original raw data. An application of the HAFT surrogates effectively means a reformulation of the null hypothesis of a Gaussian linear stochastic process (19.1) into a hypothesis of a process (19.1), realizations of which are passed through a static nonlinear transformation.

To evaluate the test, usually, and in [12], the test statistic is defined as a difference between the redundancy obtained for the original data and the mean redundancy of a set of surrogates, in the number of standard deviations (SD's) of the latter. Such an approach requires a relative small number of surrogate replications, however, a normal distribution of the test statistic is considered. The latter is not always the case, and therefore, as a more reliable approach, in Sec. 19.5 we generate a large number of surrogate realizations and estimate percentiles (e.g., 2.5th and 97.5th percentile) of the surrogate distribution. Then redundancy values are simply compared with related percentile values. It is also possible to directly estimate the “significance”, i.e., the p value of the test by counting the percentage of the surrogate realizations yielding redundancy values equal to or greater than the related redundancy obtained from analyzed data.

Note that both the redundancies and redundancy-based statistics can be

evaluated as functions of lag τ and embedding dimension n . Evaluating the redundancies and related statistics for broad ranges of lags (and several dimensions), however, can bring the problem of simultaneous statistical inference (see [12, 15] and references within for details). This approach, however, can be more reliable than single-valued tests, as was demonstrated in [12].

The redundancy $R^n(\tau)$, based on probability distributions, measures general dependences among the series $\{y(t)\}$ and its lagged versions, whereas the linear redundancy $L^n(\tau)$, based on correlations, reflects only their linear relations. Comparing the plots of $R^n(\tau)$ and $L^n(\tau)$ can provide an informal test for important nonlinearities in the studied data [14]. This approach, in [12] referred to as *qualitative testing* or *qualitative comparison*, can bring additional information to the results of the quantitative (surrogate data based) test. Moreover, one cannot always construct good surrogate data. That is, despite theoretical expectations, in numerical practice linear properties of the surrogates may differ from those of the data under study. Changes in linear properties are reflected in nonlinear statistics as well, and thus they may result in spurious detection of nonlinearity in linear data [12]. Therefore, we also evaluate a statistic based on the linear redundancy $L^n(\tau)$ (defined analogously to the statistic based on $R^n(\tau)$), which specifically reflects changes in linear properties. Then, only those significant differences in the nonlinear statistic can reliably count for nonlinearity, which are not detected in the linear statistic [12].

19.3 The Null Hypothesis of Nonlinearity Tests and Its Negations

Consider that the just-described test yielded a significant result, i.e., the null hypothesis was reliably rejected. The null hypothesis was equivalent¹ to a linear stochastic process such as that described by the ARMA model (19.1). It is very common in nonlinear dynamics literature to consider the rejection of the null (19.1) as an evidence for a process such as (19.4) with a nonlinear skeleton (19.2). This is, however, only one of the possible negations of (19.1). A number of different processes should be considered, which possess a linear deterministic skeleton², i.e., a linear AR part – the first sum in (19.1), or no deterministic skeleton at all (MA processes), however, their innovations $\{N(t)\}$ do not fulfill the conditions given earlier. Generally, one or more of the following properties could reject the null (19.1):

¹Principal equivalence and technical differences between the surrogate data constructed by using the Fourier transform and the ARMA modeling are discussed in [23, 24].

²Obviously, for a linear function F , the model (19.4) is a special case of (19.1).

1. The innovations $\{N(t)\}$ are not Gaussian.
2. The innovations $\{N(t)\}$ are not an I.I.D. process, where I.I.D. means that the innovations should be not only uncorrelated, but generally independent.
3. The variance of $\{N(t)\}$ is not constant.




FIGURE 19.1. (a): Linear redundancy $L(y(t); y(t + \tau))$ (solid line), (b): nonlinear (general) redundancy $R(y(t); y(t + \tau))$ (solid line), for a series of differences from the long term averages of the surface air pressure (Prague-Klementinum station) and for its FT surrogates (thin solid and dashed lines present mean and mean \pm SD, respectively, of a set of 30 surrogate realizations); (c): linear (L -based), and d): nonlinear (R -based) statistics; as functions of the time lag τ , measured in days.

19.4 An Example of Surface Air Pressure Series

A series of daily recordings of surface air pressure $\{P(t)\}$ ($t = 1 - 65,536$ days; i.e., 180 years) was described and analyzed in [15]. Here we ana-




FIGURE 19.2. Seasonality in mean (upper panel) and in variance (lower panel) of the surface air pressure series – long term means (upper panel) and standard deviations (square root of variance, lower panel) for each day in a year. The days are consecutively numbered, January 1 has the index 1, January 2 has the index 2, ..., February 1 has the index 32, etc.

lyze differences from long-term daily averages. This transformation of data (almost entirely) removes the oscillations with the period of one year (seasonality in mean). We ask the question about possible long-term nonlinear dependence and perform the linear redundancy – redundancy surrogate data test, described earlier.

The results are presented in Fig. 19.1. For lags larger than a few days there is only a weak linear dependence, as measured by the linear redundancy (Fig. 19.1a) and reflected in the surrogates, but the (nonlinear) redundancy detects a clear dependence as an oscillatory structure with a yearly periodicity. This difference is highly significant (10 – 30 SD's, Fig. 19.1d), while no significance in the linear statistic (Fig. 19.1c) confirm the quality of the surrogates, which reflect correctly “the linear properties” of the data (in the sense of the model (19.1)). Can this result be understood as an evidence for the model (19.4) with a nonlinear periodic skeleton F , which could provide predictability of the atmospheric pressure for several




FIGURE 19.3. (a): Linear redundancy $L(y(t); y(t + \tau))$ (solid line), (b): nonlinear (general) redundancy $R(y(t); y(t + \tau))$ (solid line), for a series of differences from the long term averages of the surface air pressure, rescaled in order to have a constant variance, and for its FT surrogates (thin solid and dashed lines present mean and mean \pm SD, respectively, of a set of 30 surrogate realizations); (c): linear (L -based), and d): nonlinear (R -based) statistics; as functions of the time lag τ , measured in days.

years in advance?

The seasonality in mean present in this data (Fig. 19.2, upper panel) was mostly removed by considering differences from the long term averages.³ The problem is that also the variance of this data is not constant, but clearly seasonal (Fig. 19.2, lower panel, the standard deviation (SD) is the square root of variance). This property is “nonlinear” in the sense that the surrogate data and the model (19.1) possess a constant variance and cannot reproduce the seasonality in variance. After rescaling the data

³The seasonality in mean can be entirely removed by filtration in spectral domain, as done in [15].

in order to obtain a constant variance, the effect of the false long-term nonlinear dependence is lost (Fig. 19.3). In the rescaled data there is no long-term dependence except of weak linear link due to not entirely removed seasonality in mean.

The preceding example of the surface air pressure clearly demonstrated the influence of variable variance on the redundancy – surrogate data non-linearity test (introduced in Sec. 19.2 and in detail in [12]). The effect of non-Gaussian innovations $\{N(t)\}$ was discussed in [12], and a possible influence of non-I.I.D. $\{N(t)\}$ (i.e., innovations containing (nonlinear) temporal structures) is understandable. It is important to note that similar effects of “defective” innovations in a process under study would effect not only this particular test for nonlinearity, but all tests which use some type of FT/ARMA surrogates and any method which contains the process (19.1) at least implicitly in its construction. Also, all entropy-related statistics, that is, not only the preceding information-theoretic functionals, but also, for instance, statistics based on correlation integral [17], are extremely sensitive to variable variance and/or to (non)Gaussianity of data/innovations. Therefore one must very carefully assess results of nonlinearity tests in order to avoid confusing this kind of effect with actual nonlinear functional dependence in the data under study.




FIGURE 19.4. (a) The series of yearly sunspot numbers (1700 – 1997). (b) A realization of the HAFT surrogate data for the “last” 256 samples. (c) A 298-sample realization of the Barnes model.

19.5 Nonlinearity in the Sunspot Cycle

An energy output of the sun as the main basis of life on Earth is nearly constant. However, the Sun is far from being uniform. The best observed solar inhomogeneities are spots on the solar surface in which the luminosity is diminished but magnetic fields appear which are stronger than usual magnetic fields on the rest of the solar surface. Appearance of the spots on the sun is quantified by so-called sunspot numbers, or index, which have been collected from the beginning of the eighteenth century.

The historical data of the sunspot index have been attracting researchers for more than a century. The now well-known eleven-year cycle was reported in [30]. Of course, the sunspot cycle is not strictly periodic, but fluctuations in its amplitude and frequency (i.e., in the cycle duration) occur. Therefore researchers have turned to stochastic models in order to make predictions of a future behavior of the sunspot cycle (see [29] and references therein). On the other hand, development in nonlinear dynamics and theory of deterministic chaos, namely methods and algorithms for analysis and prediction of (potentially) nonlinear and chaotic time series have naturally found their way into the analyses of the sunspot series. Several authors ([10, 8] and references therein) have claimed an evidence for a deterministic chaotic origin of the sunspot cycle, based on estimations of correlation dimension, Lyapunov exponents and an increase of a prediction error with a prediction horizon. The dimensional algorithms, however, have been found unreliable when applied to relatively short experimental data, and properties consistent with stochastic processes (colored noises) such as autocorrelations can lead to a spurious convergence of dimensional estimates [22]. Similar behavior has been observed for Lyapunov exponent estimators [3, 13]. The increase of a prediction error with an increasing prediction horizon is not a property exclusive to chaos; it can also be observed in systems with a deterministic skeleton and an intrinsic stochastic component (“dynamical noise”). Therefore, such results cannot be considered convincing evidence for a nonlinear dynamical origin of the sunspot cycle.

In this section we present results of the test for nonlinearity (Sec. 19.2) applied to the sunspot data. The series of yearly sunspot numbers from the period 1700 – 1997 was obtained from the Sunspot Index Data Center⁴, and is illustrated in Fig. 19.4a. As test results, here we do not present differences from surrogate mean in number of standard deviations as in the previous section, but using 15,000 surrogate replications we estimate the 2.5th and 97.5th percentiles of distributions of redundancies computed from surrogate data. If values of related redundancy obtained from the studied data lie outside the range given by these two percentile values, the

⁴Sunspot Index Data Center, Royal Observatory of Belgium, Av. Circulaire, 3, B-1180 Brussels, Internet address: <http://www.oma.be/KSB-ORB/SIDC>, file `yearssn.dat`.

null hypothesis is rejected since the test is significant with $p < 0.05$. It means that the probability of the null hypothesis (that the studied data can be explained by the surrogate model) is less than 5 percent.

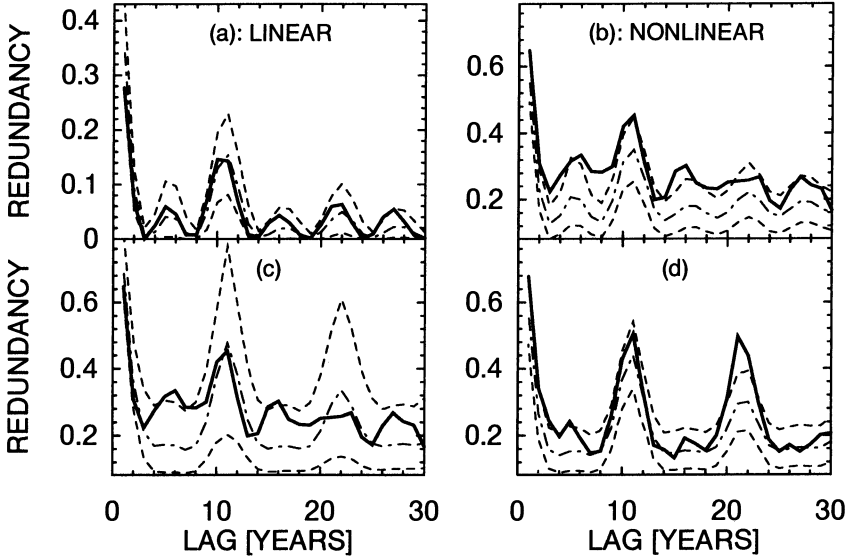


FIGURE 19.5. (a): Linear redundancy $L(y(t); y(t+\tau))$ (solid line) for the series of sunspot numbers as a function of the lag τ ; linear redundancy $L(y(t); y(t+\tau))$ for the related HAFT surrogate set; dash-and-dotted line presents the mean, dashed lines illustrate the 2.5th and the 97.5th (bottom and top line, respectively) percentile of the surrogate $L(y(t); y(t+\tau))$ distribution. (b): Nonlinear (general) redundancy $R(y(t); y(t+\tau))$ (solid line), for the series of sunspot numbers as a function of the lag τ ; nonlinear redundancy $R(y(t); y(t+\tau))$ for the related HAFT surrogate set, the same line codes as in (a). (c): Nonlinear redundancy $R(y(t); y(t+\tau))$ (solid line), for the series of sunspot numbers as a function of the lag τ ; nonlinear redundancy $R(y(t); y(t+\tau))$ for the Barnes model surrogate set, the same line codes as in (a). (d): Nonlinear redundancy $R(y(t); y(t+\tau))$ (solid line), for a realization of the Barnes model as a function of the lag τ ; nonlinear redundancy $R(y(t); y(t+\tau))$ for the related HAFT surrogate set, the same line codes as in (a).

The linear redundancy $L(y(t); y(t+\tau))$ (solid line in Fig. 19.5a) lies clearly inside the 2.5th and 97.5th percentiles of the $L(y(t); y(t+\tau))$ HAFT surrogate distribution (dashed lines in Fig. 19.5a). This test just checks the quality of the surrogate data: it says the linear properties (dependence structures) in the sunspot data do not differ from those of the HAFT surrogates (a realization presented in Fig. 19.4b), so the surrogates should not be a source of a spurious detection of nonlinearity.

The nonlinearity test itself is presented in Fig. 19.5b, where the (non-

linear) redundancy $R(y(t); y(t + \tau))$ (solid line in Fig. 19.5b) is, for the majority of studied lags, higher than the 97.5th percentile of the HAFT surrogate distribution (the upper dashed line in Fig. 19.5b). Thus the null hypothesis of a linear stochastic process (19.1), possibly passed through a static nonlinear transformation, is rejected.

Does this rejection really mean that a nonlinear dynamical system such as (19.2) or (19.4) underlies the sunspot cycle, or can this rejection be explained by any of the reasons listed in Sec. 19.3? This question is hard to answer. For instance, we cannot evaluate properties of innovations (model residuals) without a-priory knowledge of a valid model. No direct connection of the sunspot data to any physically inspired nonlinear dynamical model is known. Physical models trying to explain the variation of the solar activity come from the dynamo theory (cf. [7, 19]). The principle of such a self-exciting dynamo is that the magnetic field is amplified and maintained by the interaction of mainly three types of hydrodynamic plasma motions, namely differential rotation, turbulent convection and helicity. It is interesting to mention that there are some rather simple conceptual dynamo models which show a rich dynamical behavior and can explain several facts known from observations [4]. Such models, however, are not fitted directly to the experimental data, and are evaluated only in a qualitative way. Therefore, we should also consider concurrent linear stochastic models, such as the Barnes model [1]:

$$z_i = \alpha_1 z_{i-1} + \alpha_2 z_{i-2} + a_i - \beta_1 a_{i-1} - \beta_2 a_{i-2}, \quad (19.10)$$

$$s_i = z_i^2 + \gamma(z_i^2 - z_{i-1}^2)^2, \quad (19.11)$$

where $\alpha_1 = 1.90693$, $\alpha_2 = -0.98751$, $\beta_1 = 0.78512$, $\beta_2 = -0.40662$, $\gamma = 0.03$ and a_i are I.I.D. Gaussian random variables with zero mean and standard deviation $SD=0.4$.

The Barnes model (19.10, 19.11) is a rather simple but efficient model to mimic essential properties of the sunspot numbers. It incorporates the structure of an autoregressive moving average ARMA(2,2) model (19.10) with a nonlinear transformation (19.11) which ensures that the generated series remains asymmetric and positive and tends to increase more rapidly than it decreases. Moreover, the stochastic Barnes model can mimic some seemingly nonlinear properties such as behavior of correlation integrals [28] and phase portraits [21] obtained from the sunspot series.

We can evaluate, in the sense of this nonlinearity test, whether the Barnes model can explain the sunspot data, by using realizations of the Barnes model as surrogate data. The result of such a test is presented in Fig. 19.5c.

The nonlinear redundancy $R(y(t); y(t + \tau))$ (solid line in Fig. 19.5c) only slightly exceeds the 97.5th percentile of the Barnes surrogate distribution in three of the thirty studied lags. The rejection of the Barnes model by this test is not very convincing.

On the other hand, when we test for nonlinearity in a realization of the Barnes model (Fig. 19.5d), the HAFT surrogates are rejected. (The rejection is clear only in the lags 21 and 22, however, $R(y(t); y(t + \tau))$ of the tested series there exceeds the whole range of the surrogate values, i.e., $p = 0$ and the test is significant even considering the simultaneous statistical inference (see [12] and references therein). This result could be expected, since the nonlinear transformation (19.11) is not static; therefore the HAFT surrogates are rejected.

To summarize the last two tests, the realization of the Barnes model appeared in the nonlinearity test as nonlinear, and the rejection of the Barnes model as the null hypothesis is not very convincing. Can we find any solid argument for a nonlinear dynamical origin of the sunspot cycle, or should we accept a linear stochastic explanation, such as the Barnes model?

19.6 Amplitude-Frequency Correlation in Nonlinear Oscillators

In this section we demonstrate a typical property of nonlinear oscillators (a class of nonlinear dynamical systems), namely, the correlation between instantaneous amplitude and frequency of signals (solutions) generated by such systems.




FIGURE 19.6. (a) A solution of the nonlinear Duffing oscillator without any external driving force, and (b) the related instantaneous amplitude (thick line) and frequency (thin line).

As a demonstrative example of a nonlinear oscillator (not a model for

the sunspot cycle) we will consider the Duffing oscillator

$$\ddot{x} + 0.05\dot{x} + x + x^3 = F(t). \tag{19.12}$$

If $F(t) = 0$ and without the cubic member x^3 , Equation (19.12) represents a damped linear oscillator with a constant frequency and an exponentially decreasing amplitude. The presence of the nonlinear (cubic) member x^3 in Equation (19.12) leads to a time dependent frequency, considering again $F(t) = 0$, both the amplitude $A(t)$ and frequency $\omega(t)$ exponentially decrease and are correlated (Figs. 19.6a,b).




FIGURE 19.7. (a) A random driving force (a random walk with a few jumps). (b) A solution of the nonlinear Duffing oscillator with the random driving force $F(t)$ plotted in panel (a); (c) instantaneous amplitude (thick line) and frequency (thin line) extracted from the solution in panel (b).

Now, consider that the nonlinear oscillator (19.12) is driven by a random driving force $F(t)$. In the numerical examples presented here we consider a simple random walk with a few jumps as the driving force $F(t)$.

The relation between $A(t)$ and $\omega(t)$ is a nonlinear function and may vary in time, however, the level of the correlation between $A(t)$ and $\omega(t)$ depends on the driving force: With a relatively weak driving (Fig. 19.7a), $A(t)$ and $\omega(t)$ are almost perfectly correlated (Fig. 19.7c), with a stronger driving force $F(t)$ (Fig. 19.8a) some differences between $A(t)$ and $\omega(t)$ emerge; however, $A(t)$ and $\omega(t)$ are still correlated (Fig. 19.8c).

The amplitude-frequency correlation just presented is a property which can be tested in experimental signals, even in scalar cases (single measured time series). The instantaneous amplitude and phase of a signal $s(t)$ can




FIGURE 19.8. (a) Another example of a random driving force (“stronger”, i.e., with higher amplitude than in Fig. 19.7). (b) A solution of the nonlinear Duffing oscillator with the random driving force $F(t)$ plotted in panel (a); (c) instantaneous amplitude (thick line) and frequency (thin line) extracted from the solution in panel (b).

be determined by using the analytic signal concept of Gabor [5], recently introduced into the field of nonlinear dynamics within the context of chaotic synchronization in [20]. The analytic signal $\psi(t)$ is a complex function of time defined as

$$\psi(t) = s(t) + j\hat{s}(t) = A(t)e^{j\phi(t)}, \quad (19.13)$$

where the function $\hat{s}(t)$ is the Hilbert transform of $s(t)$

$$\hat{s}(t) = \frac{1}{\pi} \text{P.V.} \int_{-\infty}^{\infty} \frac{s(\tau)}{t - \tau} d\tau. \quad (19.14)$$

(P.V. means that the integral is taken in the sense of the Cauchy principal value.) $A(t)$ is the instantaneous amplitude and the instantaneous phase $\phi(t)$ of the signal $s(t)$ is

$$\phi(t) = \arctan \frac{\hat{s}(t)}{s(t)}. \quad (19.15)$$

The instantaneous frequency $\omega(t)$ is the temporal derivative $\dot{\phi}(t)$ of the instantaneous phase $\phi(t)$.

19.7 Amplitude-Frequency Correlation in the Sunspot Cycle

A possible amplitude-frequency correlation (AFC hereafter) in the sunspot cycle, in particular, the importance of the amplitude in determining the length of the related cycle was noted in the 1930s by [26] and recently discussed in [6]. In this section we demonstrate that the amplitude-frequency correlation found in the sunspot cycle is probably a non-random phenomenon and we propose its explanation by an underlying nonlinear dynamical system.




FIGURE 19.9. The instantaneous amplitude (thick line) and frequency (thin line) of (a) the yearly sunspot numbers series, (b) a realization of the related HAFT surrogate data, and (c) a realization of the Barnes model.

The series of yearly sunspot numbers from the period 1700 – 1997 (Fig. 19.4a) has been filtered by a simple moving average (MA) band-pass filter: First, the MA's from a thirteen-sample window have been subtracted from the data to remove slow processes and trends, and then a three-sample MA smoothing has been used to remove high-frequency components and noise. Then the discrete version of the Hilbert transform (19.14) using the window length of twenty-five samples has been applied to obtain the instantaneous amplitude $A(t)$ and the instantaneous phase $\phi(t)$. For obtaining a more robust estimation of the instantaneous frequency $\omega(t)$ than the one yielded by a simple differencing the phase $\phi(t)$, the robust linear regression [16] in a seven-sample moving window has been used. Finally, the series of $A(t)$ and $\omega(t)$ have been smoothed using a thirteen-sample MA window. The

resulting series of the instantaneous amplitude and frequency of the yearly sunspot numbers, plotted in Fig. 19.9a, yield the crosscorrelation equal to 0.505. Does this value mean that the amplitude and frequency of the sunspot cycle are correlated as a consequence of an underlying dynamics, or could this correlation occur by chance? Searching for an answer, we test the statistical significance of this correlation using the surrogate data approach.

In the first step we use the FT and HAFT surrogates, defined in Sec. 19.2, where these kinds of surrogates play the role of a linear stochastic process with the same spectrum and histogram as the studied data. Testing nonlinearity in general, it is stressed that the (HA)FT surrogates replicate the linear “properties” (more exactly, temporal dependences), while they do not contain any nonlinear dependence structure. Here, testing the significance of AFC, we consider the (HA)FT surrogates as data with cycles oscillating with the same frequencies as the sunspot cycles, but not possessing any systematic amplitude-frequency correlation. Since for generating the (HA)FT surrogates we use the Fast Fourier Transform (FFT) [16] which requires a number of samples equal to a power of two, we perform two tests, using the “first” and “last” 256 samples, i.e., the subseries of the whole 298 sample series obtained by cutting away 42 samples at the end or at the beginning, respectively, from the whole yearly sunspot number record. Thus, in each test, the surrogate data replicate the sample spectrum of the related 256-sample subseries.

The (HA)FT surrogates are generated from the raw (unfiltered) 256-sample segments of the sunspot data. Also, the 256-sample subseries are used for estimating the amplitude-frequency correlation related to the particular subseries, applying the procedures described earlier. Then, each realization of the (HA)FT surrogates, generated with respect to the raw data, undergoes the same processing as the raw data, i.e., the MA bandpass filtering, the Hilbert transform and the robust linear regression for the $\omega(t)$ estimation, and the final $A(t)$ and $\omega(t)$ smoothing are performed before computing the AFC for each surrogate realization. Then the *absolute* values of the AFC’s for 150,000 surrogate realizations are evaluated to assess the significance of the related AFC value found in the sunspot data. The first 256-sample subseries of the sunspot yearly number yields AFC equal to 0.605, while the mean value of the *absolute* AFC for the HAFT surrogate set is 0.26 with the standard deviation (SD) equal to 0.17.

In usual surrogate tests the significance is derived from the difference between the data value and the surrogate mean, divided by the surrogate SD, provided normal distribution of the surrogate values. Having generated a large amount of the surrogate replications, here we directly estimate the *p*-value of the test, i.e., the probability that the assessed correlation occurred by chance (randomly) within the chosen null hypothesis (surrogate model), by simply counting the occurrences in the surrogate set of absolute AFC values greater than or equal to the assessed raw data value, i.e., 0.605 in

this case. The number obtained is 3637, which is equal to 2.43 percent. Statistically speaking, the test result is significant on $p < 0.03$, or, in other words, the probability that the amplitude-frequency correlation found in studied segment of the sunspot data occurred by chance (as a random event) is smaller than 3 percent.

Processing the “last” 256-sample segment of the yearly sunspot number, the obtained AFC is equal to 0.532, while the values from the HAFT surrogates are the same as earlier, however, the p -value in this case is 6.58 percent. Still, we can conclude that the test result is significant on $p < 0.07$. An example of the HAFT surrogate realization is plotted in Fig. 19.4b, its instantaneous amplitude and frequency in Fig. 19.9b.

The results from the tests using simple FT surrogates (i.e., without the histogram adjustment) are practically the same as those from the HAFT surrogates. Testing the monthly sunspot numbers⁵, the segments of (“first” and “last”) 2048 samples were used. The same data processing has been applied as described earlier in the case of the yearly data with the windows lengths equivalent in real time, i.e., multiplied 12 by the number of samples. The obtained results are perfectly equivalent to those yielded by the yearly data, i.e., $p < 0.03$ and $p < 0.07$ for the “first” and “last” 2048-sample segments, respectively.

In the second step of testing the significance of the sunspot cycle AFC, we use realizations of the Barnes model as surrogate data. As noted earlier, the Barnes model mimics some important properties of the sunspot data; it is hard to reject it by standard nonlinearity tests, however, realizations of the Barnes model do not have any systematic amplitude-frequency correlation. A realization of the Barnes model is plotted in Fig. 19.4c, and its instantaneous amplitude and frequency are plotted in Fig. 19.9c.

In the test, 150,000 298-sample realizations of the Barnes model have been generated and processed in the same way as the sunspot series. The mean absolute AFC is equal to 0.21, $SD=0.15$; comparison with the AFC obtained for the whole 298-sample yearly sunspot series ($AFC=0.505$) yields the p -value equal to 4.36 percent. Thus, considering the Barnes model, the probability that the whole yearly sunspot series $AFC=0.505$ occurred by chance is $p < 0.05$.

19.8 Discussion

Recent development in the art of identifying nonlinear dynamics underlying experimental time series has led several authors to find ways of preventing a spurious detection of nonlinearity due to flawed surrogate data. In [12] we proposed testing whether linear dependences contained in the data

⁵Monthly sunspot numbers from the period 1749 – 1997 were also obtained from the Sunspot Index Data Center, <http://www.oma.be/KSB-ORB/SIDC>, file `monthssn.dat`.

under study were perfectly reproduced in the surrogates. Other authors have proposed sophisticated methods for constructing special constrained surrogate data. Nevertheless, we should be aware of the fact that a formal rejection of a linear stochastic null hypothesis does not automatically mean evidence for nonlinear dynamical origin of studied data. Looking for a more concrete explanation of a process under study, we propose to test specific features of nonlinear dynamical systems. Such a feature, present in nonlinear oscillators, is their amplitude-frequency correlation. Using the analytic signal approach and the Hilbert transform we can estimate the instantaneous amplitude and frequency even from scalar signals and test significance of obtained amplitude-frequency correlation using a proper null model, as demonstrated here in the case of the sunspot index series.

Using two different types of stochastic models (scaled isospectral surrogates and the Barnes model) which replicate some properties of the sunspot cycle, we have obtained a statistical support for the hypothesis that the amplitude-frequency correlation observed in the sunspot cycle did not occur by chance (as a random event) but is probably a property of an underlying dynamical mechanism. Well-known systems, possessing this property, are nonlinear oscillators, in which a significant AFC can be observed also in cases of external, even random, driving force. Therefore the presented results can be considered as a statistical evidence for a nonlinear oscillator (with an external, possibly random driving force) underlying the dynamics of the sunspot cycle, unless the amplitude-frequency relation is explained by a different mechanism.

Although no particular model for the solar cycle has been proposed here, the presented statistical evidence for a nonlinear dynamical mechanism underlying the sunspot cycle can be understood as a first step in bridging the gap between reliable statistical analyses of the experimental sunspot data (dominated by linear stochastic methods) and physical models such as nonlinear dynamo models [4, 9] (compared with data only on a qualitative level). In the next step, we can consider statistical comparison of the data with the dynamo models, with the aim of constructing a realistic data-driven model for the solar cycle. Another interesting problem is identification of a force driving the sunspot cycle.

In conclusion we state that the proposed method for identification in experimental time series of amplitude-frequency correlation, a feature of nonlinear oscillators, have brought interesting information about the sunspot cycle. The method can also be applied in different fields where time series generated by possibly nonlinear processes are registered and studied.

Acknowledgments

The author would like to thank J. Kurths, D. Novotná and I. Charvátová for valuable discussions and cooperation in research of solar and climate variability. This research was supported by the Grant Agency of the Czech Republic (grant No. 205/97/0921).

References

- [1] J.A. Barnes, H.H. Sargent, and P.V. Tryon. Sunspot cycle simulation using random noise. In R.O. Pepin, J.A. Eddy, and R.B. Merrill, editors, *The Ancient Sun*, pages 159–163, New York, 1980. Pergamon Press.
- [2] G. Cini Castagnoli and A. Provenzale, editors. *Past and Present Variability of the Solar-Terrestrial System: Measurement, Data Analysis and Theoretical Models*, Proc. Intl. School of Physics “Enrico Fermi”, Amsterdam, 1997. IOS Press.
- [3] M. Dämmig and F. Mitschke. Estimation of lyapunov exponent from time series: the stochastic case. *Phys. Lett. A*, 178:385–394, 1993.
- [4] U. Feudel, W. Jansen, and J. Kurths. Tori and chaos in a nonlinear dynamo model for solar activity. *Int. J. Bif. & Chaos*, 3:131–138, 1993.
- [5] D. Gabor. *J. IEE London*, 93:429, 1946.
- [6] D.H. Hathaway, R.M. Wilson, and E.J. Reichman. The shape of the sunspot cycle. *Solar Physics*, 151:177–190, 1994.
- [7] F. Krause and K.-H. Raedler. *Mean-Field Magnetohydrodynamics and Dynamo Theory*. Berlin, 1980.
- [8] M.N. Kremliovsky. Can we understand time scales in solar activity? *Solar Physics*, 151:351–370, 1994.
- [9] J. Kurths, U. Feudel, W. Jansen, U. Schwarz, and H. Voss. Solar variability: simple models and proxy data. In Cini Castagnoli and Provenzale [2], pages 247–261.
- [10] M.D. Mundt, W.B. Maguire II, and R.R.P. Chase. Chaos in the sunspot cycle: analysis and prediction. *J. Geophys. Res.*, 96(A2):1705–1716, 1991.
- [11] M. Paluš. Identifying and quantifying chaos by using information-theoretic functionals. In Weigend and Gershenfeld [27], pages 387–413.
- [12] M. Paluš. Testing for nonlinearity using redundancies: Quantitative and qualitative aspects. *Physica D*, 80:186–205, 1995.
- [13] M. Paluš. Chaotic measures and real-world systems. In H. Kantz, J. Kurths, and G. Mayer-Kress, editors, *Nonlinear Analysis of Physiological Data*, pages 49–66, Heidelberg, 1998. Springer.
- [14] M. Paluš, V. Albrecht, and I. Dvořák. Information-theoretic test for nonlinearity in time series. *Phys. Lett. A*, 175:203–209, 1993.
- [15] M. Paluš and D. Novotná. Testing for nonlinearity in weather records. *Phys.Lett. A*, 193:67–74, 1994.
- [16] W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling. *Numerical Recipes: The Art of Scientific Computing*. Cambridge Univ. Press, Cambridge, 1986.

- [17] D. Prichard and J. Theiler. Generalized redundancies for time series analysis. *Physica D*, 84:476–493, 1995.
- [18] M. Priestley. *Spectral Analysis and Time Series*. Academic Press, New York, 1981.
- [19] M.R.E. Proctor and A.D. Gilbert. *Lectures on Solar and Planetary Dynamios*. Cambridge Univ. Press, Cambridge, 1994.
- [20] M.G. Rosenblum, A.S. Pikovsky, and J. Kurths. Phase synchronization of chaotic oscillators. *Phys. Rev. Lett.*, 76:1804–1807, 1996.
- [21] L.A. Smith. The maintenance of uncertainty. In Cini Castagnoli and Provenzale [2], pages 177–246.
- [22] J. Theiler. Spurious dimension from correlation algorithm applied to time series data. *Phys. Rev. A*, 34:2427–2432, 1986.
- [23] J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, and J.D. Farmer. Testing for nonlinearity in time series: the method of surrogate data. *Physica D*, 58:77–94, 1992.
- [24] J. Theiler and D. Prichard. Constrained-realization monte-carlo method for hypothesis testing. *Physica D*, 94:221, 1996.
- [25] H. Tong. A personal overview of nonlinear time series analysis from a chaos perspective. *Scandinavian Journal of Statistics*, 22(2):399–445, 1995.
- [26] M. Waldmeier. *Astron. Mitt. Zürich*, 14(133):105, 1935.
- [27] A.S. Weigend and N.A. Gershenfeld, editors. *Time Series Prediction: Forecasting the Future and Understanding the Past*, volume XV of *Santa Fe Institute Studies in the Sciences of Complexity*, Reading, Mass., 1993. Addison–Wesley.
- [28] N.O. Weiss. *Philos. Trans. R. Soc. London, Ser. A*, 330:617, 1990.
- [29] G.L. Withbroe. *Spacecraft*, 26:394, 1989.
- [30] R. Wolf. *Acad. Sci. Comp. Rend.*, 35:704, 1852.