

Lecture Notes of the Institute
for Computer Sciences, Social Informatics
and Telecommunications Engineering

91

Editorial Board

Ozgur Akan

Middle East Technical University, Ankara, Turkey

Paolo Bellavista

University of Bologna, Italy

Jiannong Cao

Hong Kong Polytechnic University, Hong Kong

Falko Dressler

University of Erlangen, Germany

Domenico Ferrari

Università Cattolica Piacenza, Italy

Mario Gerla

UCLA, USA

Hisashi Kobayashi

Princeton University, USA

Sergio Palazzo

University of Catania, Italy

Sartaj Sahni

University of Florida, USA

Xuemin (Sherman) Shen

University of Waterloo, Canada

Mircea Stan

University of Virginia, USA

Jia Xiaohua

City University of Hong Kong, Hong Kong

Albert Zomaya

University of Sydney, Australia

Geoffrey Coulson

Lancaster University, UK

Patty Kostkova Martin Szomszor
David Fowler (Eds.)

Electronic Healthcare

4th International Conference, eHealth 2011
Málaga, Spain, November 21-23, 2011
Revised Selected Papers



Springer

Volume Editors

Patty Kostkova
Martin Szomszor
David Fowler

City University London, School of Health Sciences
City eHealth Research Centre
Northampton Square, London EC1V 0HB, UK
E-mail: patty@soi.city.ac.uk; {martin.szomszor.1, david.fowler.1}@city.ac.uk

ISSN 1867-8211 e-ISSN 1867-822X
ISBN 978-3-642-29261-3 e-ISBN 978-3-642-29262-0
DOI 10.1007/978-3-642-29262-0
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2012934259

CR Subject Classification (1998): K.4, I.2, J.3-4, H.4, C.2, H.5

© ICST Institute for Computer Sciences, Social Informatics and Telecommunications Engineering 2012
This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.
The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

It is our great pleasure to introduce the special issue of LNICST comprising the scientific presentations at the 4th ICST International Conference on eHealth (eHealth 2011) that took place in Málaga, Spain, during November 21–23, 2011.

Following the great successes of eHealth 2008 in London, 2009 in Istanbul, and 2010 in Casablanca, a specific aim of eHealth 2011 was research into social networks for epidemic intelligence, a new rapidly emerging field bringing together data mining, Web intelligence, social networks, and Web 2.0 technologies to assist public health and field epidemiology.

We selected 20 full and 8 short technical presentations by speakers from five continents. Revisited papers appear in these proceedings and cover a wide range of topics, including: social media analysis, personalization and patient support systems, early warning systems and mobile monitoring, games and learning, online support for professionals and patients, and agents in eHealth.

As in previous years, we were delighted to welcome two prestigious keynote speakers. Denis Coulombier, the head of unit for Surveillance and Response Support at ECDC, discussed epidemic intelligence from a European perspective, while S. Yunkap Kwankam delivered a talk on realizing the full potential of information and communication technology on global health.

In addition to our keynote speakers we invited a wide range of speakers from industry, hospitals, and public health organizations. These were: Johan Muskens (Philips Research Europe), Caroline Crousillat (Orange Healthcare), Jeremy Nettle (Oracle Corporation), Corinne Marsolier (Cisco Europe), Frederic Lievens (Med-e-Tel and ISfTeH), Anouk Berger (WHO), Nuria Oliver (Telefónica Digital), Josette Najjar-Pellet (Fondation Mérieux), Alberto E. Tozzi (Bambino Gesù Hospital, Rome, Italy), Rossana Alessandrello (Barcelona Digital Technology Centre) and Ciro Cattuto (ISI Foundation).

In the current economic situation it is essential to understand better the financial landscape supporting eHealth research, translational research, and health-care delivery. It has become ever more challenging to secure funding for eHealth and ICT services to improve national and global health. We created a special funding session at eHealth 2011 and were delighted to have Jenny Tooth (Angel Capital Group) as invited speaker to discuss these issues and challenges with the audience.

In a panel discussion chaired by Connie St. Louis entitled “Panic or Risk?”, Peter Openshaw from Imperial College and Ben Duncan of ECDC gave their views and answered audience questions on what health care agencies and the media can do to communicate risk without causing panic using the H1N1 outbreak in 2009 as an example. As a novel (and successful) experiment, both panellists took part by video conference, from the UK and Sweden, respectively.

Ed de Quincey and Gawesh Jawaheer chaired a joint posters and demo session at eHealth 2011, which was very well attended and provoked lively discussions. Twelve posters and eight demos were accepted for presentation. This year the session was “gamified” to encourage participation - in addition to awards for best poster and best demo, we also gave a prize for “best participant.”

Organized and chaired by Tim Meldrum on the first evening of the conference, the Speed Networking Session was an enjoyable and informal event to bring together medical practitioners, industry delegates, and eHealth researchers in a unique forum to explore the cutting edge of eHealth knowledge, technology, and practice.

This year’s edition of the conference made more use of social media than ever before. In addition to a conference Twitter account (@ehealthConf, with almost 600 followers at the time of writing) and hashtag (#ehealth2011), we had a Facebook page (www.facebook.com/eHealthConf, with 79 “fans”) and a Flickr poster gallery (www.flickr.com/photos/eddequincey/galleries/72157628028815537/, 374 views at the time of writing). We covered the conference using a live blog written by Gozde Zorlu, with descriptions of the presentations and other sessions (<http://electronic-health.org/>). The blog had over 400 readers, and attracted 375 comments and questions during the conference.

We would like to thank all those who helped make eHealth 2011 such a success - including the authors of all submitted papers, demos and posters; the speakers; keynote and invited speakers; the panellists; the members of the Program and Organizing Committees; the Session Chairs; and the conference participants (both in person and online). Special thanks go to the Local Chair, Cristina Urdiales, together with student volunteers and local organizers. Finally, we thank ICST, CREATE-NET, and IEEE EMB for sponsoring the conference, and Springer for publishing this LNICST book.

Patty Kostkova
Martin Szomszor
David Fowler

Organization

General Co-chairs

Patty Kostkova	City eHealth Research Centre, City University London, UK
Martin Szomszor	City eHealth Research Centre, City University London, UK

Clinical Chair

Femida Gwadry-Sridhar	Lawson Health Research Institute, Canada
-----------------------	--

Industry Chair

Corinne Marsolier	Cisco Systems Inc.
-------------------	--------------------

Knowledge Transfer Chair

Tim Meldrum	Enterprise Office, City University London, UK
-------------	---

Scientific Media Chair

Connie St Louis	Graduate School of Journalism, City University London, UK
-----------------	--

Social Media Chair

Marcel Winandy	Ruhr University Bochum, Germany
----------------	---------------------------------

Publications Chair

David Fowler	City eHealth Research Centre, City University London, UK
--------------	---

Poster Chair

Ed de Quincey	University of Greenwich, UK
---------------	-----------------------------

Demo Chair

Gawesh Jawaheer City eHealth Research Centre, City University
London, UK

Web Chair

Simon Hammond City eHealth Research Centre, City University
London, UK

Local Chair

Cristina Urdialis University of Malaga, Spain

IEEE EMB Representative

Laura Roa Romero University of Seville, Spain

Conference Coordinator

Justina Senkus European Alliance for Innovation

Senior Program Committee

Patty Kostkova City eHealth Research Centre, City University,
London, UK

Martin Szomszor City eHealth Research Centre, City University,
London, UK

Toni Moreno Universitat Rovira i Virgili, Tarragona, Spain

Floriana Grasso University of Liverpool, UK

Femida Gwadry-Sridhar Lawson Health Research Institute, Canada

Ulises Cortez Technical University of Catalonia, Spain

Rob Procter University of Manchester, UK

Albert Burger Heriot-Watt University, UK

John Powell University of Warwick, UK

Erik van der Goot Joint Research Centre, Ispra, Italy

Cecile Paris CSIRO ICT Centre, New South Wales,
Australia

Robert Stevens University of Manchester, UK

Ciro Cattuto ISI Foundation, Turin, Italy

Laurent Billonnet Université de Limoges, France

Technical Program Committee

Jan Aidemark	Linnéuniversitet, Sweden
Pedro Alexandre	
Da Costa Sousa	UNINOVA, Portugal
Dimitra Alexopoulou	The Biotechnology Center, Technische Universität Dresden, Germany
Margarita Anastassova	CREATE-NET, Trento, Italy
Bill Andreopoulos	The Biotechnology Center, Technische Universität Dresden, Germany
Roberta Annicchiarico	IRCCS Fondazione Santa Lucia, Italy
Raphael Bahati	The University of Western Ontario, Canada
Martin Beer	Sheffield Hallam University, Sheffield, UK
Ayse Bener	Bogazici University, Istanbul, Turkey
Elizabeth Borycki	University of Victoria, Canada
Miroslav Bursa	Czech Technical University, Czech Republic
Tito Castillo	University College London, UK
Olivier Corby	INRIA, France
Joanne Curry	University of Western Sydney, Australia
Foad Dabiri	University of California Los Angeles, USA
Ed De Quincey	University of Greenwich, London, UK
Danilo De Rossi	University of Pisa, Italy
Gayo Diallo	University of Bordeaux, France
Mark Donnelly	University of Ulster, Belfast, UK
Mariusz Duplaga	Jagiellonian University Medical College, Poland
Jonathan Elford	City University, London, UK
Henrik Eriksson	Linköping University, Sweden
Elisabetta Farella	University of Bologna, Italy
Anthony Fleury	École des Mines de Douai, France
David Fowler	City University, London, UK
Monique Frize	Carleton University, Ottawa, Canada
Leo Galway	University of Ulster, Belfast, UK
Xiaohong Gao	Middlesex University, London, UK
George Ghinea	Brunel University, Uxbridge, UK
Amir Hajjam-el-hassani	University of Technology of Belfort-Montbéliard, France
Simon Hammond	City University, London, UK
Ali Hamou	The University of Western Ontario, Canada
David Hansen	eHealth Research Centre Australia
Mark Hawker	University of Leeds, Leeds, UK
Anja Henner	Oulu University of Applied Sciences, Finland
Hermie Hermens	University of Twente, The Netherlands
Jesse Hoey	University of Dundee, UK
Alexander Hörbst	University for Health Sciences, Medical Informatics and Technology, Austria
Roberto Hornero	University of Valladolid, Spain

David Isern	Universitat Rovira i Virgili, Tarragona, Spain
Gawesh Jawaheer	City University, London, UK
Malina Jordanova	Solar-Terrestrial Influences Institute, Bulgaria
Simon Jupp	University of Manchester, UK
Anastasia Kastania	Athens University of Economics and Business, Greece
Sabine Koch	Karolinska Institutet, Stockholm, Sweden
Waldemar W. Koczkodaj	Laurentian University, Ontario, Canada
Peter Kokol	University of Maribor, Slovenia
Jorma Komulainen Ylilääkäri	National Institute for Health and Welfare, Finland
Patty Kostkova	City University, London, UK
Vaclav Kremen	Czech Technical University, Czech Republic
Elaine Lawrence	University of Technology, Sydney, Australia
Athina Lazakidou	University of the Peloponnese, Greece
Lenka Lhotska	Czech Technical University, Czech Republic
Panos Liatsis	City University, London, UK
Magí Lluch-Ariet	MicroArt, Barcelona, Spain
Paul Marshall	University of Warwick, UK
Maria G. Martini	Kingston University, London, UK
Kenneth Mcleod	Heriot-Watt University, Edinburgh, UK
Andrew Miller	University of Wollongong, NSW, Australia
Maria Mirto	University of Lecce, Italy
Fleur Mougin	Université Bordeaux, France
Bruce Moulton	University of Technology Sydney, Australia
Farid Naït-abdesselam	University of Paris Descartes, France
Sonja O'Neill	University of Ulster, Belfast, UK
Silvia Olabarriaga	Universiteit van Amsterdam, The Netherlands
Venet Osmani	CREATE-NET, Ireland
Julian Padget	University of Bath, UK
Daniela Paolotti	Institute for Scientific Interchange Foundation, Italy
Dorina Petriu	Carleton University, Ottawa, Canada
Joao Pimentao	UNINOVA, Portugal
Edwige Pissaloux	Institut des Systèmes Intelligents et de Robotique, Université Pierre et Marie Curie, Paris, France
George Polyzos	Athens University of Economics and Business, Greece
Alain Pruski	Lasc Université de Metz, France
Dietrich Rebholz-Schuhmann	European Bioinformatics Institute, Cambridge, UK
Blaine Reeder	University of Washington, Seattle, USA

Pekka Ruotsalainen	National Institute of Health and Welfare, Finland
Sandro Scattareggia	Marchese Signo Motus Srl, Messina, Italy
Andreas Schrader	University of Lübeck, Germany
Michael Schumacher	The Institute of Business Information Systems, Switzerland
Tacha Serif	Yeditepe University, Istanbul, Turkey
Gregor Stiglic	University of Maribor, Slovenia
Kåre Synnes	Luleå University of Technology, Sweden
Martin Szomszor	City University, London, UK
Adel Taweel	Kings College, London, UK
Manolis Tsiknakis	Centre for eHealth Technologies, Heraklion, Crete, Greece
Alexey Tsymbal	Siemens AG, Germany
Asli Uyar	Bogazici University, Turkey
Aïda Valls	Universitat Rovira i Virgili, Tarragona, Spain
Bert-jan van Beijnum	University of Twente, The Netherlands
Erik Van Der Goot	Joint Research Centre, Ispira, Italy
Mateja Verlic	University of Maribor, Slovenia
Nadine Vigouroux	IRIT, France
Vivian Vimarlund	Linköping University and Jönköping International, Sweden
Peter Weller	City University, London, UK
Tatjana Welzer	University of Maribor, Slovenia
Jana Zvarova	EuroMISE Center, UK and AV ČR, Prague, Czech Republic

Table of Contents

Session 1: Social Media Analysis

Social Media: A Systematic Review to Understand the Evidence and Application in Infodemiology	1
<i>Stacey Guy, Alexandria Ratzki-Leewing, Raphael Bahati, and Femida Gwadry-Sridhar</i>	

An Examination of the Behaviour of Young and Older Users of Facebook	9
<i>Darren Quinn, Liming Chen, and Maurice Mulvenna</i>	

Session 2: Knowledge Integration and EPR

An Agile Methodology for IHE-Enabled Deployments	17
<i>Bruno Alves and Michael Schumacher</i>	

Obstetric Medical Record Processing and Information Retrieval	26
<i>Miroslav Bursa, Lenka Lhotska, Vaclav Chudacek, Michal Huptych, Jiri Spilka, Petr Janku, and Martin Huser</i>	

Session 3: Personalisation and Patient Support Systems

Health Care Performance Monitoring Using an Inpatient Reference Process Model and Clinical KPIs	34
<i>Eva Gattnar, Okan Ekinci, and Vesselin Detschew</i>	

A Model for a Motivational System Grounded on Value Based Abstract Argumentation Frameworks	43
<i>Eugenio Di Tullio and Floriana Grasso</i>	

An Assessment of the Potential for Personalization in Patient Decision Aids	51
<i>Øystein Eiring and Laura Slaughter</i>	

Session 4: Early Warning Systems and Mobile Monitoring

An Analysis of Twitter Messages in the 2011 Tohoku Earthquake	58
<i>Son Doan, Bao-Khanh Ho Vo, and Nigel Collier</i>	

Web-Based Stress Management System Goes Mobile:
 Considerations of the Design of a Mobile Phone App 67
Åsa Smedberg and H el ene Sandmark

Use of SMS for Tsunami Early Warnings at a Table Top Exercise 76
Vincenzo Lagani, Vasilis Kontogiannis, Panagiotis Argyropaidas, and Catherine Chronaki

Session 5: Games and Learning

Seamless Evaluation of Interactive Digital Storytelling Games:
 Edugames4All 80
Patty Kostkova

Patient-Centered Care as a Learning Process 85
Jan Aidemark and Linda Asken us

Wireless Computer Games and Applications in the Medical Education
 Curriculum: Adventures in Pedagogy 93
Juanita Fernando and Nicolette Peters

Session 6: Security, Privacy and Prevention

The Process of Policy Authoring of Patient-Controlled Privacy
 Preferences 97
Thomas Trojer, Basel Katt, Thomas Schabetsberger, Richard Mair, and Ruth Breu

Modelling a User Authorisation and Data Access Framework
 for Multi-specialty Research Systems in Secondary Health Care 105
Ire Ogunsina, Sarah N. Lim Choi Keung, Lei Zhao, Gavin Langford, Edward Tyler, and Theodoros N. Arvanitis

Software Engineering-Inspired Approach to Prevention Healthcare 109
Victor Rentea, Andrei Vasileteanu, Radu Ioanitescu, and Luca Dan Serbanati

Session 7: Online Support for Professionals and Patients

Evaluation of a Web-Based Patient Portal for Chronic Disease
 Management 114
Stacey Guy, Alexandria Ratzki-Leewing, and Femida Gwadry-Sridhar

Cardio Online Reader/COR: A Web 2.0-Based Tool Aimed at Clinical
 Decision-Making Support in Cardiology 122
Vendula Pap ikov a and Miroslav Zvolensk y

Session 8: Agents in eHealth

- An Agent Based Pervasive Healthcare System: A First Scalability Study 128
Johannes Krampf, Stefano Bromuri, Michael Schumacher, and Juan Ruiz
- An Agent-Based Approach to Real-Time Patient Identification for Clinical Trials 138
Gareth Tyson, Adel Taweel, Simon Miles, Michael Luck, Tjeerd Van Staa, and Brendan Delaney

Session 9: Online Communities of Practice

- The Use of Social Bookmarking by Health Care Students to Create Communities of Practice 146
Ed de Quincey, Avril Hocking, Josephine O’Gorman, Simon Walker, and Liz Bacon
- Engagement in Online Medical Communities of Practice in Healthcare: Analysis of Messages and Social Networks 154
David Fowler, Martin Szomszor, Simon Hammond, John Lawrenson, and Patty Kostkova
- Towards Delivering Disease Support Processes for Patient Empowerment Using Mobile Virtual Communities 158
Bert-Jan van Beijnum, Pravin Pawar, Lamia Elloumi, and Hermie Hermens

Session 10: eHealth Solutions

- E-Health Readiness Assessment for E-Health Framework for Africa: A Case Study of Hospitals in South Africa 162
Alfred Coleman, Marlien E. Herselman, and Dalenca Potass
- The Physical Activity Loyalty Card Scheme: Development and Application of a Novel System for Incentivizing Behaviour Change 170
Ruth F. Hunter, Michael Davis, Mark A. Tully, and Frank Kee

Session 11: Social Media Surveillance

- Tracking Media Reports on the Shiga Toxin-Producing *Escherichia coli* O104: H4 Outbreak in Germany 178
Jens P. Linge, Jas Mantero, Flavio Fuart, Jenya Belyaeva, Martin Atkinson, and Erik van der Goot

Syndromic Classification of Twitter Messages	186
<i>Nigel Collier and Son Doan</i>	
Towards Spatial Description-Based Integration of Biomedical Atlases . . .	196
<i>Nurzi Juana Mohd Zaizi and Albert Burger</i>	
Author Index	205

Social Media: A Systematic Review to Understand the Evidence and Application in Infodemiology

Stacey Guy, Alexandria Ratzki-Leewing, Raphael Bahati,
and Femida Gwadry-Sridhar

Lawson Health Research Institute,
Commissioners Rd E. 801, N6C 5J1 London, Canada
stacey.guy@sjhc.london.on.ca
{alexandria.ratzkileewing, raphael.bahati}@lawsonresearch.com
femida.gwadry-sridhar@lhsc.on.ca

Abstract. Social media represents a new frontier in disease surveillance. Infoveillance allows for the real-time retrieval of internet data. Our objective was to systematically review the literature utilizing social media as a source for disease prediction and surveillance. A review of English-language conference proceedings and journal articles from 1999 to 2011 using EMBASE and PubMed was conducted. A total of 12 full-text articles were included. Results of these studies show the use of open-source micro-blogging sites to inform influenza-like-illness monitoring. These results inform recommendations for future research directions.

Keywords: social media, review, population surveillance, data mining.

1 Introduction

Today, the number of social media users continues to skyrocket with rates of participation on social networking sites already quadrupling from 2005-2009 [1]. As an easily accessible, highly cost-effective and interoperable system, social media opens doors to a better understanding of community creation, providing fast access to information anywhere in the world, 24 hours a day [2,3].

Due to the popularity of online communication, open-source social media platforms present excellent opportunities in health research [4]. Using a strategy called infoveillance, real-time online data can be systematically mined, aggregated and analyzed to inform public health and policy [5,6]. More specifically, social media can be used as a relevant and real-time source of epidemic intelligence [4].

Mining online information can provide insight to abnormal patterns of disease and aid in predicting disease outbreaks. Various studies have confirmed the potential of infoveillance to advance epidemic intelligence.

The aim of this paper is to illustrate how data generated through social media can be used to inform planning and implementation of strategies to address communicable disease emergence - in turn, changing the future of health research.

This paper is organized as follows. The methodology used in this systematic review is described in section 2. The results of the review are presented in section 3. A discussion of these findings and their implications are debated in section 4.

2 Method

2.1 Data Sources

We conducted a systematic review of the literature utilizing the bibliographic databases EMBASE and PubMed in June 2011. The following keywords were used to search EMBASE; keywords were divided into three categories: (1) *Disease (Early detect\$, Pandemic\$, Epidemic\$, Communicable disease\$, Early diagnosis)*, (2) *Medium (Information technol\$, Internet, Mass medium, Medical computing, Social Media, Social network\$, Geolocation)*, and (3) *Methodology (Disease surveillance, Monitor\$, Disease control, Algorithm\$, Data min\$, Query process\$, Information retrieval\$)*. MeSH terms used to conduct the search in PubMed were also divided into three categories: (1) *Disease (Pandemics, Communicable diseases, Disease outbreaks, Early diagnosis)*, (2) *Medium (Communications media, Databases, factual, Internet*, User-computer interface*)*, and (3) *Methodology (Population surveillance, Information storage and retrieval, Forecasting, Data mining, Sentinel surveillance)*. The keyword in which the article was indexed, the title of the article, and the abstract were searched for these categories described above. Unfortunately, terms such as ‘infoveillance’ have not yet been coded in these databases.

Search terms were chosen to reflect our objective – to review published research on disease surveillance using open-source social media. Conducting a generalized search in Google, elicited a few relevant publications to inform our search term selection. Index terms of these publications informed our formalized search strategy.

EMBASE and PubMed were chosen as they limit our search to published material in the arena of health care as opposed to using a computer-focused bibliographic engine that would require more terms to narrow the search to health. By including both databases we were able to search both North American and European published literature. The search was limited to both journal articles and conference proceedings – specifically, the terms ‘conference paper’ and ‘proceeding’ was used for EMBASE, while ‘clinical conference’ was used for PubMed) written in the English-language and published between 1999 and 2011 (for EMBASE, ‘current’ was chosen; for PubMed we used 1999/01/01 to 2011/12/01). The time period chosen reflects the addition of the term ‘internet’ to the PubMed index – no such MeSH term exists for social media however the term ‘internet’ is broad and encompassing.

2.2 Data Extraction and Synthesis

Using the search terms described above, a total of 384 journal articles and 1 conference proceeding was retrieved from PubMed and a total of 484 journal articles from EMBASE (with limits applied). In addition, our hand search (informed by the

references within the retrieved publications and Google searching) revealed 16 possible publications. The combined databases of 885 publications were searched for duplicates which resulted in a total of 287 publications being eliminated. Three of the authors (R.B, S.G, A.R-L) reviewed the publication abstracts to further establish relevance. Publications without abstracts were excluded. Publications that mined RSS feeds, survey data, physician records, medical records and search engines were excluded. Twelve publications focusing on data mining social media were reviewed.

3 Results

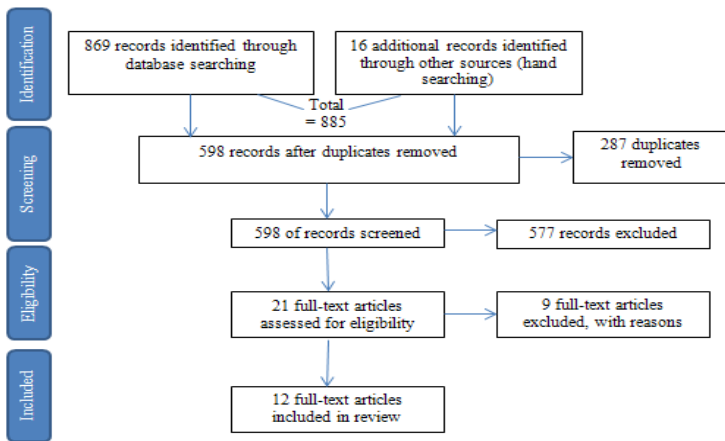


Fig. 1. This flow diagram illustrates the study selection process. A total of 885 publications were identified by bibliographic and hand searching. Through application of inclusion and exclusion criteria, 12 full-text publications were included in this review. This diagram is based upon The PRISMA Statement [7].

3.1 Study Characteristics

All publications reported focused on mining social media for the purpose of disease surveillance and prediction. As infodemiology is a relatively novel and emerging field, most studies were exploratory in design. The diseases of choice include influenza-like-illness (ILI) and H1N1.

3.2 Results of Individual Studies

Social Network Enabled Flu Trends (SNEFT) [8,9] is used to track and predict ILI activity. Tweets were retrieved (including metadata) with ILI content over approximately a 12 month period. SNEFT consists of separate data repositories where aggregated anonymous data is stored, and uses the ARMA model to predict ILI incidence. The web crawler was developed using Twitter's Search API. Regressive models were built and

evaluated with CDC (Centres for Disease Control) data. The system retrieved 4.7 million – retweets were extracted. To validate the observed trends, this data was compared to CDC data. The dataset containing no retweets and no tweets from the same user, resulted in the highest correlation ($r=0.9846$) with CDC data.

Corley et al., [10] evaluated blog posts containing ILI keywords. Flu-related posts were extracted from 44 million posts collected over a 3 month period. A seven day period in posting was identified and verified. Categorized baseline trends were compared to CDC data to identify anomalies. Results show a significant correlation between the frequency of ILI posts per week and CDC data. In addition, Corley et al. [4750 Corley 2010; 4028 Corley 2010;] collected 97,955,349 weblogs, micro-blogs and social media items pertaining to ILI data over a 20 week period. English language items containing relevant keywords were retrieved and grouped by month, week and day of the week. Flu-related data was compared to CDC data, and found to be highly correlated ($r=0.626$ at 95% confidence).

Culotta et al., [4] have developed regression models using 574, 643 tweets collected over a 10 week period to predict ILI. To obtain a random sample of tweets the authors searched for common words. The percentage of tweets that reported ILI was estimated. Their findings indicated multiple regression out performs simple regression, keywords selected based on residual sum of squares is more effective than selecting keywords based on a correlation coefficient, and the best model of prediction ($r=0.78$) was one where “a simple bag-of-words classifier trained on roughly 200 documents can effectively filter erroneous document matches”.

de Quincey & Kostkova [13] conducted a study to identify ILI trends present in tweets. The Twitter Search API was used to retrieve 100 tweets (including metadata). A PHP code parsed returned tweets (every minute) which were then saved to a MYSQL database. The system ran for 1 week in May 2009. A total of 135, 438 tweets containing ‘flu’ terminology were retrieved. The content was analyzed for trends using ‘flu’ in conjunction with other keywords (‘swine’ + ‘flu’). Future plans involve the use of collocation analysis to identify trends.

Infovigil is an open-source infoveillance system which mines, analyzes and visually represents textual health-related data from Twitter [6,14]. Infovigil was used to plot term prevalence, and provide content analysis of tweets pertaining to H1N1. Two million tweets were retrieved over an 8 month period. English-language tweets were selected for, and retweets were excluded. Tweet patterns were influenced by media with the most commonly tweeted material being news (52.6%) Original tweets, as opposed to retweets, contained more personal experiences. There was very little misinformation found in the tweets (4.5%). The majority of automated queries correlated with manual coding results.

Lamos & Cristianini [15] developed a monitoring tool to track ILI patterns using UK specific Twitter data. Tweets containing symptom-related keywords were collected over 6 months during 2009. A daily average of 160,000 tweets were retrieved. This data, converted to a flu-score, was compared to weekly H1N1 reports from the Health Protection Agency. This resulting score correlated highly with reports (>95%). This method works independently of language, can determine self-diagnostic statements in tweets, and uses time series geolocated data.

Data Collector [16] is a system that uses social media as a source for real-time data. The backend consists of a web crawler, written in PHP and utilizing the Twitter Search API. Tweets are stored in a relational database according to UML class with 2 main categories: disease and location. Data Collector supplies a RESTful API divided into location, disease, and occurrence, which specifies methods and parameters through which one can access the database. The frontend web interface uses AJAX to produce real-time graphs and maps. A dataset containing H1N1 tweets was collected from 6 European countries between May and July 2009. This system collected an average of 3200 tweets per day; 700 pertained to H1N1.

Signorini et al., [17] tracked sentiment and H1N1 activity using tweets. Keywords were used to retrieve tweets from the US. Prediction models were trained using CDC ILI values. Results are divided into 2 sets of data. The first saw 951,697 tweets from 334,840,972 retrieved over 34 days and was used to plot spikes of public interest. The second set represented over 4 million tweets from 8 million over a 3 month period. When analyzed, this showed no sustained interest in vaccine-related issues by the public. ILI estimates were gathered using a model trained on 1 million ILI tweets for 8 months. This model produced estimates that the authors believed were fairly accurate (average error = 0.28%; SD 0.23%). In order to garner real time estimates of ILI according to region, Signorini et al., [17] developed a model with region readings fitted to geolocated tweets. This model was less precise (average error = 0.37%) than the national weekly model.

The majority of articles included in the systematic review collected and examined data from Twitter (n=7); only 1 article looked at weblogs, micro-blogs, and social media. Studies demonstrated that a general correlation exists between ILI content in Tweets and CDC data. Research also indicated that Tweet patterns are strongly influenced by media, with news being the most commonly tweeted material. Finally, additional research is needed to determine the effectiveness of geo-location in garnering real-time estimates of ILI according to region.

4 Discussion

In realizing the potential of infodemiology in healthcare, it is important to consider the advantages and disadvantages of mining social media. Additionally, researchers must acknowledge and identify key target audiences; indeed, social media has specific target audiences with unique engagement behaviour specific to a platform that may or may not be representative of the population at large. Assessing the uses and potential of infodemiology in healthcare can improve user interaction, information access, evidence-based medicine, and knowledge representation.

4.1 Advantages and Disadvantages of Mining Social Media

Intelligence can provide real-time, immediate and relevant information [6,14]. This is particularly useful when seeking timely and reliable data on the spread or severity of influenza [18]. Analyzing and disseminating real-time information can also

improve public access to health surveillance information. As data-mining sources utilize open-source information, the operating costs of these systems can be extremely low [6,14]. In addition ‘mashups’ are the new multi-taskers, capable of mining, categorizing, filtering, and visualizing online, real-time data on epidemics [19].

Certain pitfalls to mining social media exist. First, textual data can be difficult to classify and interpret since harvested data (e.g., a tweet) may not provide enough information and meaning to facilitate automatic classification [17]. Second, the collected data may not be representative of the entire population—this challenge is especially pertinent as social media users are often younger, more educated, and urban-dwelling with higher incomes [17]. Furthermore, while coding for geographic origin may resolve certain limitations, not all profile accounts on networking sites contain geographic information; even so, visible geographic information cannot be verified for accuracy.

It is thus worthwhile to explore data mining sources that track IP addresses, or techniques to monitor social media activity on mobile phones [17]. Alternatively using GPS monitoring – using GIS systems that are either embedded in smart phones or attached to independent devices can provide supplemental information. One question that has not been resolved is whether participants need to provide explicit consent or whether the fact they are using publically available communication tools renders their information available and subsequently usable.

4.2 Study Limitations

Our search strategy was not as streamlined as we had planned. Not having the ability to choose a MeSH term as distinct as “social media”, “infoveillance”, or “Twitter”, meant that time was wasted sifting through irrelevant publications. While we extended our search to include conference proceedings as well, the publications retrieved were not of value. Rather, the references in retrieved publications provided direction to relevant crucial proceedings. However, some of the publications we retrieved through hand searching were not retrieved through direct bibliographic database searching.

4.3 Research Directions

In this section, we outline key recommendations that we believe are essential to forging new opportunities in data mining and collaborative analysis within user-driven content sharing paradigms. This will enable the full realization of the significant potential of patient engagement and information sharing and may help transform healthcare as a whole.

Target Audiences. The initial task for any text mining solutions is identifying entities of interest from the relevant textual content. This is often achieved through the use of natural language processing techniques. Each social media has a specific target audience with unique engagement behavior specific to that platform. Information harvested through social networking sites may not be representative of the population

at large. It is therefore important that data mining solutions take into account demographic characteristics of audiences within individual platforms.

User Interactions: Every social networking platform has a set of rules governing how its users interact with one another. For example, some social media platforms such as Twitter and Facebook enable real-time interactions between users while YouTube tends to be less interactive. As such, the types of interactions will determine when and how often data must be collected in order to derive any meaningful information.

Information Access: Social media platforms are already compiling fine-grained user-generated content based on individuals' online activities. While the means for deciphering what is relevant through information mining already exists and have proven extremely successful considering the amount of money companies are willing to pay to have such kind of access, such personal and information-rich content is not often publicly available. More open-source social networking solutions are therefore needed to facilitate any meaningful data mining solutions beyond the basic alert systems discussed in this paper.

Evidence-Based Medicine: Research is needed to identify effective ways of embedding evidence within social media platforms that could support monitoring positive impact on desired behavior changes and allowing users to share/compare experiences and provide support. Given the broad range of users, there is also a need to provide levels of detail regarding the evidence itself so that meaning information could be mined.

Knowledge Representation: Finally, mining social media content for medical information can only succeed if we recognize the role of ontologies in knowledge management and knowledge discovery. Ontology offers significant benefit to knowledge harvesting in social networking platforms as it facilitates data pruning and can help accelerate the discovery of meaningful information.

Summary. Evolving the use of infodemiology in healthcare will involve the examination of: (1) user interactions, which may determine the time and frequency of data collection; (2) information access, which may require the creation of more open-source social networking sites to facilitate more meaningful data mining solutions; (3) evidence-based medicine, which will allow users to share and compare experiences and provide support from within a social media platform; and (4) knowledge representation, which can allow the mining of social media for medical information, knowledge management and knowledge discovery. Health and healthcare might soon be achieved at the click of a mouse.

References

1. Evans, D.: *Social Media Marketing: An Hour a Day*. Wiley Publishing, Indiana (2008)
2. McNab, C.: What Social Media Offers to Health Professionals and Citizens. *Bulletin of the WHO* 87, 566 (2009)
3. Bacigalupe, G.: Is There a Role for Social Technologies in Collaborative Healthcare? *Families, Systems & Health* 29, 1–14 (2011)

4. Culotta, A.: Towards Detecting Influenza Epidemics by Analyzing Twitter Messages. In: 1st Workshop on Social Media Analytics (SOMA 2010), Washington DC (2010)
5. Wilson, N., Mason, K., Tobias, M., Peacey, M., Huang, Q.S., Baker, M.: Interpreting Google Flu Trends Data for Pandemic H1N1 Influenza: The New Zealand Experience. *Euro. Surveill.* 14, 19386 (2009)
6. Chew, C., Eysenbach, G.: Pandemics in the Age of Twitter: Content Analysis of Tweets During the 2009 H1N1 Outbreak. *PLoS ONE* 5 (2010)
7. Moher, D., Liberati, A., Tetzlaff, J., Altman, D.G., PRISMA Group: Preferred Reporting Items for Systematic Reviews and Meta-Analyses: The PRISMA Statement. *Annals of Internal Medicine* 151, 264–269 (2009)
8. Achrekar, H., Avinash, G., Lazarus, R., Yu, S., Liu, B.: Predicting Flu Trends using Twitter Data. In: International Workshop on CPNS, Shanghai, China (2011)
9. Chen, L., Achrekar, H., Liu, B., Lazarus, R.: Vision: Towards Real Time Epidemic Vigilance through Online Social Networks: Introducing SNEFT. In: 1st ACM Workshop on Mobile Cloud Computing, San Francisco, California (2010)
10. Corley, C.D., Mikler, A.R., Singh, K.P., Cook, D.J.: Monitoring Influenza Trends through Mining Social Media. In: International Conference on Bioinformatics and Computational Biology, Las Vegas, Nevada (2009)
11. Corley, C.D., Cook, D.J., Mikler, A.R., Singh, K.P.: Text and Structural Data Mining of Influenza Mentions in Web and Social Media. *International Journal of Environmental Research and Public Health* 7, 596–615 (2010)
12. Corley, C.D., Cook, D.J., Mikler, A.R., Singh, K.P.: Using Web and Social Media for Influenza Surveillance. *Advances in Experimental Medicine and Biology* 680, 559–564 (2010)
13. de Quincey, E., Kostkova, P.: Early Warning and Outbreak Detection Using Social Networking Websites: The Potential of Twitter. In: Kostkova, P. (ed.) *eHealth 2009*. LNCS, vol. 27, pp. 21–24. Springer, Heidelberg (2010)
14. Eysenbach, G.: Infodemiology and Infoveillance Tracking Online Health Information and Cyberbehavior for Public Health. *American Journal of Preventive Medicine* 40, S154–S158 (2011)
15. Lamos, V., Cristianini, N.: Tracking the Flu Pandemic by Monitoring the Social Web. In: Int. Workshop on CIP, Elba Island, Italy (2010)
16. Lopes, L.F., Zamite, J.M., Tavares, B.C., Couto, F.M., Silva, F., Silva, M.J.: Automated Social Network Epidemic Data Collector. In: *INForum* (2009)
17. Signorini, A., Segre, A.M., Polgreen, P.M.: The Use of Twitter to Track Levels of Disease Activity and Public Concern in the U.S. during the Influenza A H1N1 Pandemic. *PLoS One* 6, e19467 (2011)
18. Wilson, K., Brownstein, J.S.: Early Detection of Disease Outbreaks Using the Internet. *CMAJ: Canadian Medical Association Journal* 180, 829–831 (2009)
19. Brownstein, J.S., Freifeld, C.C., Madoff, L.C.: Digital Disease Detection—Harnessing the Web for Public Health Surveillance. *The New England Journal of Medicine* 360, 2153–2155, 2157 (2009)

An Examination of the Behaviour of Young and Older Users of Facebook

Darren Quinn, Liming Chen, and Maurice Mulvenna

School of Computing and Mathematics, University of Ulster,
Jordanstown. BT37 0QB, Northern Ireland
quinn-d15@email.ulster.ac.uk,
{l.chen,md.mulvenna}@ulster.ac.uk

Abstract. Increased use in recent years of social networking applications such as Facebook has resulted in a rapid growth of user generated content. As a communication tool for social engagement, the online patterns and characteristics of individuals and user types can now be better understood by monitoring user behaviour, activity frequency and modality of use. The aim of this paper is to present an analysis of the behaviour and characteristics of young and older user groups in an online social network. Our findings show that correlations exist between user behaviour and age, based upon activities as evidenced within the data. Visualisation of full user activity has further helped our understanding of user behaviour by identifying interaction patterns, which may be used for the monitoring of social engagement.

Keywords: Interaction Analysis, Social Networking, Social Computing.

1 Introduction

The growth and adoption of Internet technologies has increased markedly over the past decade. From 2000 to 2011, growth in Internet usage has witnessed an increase of 353% and as of March 2011, 58% of Europe's 816 million population were Internet users [1]. As a subsequent development of the Internet, online social networks emerged through a desire by users to share data, experiences, events and emotions. As a communication tool they performed a major role in the increased adoption rates of Internet technologies. As such a diverse range of social networking sites have been developed, from generic social networks such as Facebook¹, to specialist content sharing sites such as the photo sharing network of flickr². The impact of social networking has been acclaimed as rendering a new global footprint [2], due to its profound social and economic impact, on both business and individual users. With increased access to web services, and ever more innovative developments for end users, high profile social networking applications such as Facebook and Twitter³ are being routinely applied as a core communication approach for users desiring to share

¹ www.facebook.com

² www.flickr.com

³ www.twitter.com

information on a range of topics. An increased use of social networking sites has led to a rise in research interest, particularly within the domain of Social Network Analysis (SNA), a research area with a traditional role of analysis and visualisation of networks [3], discovering details such as a networks structure, or the strength and cohesiveness of community structures [4].

More recently within SNA a trend is emerging which is extending traditional concepts towards a more contemporary approach of Social Network Interaction Analysis (SNIA). As an example, Facebook is currently the most popular Social Networking Site with an estimated 700 million unique monthly visitors [5]. The aim of SNIA is to understand the "*individual*" and their use of functions in these vast networks. Through a variety of techniques such as web crawlers or bespoke programs, SNIA seeks to discover the behavioural patterns of each individual by analysing their user generated content. Contemporary studies have been interested in aspects such as user profiles, communication patterns and interactions in online networks [6], [7], [8]. As an approach it not only allows for the analysis of users and groups etc, but may also be applied to develop a greater understanding of the behavioural patterns and related characteristics of each user.

1.1 Aim

The aim of our current research is to harness and mine user generated content, through data collection approaches applied to Facebook, exploring the behaviour and characteristics of two distinct user groups, young (15-30) and older (50+) users, comprised of equal numbers of males and females. This study is made up in two parts; firstly we investigated user activity, defined as the posting of comments or replies to comments. As a study, its aim is to define the individual interaction metrics of any given user's activity. These measures are applied to compare interaction rates of *a user*, and across two *age categories*. Secondly we explore activity visualisation, to disclose interaction patterns. Illustrating user interaction data may help to understand previously undisclosed patterns, detailing the: *who*, *what* and *when* of a user's social network engagement, illustrating patterns such as network joining or the activity evolution of a user, defining their high and low periods of user activity.

The remainder of the paper is organised as follows. Section 2 discusses the concept, exploring in detail the chosen social network structure, defining interactions metrics and analysis methods applied in the study. Section 3 describes the user data and the data collection process, with results of activity frequencies and activity visualisation presented. Section 4 discusses the patterns that can be observed from the results. Section 5 summarises current work, discussing future work and limitations.

2 Concept and Methodology

The aim of the research is to develop an understanding into the individual activity of users. The research is aimed at determining individual metrics, and also to establish if correlations are identifiable between age groups and interaction frequency. Studies are

designed with the primary objectives of: quantifying user activity, visualising user engagement, identifying behavioural patterns and characteristics of disparate user groups, based on real world observations in an online social network. Metrics are subsequently applied to categorise users based on measures of activity, based upon frequency of posting comments and replies to comments. The primary motivation for investigating young and older online users is that it helps to establish if different behaviours (e.g., high or low activity patterns) are present and identifiable. Secondly the research may help to determine if these behaviours are attributable to the age of online social networks users. Future investigations will focus on establishing if indicators in the behaviour of people in online social networks relate to user well being (e.g., low activity frequency corresponding to low feelings of well being).

2.1 Methodology

"Facebook is a social utility that connects people with friends and others who work, study and live around them. People use Facebook to keep up with friends, upload an unlimited number of photos, post links and videos and learn more about the people they meet" [9]. Facebook was selected to investigate user's online social network engagement for the following reasons: prominence as the leading social network, user volume, data accessibility and a range of interaction features. Essentially Facebook users engage through their 'Wall', a facility which controls user content, such as for the posting of comments. As users contribute to their wall a chronicle of interactions is amassed. User's security options are self-imposed and dictate the accessibility of data by other Facebook network users. Privacy settings are set and controlled by the user, and information can be disclosed to: 'everyone', 'friends of friends' or 'friends only'. We class information as publically available if a user's privacy settings are set to everyone. It is only these profiles with which the study was concerned.

As is dependent on the information provided by the user, the age of a user is not always explicitly given. However, provided certain information is available it then becomes possible to determine a user's age (within a close proximity), in one or a combination of three ways, either by: date of birth, school leaving year or university leaving year. Interactions were defined within the context of this study as being any measurable user activity occurring on the user wall. A full range of potential Facebook interactions were assessed for their inclusion within the scope of the current study. However, the remit of the current study is to assess only the user's personal contribution to their network, therefore all non-user elements (e.g., non-user comments) were omitted. Due to the subjective nature of particular wall features our study concerned itself solely on user comments and user replies as core application components to provide a measurement of user contribution. It is viewed that later studies may provide an opportunity to provide a broader scoped analysis. As a specific example, the following comment scenario is provided; a user posts a comment and another user posts a reply in response to this comment, both interactions are recorded. However, only the first interaction will be of interest in the scope of this study as this is the only interaction attributable to the user in question. Corresponding metadata of date and time was also captured for each user element.

2.2 User Frequency

User metrics were derived using the following approach. Activity frequency was calculated as a measure of each individual's activity count, which was then divided over the length duration of available user data. In reality this means extraction of the first and last activity dates of any particular user, resulting in a value for the total number of active days. This value was then divided by the number of activities to obtain an overall frequency metric for the entire active period. Activity frequency calculation:

$$af = n / (x1 - x2) \quad (1)$$

Where; activity frequency (af), first date stamp (x1), last date stamp (x2), total number of user activities (n).

3 Data Collection and Trends

This section describes the data collection process, and provides specific details of the user data obtained for young and older users. The process was repeated in both instances altering only for the setting of profile ages. A user profile was created and using the '*find classmates*' tool, a seeding point of local universities was applied. Based on the defined profile attributes (age, gender, etc), a listing of potential friends were identified and returned by Facebook. User suitability for inclusion in the study was then determined on the two following issues; firstly a user's wall must be publically viewable; secondly age of the user must be identifiable from the stated user profile. Profiles of each suitable user's '*friends*' were then assessed using the same approach. Process was repeated until a sufficient amount for each user group and gender was acquired. Data was acquired over two months from December 2010, and subjected to a three stages of analysis process. (1) User walls were manually extracted through generated source code, containing all visible wall data. (2) A bespoke parsing program was developed for the identification and extraction of required user interaction data. (3) Comparison evaluations were applied for each user group and gender set. In total 500 full user profiles were obtained (125 young males, 125 young females, 125 older males and 125 older females). User data consisted of the entire wall history for each user, with walls exhausted until posts were no longer available. User age data was determined on a users '*info*' page, and extracted in conjunction with wall data.

3.1 Activity Frequency

Based on initial analysis, user groups were categorised into the following five activity groupings: '**G1**': Zero activity, '**G2**': Greater than 0 and less than 5 days, '**G3**': Greater or equal to 5 days and less than 10 days. '**G4**': Greater or equal to 10 days and less than 20 days. '**G5**': Greater than 20 days. These classifications are designed to represent increasing activity ranges, from G1 to G5. The first group 'G1' is the lowest possible activity value of zero, applied when no user activity was recorded on the user wall. As detailed the majority of users can then be subsequently categorised within

the remaining groupings of G2 to G5. The user comment frequency for both the older and younger users (Fig.1), and also the user reply frequencies for older and younger users (Fig.2) have been illustrated below. Further detailed is the activity frequency comparison table of young and older users (Table 1).




Fig. 1. User Comment Frequency

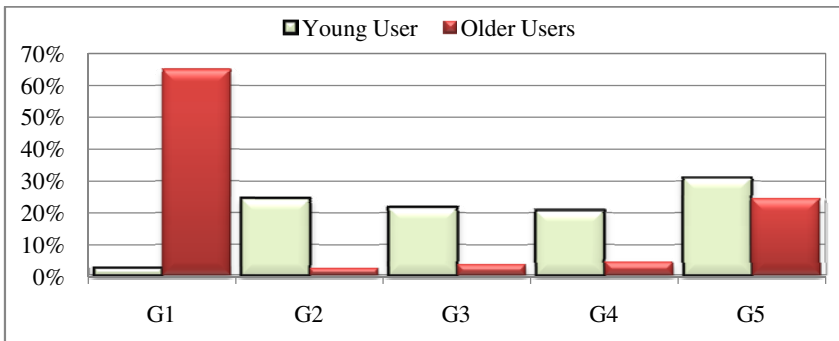


Fig. 2. User Reply Frequency

Table 1. Activity Frequency Table

Group	Comment Frequency (%)		Reply Frequency (%)	
	Younger	Older	Younger	Older
G1: Zero days	1 %	52 %	2 %	65 %
G2: 0 - 4 days	52 %	10 %	24 %	2 %
G3: 5 - 9 days	23 %	6 %	21 %	4 %
G4: 10 - 19 days	15 %	10 %	21 %	4 %
G5: Over 20 days	10 %	21 %	31 %	24 %

3.2 Activity Visualisation

In order to visualise user activity a 'representative user' was selected (Fig.3). Analysis of activities for both groups indicates that younger users generate greater volumes of

activity, therefore the younger cohort was chosen. Of the active groups, statistically 'G2' (Frequency values 0-5) was shown to be the most representative for younger users. In order to isolate and visualise a user, an individual was selected which contained the value closest to the average number of days value. The average number of active user days for younger users was 368 days. To ensure objectivity the foremost G2 young user to be found either side of the average number of days was selected. In this instance (whilst ensuring user privacy) we can reveal a number of user elements; the user was male, active for 370 days, retaining a frequency for their user comments of 2.03 and for their replies 3.06. As shown below (Fig.3):




Fig. 3. User Activity Chart

4 Analysis and Discussion

Analysis of comment activity frequencies discloses the following. Relating specifically to older users, the most dominant grouping was G1, containing 52% of all older users activity frequency values. As a group it identifies users whose walls contain no user activity. The second largest grouping was 'G5', identifying all user activity frequencies greater than 20, with 21%. A decreasing trend further continued through groups G4 and G3 with 10% and 6% respectively. 'G2' observed 10%. Within younger users notably only 1% of user values were contained within 'G1'. The largest younger user grouping was 'G2' retaining 52% of all younger frequency values. Younger user values declined sharply throughout the remainder of the groupings with; G3 containing 23%, G4 15% and G5 10%. Comparisons of the two groups show clear disparities are identifiable between the activity frequencies of young (Fig.4a) and older groups (Fig.4b). It can be shown that the majority of young users engage frequently, in contrast to older users. Older user engagement can now be said to be much less frequent with markedly lower levels of activity.




Fig. 4a. Younger Frequencies




Fig. 4b. Older Frequencies

As illustrated (Fig.3) in plotting user activity patterns are identifiable providing and understanding of a users habits. A total of 303 activities were observed made up of 182 user comments and 121 user replies. Within Fig.3 a series of groupings have been highlighted (A-C). Band A illustrates early social network engagement were a cluster of low volume recurring activity was recorded. A period of consistent activity can be seen over approximately the first 16 weeks from January 2010 to April 2010, were the majority of activity focused on user comments. Band B is identifiable as a period of heightened and concerted activity, were a significant increase in activity is visible, with a nucleus of activity between May 2010 and June 2010. Band C is identified as current behaviour in the final quarter of the year, with consistent activity. Spiking patterns are observed in this period resulting in increased user volumes. Given the length of engagement (370 days), it is viewed activity is non random and that patterns are evident. Notable however is the observation of low or non engagement periods, such as that of late July 2010 or early November 2010.

5 Conclusion

This study makes a number of contributions to the knowledge of online user interactions and user behaviour. This paper has shown that in the online social network of facebook.com the activity frequency at which younger and older users engage is varied. By harvesting user generated content, contained within a user's wall profile, the frequency at which a user engages can be determined. The development of activity frequency metrics now quantifies the frequency at which a user contributes to a social network. As is shown, such approaches can enable users to be classified based on their activity. Subsequently it is possible for the interaction rates to be compared across different users and against different user types in age groups. These results show that younger users can be classified as high frequency users, producing large volumes of online user data. Older users have been shown to engage less with a more infrequent online social network engagement. Through the plotting of a user's online history, results show that the behavioural patterns of a user can be observed, particularly in regard to high, low and non engagement.

6 Summary of Results, Future Work and Limitations

In determining the activity frequency of a user's data, and its subsequent visualisation, it has helped further the understanding of user's behaviour, detailing when and how users engage, identifying patterns of early, mid and current engagement. Importantly it defines high and low periods of activity. As with just one user example, it may be possible to hypothesize as to the reasoning of activity patterns. They are observed throughout the spring and early summer months as having a concerted period of activity. One tentative hypothesis is that behavioural patterns may have been stimulated by the reaction to increased real world social activities (holidays, etc.). Although conjecture for this particular user, patterns may realistically be applied to specific user types for monitoring social engagement, as early warnings relating to increasing social isolation of vulnerable individuals or groups. Online social networks have been shown as a means to track a user's social engagement levels, which may be applied in epidemiology for pattern anomaly detection.

It is envisioned that our future work in this area will extend current research, establishing if correlations exist between online activity and user well being. Limitations of current work are concerned with user data, specifically that the study is limited to only users that defined privacy settings to 'everyone'. It is hoped that a tool or process can be developed to overcome this issue, whilst acquiring a larger data set.

References

1. Miniwatts Marketing Group, Internet World Stats. Internet (March 2011), <http://www.internetworldstats.com/stats4.htm> (accessed July 12, 2011)
2. Charron, C., Favier, J., Li, C.: Social Computing, <http://www.forrester.com> (accessed April 02, 2010)
3. Lewis, K., Kaufman, K., Gonzalez, Christakis, M.: Tastes, ties and time: A new social network dataset using facebook.com. *Social Networks* 30(4), 330–342 (2008)
4. Porter, M.A., Onnela, J.P., Mucha, P.J.: Communities in networks. *Notices of the American Mathematical Society* 56(9), 1082–1097 (2009)
5. eBizMBA. Top 15 most popular Social Networking Sites. Internet, <http://www.ebizmba.com/articles/social-networking-websites> (accessed July 12, 2011)
6. Massari, L.: Analysis of MySpace user profiles. *Information Systems Frontiers* 12(4), 357–359 (2010)
7. Huang, L., Xia, L.: User Character and Communication Pattern Detecting on Social Network Site. In: International Conference on Engineering Computation, ICEC 2009, May 2-3, pp. 261–264 (2009)
8. Wilson, C., Boe, B., Sala, A., Puttaswamy, K.P.N., Zhao, B.Y.: User interactions in social networks and their implications. In: Proceedings of the 4th ACM European Conference on Computer Systems (Eurosys 2009), pp. 205–218 (2009)
9. Facebook. Facebook Press Room, About Facebook. Internet, <http://www.facebook.com/press.php> (July 12, 2011)

An Agile Methodology for IHE-Enabled Deployments

Bruno Alves and Michael Schumacher

University of Applied Sciences Western Switzerland, 3960 Sierre, Switzerland
`bruno.alves@hevs.ch`

Abstract. Recent history of healthcare software is littered with failures, mostly attributable to bad design and the inability to capture user community's needs and workflows[1]. This paper presents a new agile methodology aiming at improving engineering practises in Integrating the Healthcare Enterprise (IHE)-based projects in small-to-medium scale hospitals and healthcare organizations. The work described here is a compilation of experiences in the field and is based on the successful Scrum methodology, centered on IHE specifications that aims at providing value by making IHE developments easier and more predictable.

Keywords: Integrating the Healthcare Enterprise, IHE, security, methodology, planning, Scrum, interoperability.

1 From Hailing to Failing

The Therac-25 debacle is certainly the most serious case of computer-related accident to date involving patient death. A study published in [2] pointed to severe deficiencies in the engineering process. Ross Koppel revealed in an article published in 2005 [3] that *computerized physician order entry systems* (CPOE) can unexpectedly and contrary to conventional belief increase the number of medication errors, because they fail at capturing user requirements and because *"machine rules do not correspond to work organization or usual behaviors"*.

Often, responsibility for errors is not only attributable to a bad design, but also partly to the users, who are *"unlikely to accept blame for their own error or acknowledge their own inadequacies with respect to using the system"*[1]. In 1999, Kohn and his colleagues suggested in *To Err Is Human: Building a Safer Health Care System* [4] that healthcare professionals should increase awareness on how information technology could be applied to deliver safer care. Bad design, user misinformation and bad practises along with uncontrolled system growth and fast moving technology can actually lead to a new range of unexpected consequences. Awareness to the value of healthcare information exchange and interoperability[5] is probably the key to building better and safer systems.

Integrating the Healthcare Enterprise (IHE)¹ is an initiative led by professionals and the industry to facilitate information sharing in healthcare. IHE promotes the coordinated use of well established standards (such as DICOM² or HL7³) to address particular clinical needs in specifications known as Profiles described both in terms of actors and transactions in a Technical Framework[6]. Each Profile solves an unique clinical use-case (document sharing, patient management, ...), but can be integrated in more general framework. Profiles are systematic implementation guides that allow developers to focus more on the actual user experience and less on the integration problems, leading thus to less frustration for the developers and more satisfaction to the users. The amount of available Profiles however and the high number of elements to care about may be a bit overwhelming as more Profiles are integrated into one product. The problem can be largely mitigated with a careful and methodic development approach.

Healthcare software has usually a rigid set of fixed requirements and one may think development in healthcare does not require the agility that is essential in more business-oriented software. Healthcare IT requires high connectivity, a great number of overlapping subsystems and presents challenging integration issues, difficult to foresee. Development teams are expected to be extremely reactive to changes and may have to respond to notable modifications in the environment variables (timeframe, team, ...). Agile methods provide the kind of flexibility necessary to handle these cases, and especially Scrum.

Scrum [7] aims at delivering as much quality software as possible in small frames of time. It improves on existing engineering practises by involving frequent management activities tracking any deficiencies or impediments in the development process. Scrum particularly adapted to reduced-size regional organizations, because its process does not require large teams. And, since the customer is usually the organization itself, the proximity allows for a better communication and helps in better serving its needs.

The IheTools project is a follow up of the Medicoordination project published in [8],[9] and [10]. Its main objective is to take the exploration of IHE Profiles started with Medicoordination a step further. This exploratory research is necessary to gain enough knowledge so to be able to assist regional hospitals and healthcare organizations in their IHE-related tasks.

This paper presents a new methodology based on the successful Scrum that aims at assisting small-to-medium scale regional hospitals and healthcare organizations in their IHE-related development tasks, by providing a consistent and systematic approach on building up frameworks based on IHE Profiles. A motivation case will be first introduced; the methodology will be then presented and discussed at the end of this work.

¹ IHE International, <http://www.ihe.net>

² DICOM Homepage, <http://medical.nema.org/>

³ Health Level Seven International, <http://www.hl7.org/>

2 Motivation Case Study

2.1 On the Importance of a Good Plan

Our experience with IHE technologies was mostly acquired in the context of Medicoordination project. Although the project was successful in several aspects and the experiment was positively evaluated, the first prototype suffered from a few important design flaws that could have been avoided through a more careful design. This experience emphasized the importance of having good planning and a good methodology. Here are some experiences from the field:

As we came to learn, security is a topic that needs planning. It is in general a good practice to consider it already in the early stages of the product's development lifecycle. Although it may sound like a good idea to first develop an unsecured version and then add security to it, it is not. Security planning has many implications on the very foundations of the architecture and can lead to the creation of new unplanned breaches. IHE solves the problem partly, by providing a methodic and well documented security section in every Profile description. Common security concerns and elements to care about are well documented.

Furthermore, because healthcare software is usually provisionned with patient-centric sensitive data, not having a risk assessment and mitigation plan is unacceptable. Without it, the risk implications are not well understood, users not trained at its consequences and the final product may go well against local policies, against law or even against ethics.

Finally a mistake commonly found in research projects is to believe that a product will integrate flawlessly in the target operating environment. People often tend to believe there's no need to take the environment in consideration, even more in the early stages of the development. Murphy's Law⁴ states that *"if something can go wrong, it will eventually go wrong"*. Taking the operating environment into account in the planning allows avoiding most nasty unforeseen consequences.

2.2 Price Is What You Pay. Value Is What You Get⁵

Small-to-medium scale regional hospitals and healthcare organizations often have limited financial resources and assets. Although Return on Investment (ROI) is an important metric in decision making, the primary focus is on the value. Agile methods concentrate on the value by delivering functional iterations in short periods of time and also by placing the customer to the center of the discussion.

IHE can also provide value. A set of Open-Source IHE Profile implementations and frameworks was evaluated and a state of the art was published in [11]. All the selected implementations are free and IHE-compliant (most of them provide *compliance statements*) and may prove of good value for enterprises not willing to spend millions of euros in commercial licenses.

⁴ Murphy's Law, http://en.wikipedia.org/wiki/Murphy's_law/

⁵ Warren Buffet 1930, American Investment Entrepreneur.

2.3 The Waterfall Falls, The Scrum Scrums

Unlike the traditional waterfall development model, agile methods do not assume that the requirements will not change. On the contrary, they embrace change as a mean to constantly improve the software architecture. Scrum was first introduced




Fig. 1. The *Scrum* methodology process

in 1986 by Hirotaka Takeuchi and Ikujiro Nonaka [12] as a new approach to commercial product development. The primary objective was to increase the speed and flexibility of developments, by enforcing a single team with cross-functional skills to the whole process.

The Scrum methodology contains a set of predefined roles:

1. the "**ScrumMaster**" makes sure the process goes as intended
2. the "**Product Owner**" represents the customer
3. the "**Team**" works on the delivery of the product

The Scrum methodology does not make assumptions on the model used for the implementation, but rather limits itself to the description of how to drive the development. Scrum consists in three main phases: pre-game, development and the postgame.

The *pre-game phase* starts with a meeting to define a set or prioritized high-level requirements, called *product backlog list*. A thorough planning is achieved, including risk assessment, security consideration, project team and so on. In a second step, an architecture model is designed from the backlog items and problems, which may arise are identified and mitigated.

The *development phase* undergoes successive iterations, called *sprints* that can last up to 30 days until the release is ready for distribution. Each *sprint* starts with a meeting in which several items are taken from *product backlog list* to work on. Items are often splitted in smaller tasks and inserted into the *sprint backlog*. Some environment variables, such as timeframe, quality, resources and so on are constantly re-evaluated. The Scrum Master is tied to management activities to assess that no impediment will ever affect the ongoing development.

The *post-game phase* leads to a fully working release. The development enters this phase as soon as all the environment variables (requirements) are completed. The system is ready to be released and preparation includes tasks such as integration, user training and documentation.

3 Description of the IHE-Agile Methodology

This section describes the three phases of the Scrum methodology including elements from the IHE documentation.

3.1 Preparing the Game

The pre-game phase contains all the required groundwork before the actual development can start. The phase starts with a planning, in which the central element is the generation of items in the *product backlog* by the members of the whole team, including customers, sales and marketing division, software developers and a few other roles, that may vary from one project to the other. The initial planning generally includes elements from security, risk assessment, team foundation, training needs, integration and testing plans and many others that have been detailed in other works on agile development.

IHE documents state that a risk assessment and mitigation plan should be written for each profile. IHE provides a Security Considerations section for every Profile, which is based on the mitigations identified in each risk assessment. The Security Considerations section is not a thorough standalone security assessment, but just deals with issues specifically relevant to interoperability. There is no pre-cooked recipe for a risk assessment and mitigation plan, since risks usually vary from one project to the other. However, elements of the method used to write the Security Considerations section, which is described in the *IHE Cookbook: Preparing the IHE Profile Security Section*[13] can be used to help writing the global risk assessment plan. The method consists in identifying lists of risks, by imagining different scenarios and then assess their level of impact and probability of occurrence. Finally, mitigation of relevant risks for each profile is proposed.

Ethical and legal considerations need to be accounted for in the early stages of the development process. Ethical issues have to be well understood and communicated to the future users of the system. It should be clearly stated what is going to be done with sensitive data, who has access to it, what are the mechanisms to protect the data and what should be done in case of unauthorized disclosure. Legal aspects need to be analyzed transversely taking into account the laws of the country as well as the laws and policies of the lower levels (region, organization or business unit).

Security planning should be also partly done in the early phases of the development and improved during the subsequent *sprints* as the technical requirements become more and more clear. Some risk mitigation solutions proposed in the Technical Framework, must be analyzed and described in the security planning. IHE proposes some good articles on Profile-centric access control mechanisms[14]

and on the management of security and privacy. The *Template for XDS Affinity Domain Deployment Planning Handbook*[15] is usually used in planning the deployment of a XDS domain, but can also be used to plan organizational, operational and membership rules, as well as patient privacy and consent matters. System developers should not forget to attach the audit logging planning to the security section.

3.2 The Scrum Process

The subsections below describe the different phases of a single *sprint*. This is where the process shows agility. Indeed, in each *sprint*, environmental variables, such as requirements, technology and timeframes, are tracked and the development adapted if necessary.

Requirements. During the *Sprint Planning Meeting* punctuating the beginning of the next *sprint*, the Scrum Master, customers, users, Product Owner and Scrum team decide upon the goals and functionality. A second phase of the meeting gathers the Scrum Master and Scrum Team to decide how the product increment is going to be implemented and integrated into the existing product.

Analysis. The analysis phase includes elements such as security planning, usability and testing procedures.

Risk assessment must typically be done for every profile integrated in the final product, because different profiles generate different risks. Furthermore integrating with other Profiles generates even more risks.

Security planning must include all important information to help making security decisions. The team should be aware of topics such as actors, roles, authorizations, secure protocols, privacy, confidentiality and patient consent (involvement). It is important to define here what type of data is stored, where it is stored, who has access to it, for how long. This section may also contain emergency and bypass protocols as well as audit logs format and auditing placement.

Design. A typical IHE design should include low level aspects, such as IHE actors, transactions as well as security, interoperability annotations, communication protocols, encryption standards, protection mechanisms, role-based access controlling schemes and so on. All the links between the actors should be clearly identified and annotated with the type of security they require (mutual TLS over HTTP, for example) and with used protocols at both ends. For example, if one link is used to send data from a system that produces HL7 CDA R2 documents, it is a good idea to write HL7 CDA R2 next to the source actor. This kind of annotations helps making sure that two integrated systems are talking the same language and the same version.

Implementation. Demos are important, because they help them assessing the progress of an ongoing development. The implementation process may consist in writing modules or deploying existing frameworks, but the general rule is that implementation efforts should always result in output, which is visible and presentable at the end of the *sprint*. Customers do not live on promises, they want value. IHE projects often tend to mix both coding and deployments. It is thus important to limit the scope of the current *backlog item* in order to produce some visible output until the end of the next deadline, even if that means to split the item for the next iterations.

Testing. Unit testing is a great tool for assessing the correct behaviour of the iteration, even though it is far from sufficient. There are other key aspects that must be tested and validated such as: interoperability, security and usability.

Interoperability testing comes down to assessing that your system is using the right communication protocols and standard. Sometimes, using a different version of a same protocol, for example HL7 v2.5 over HL7 2.4, may invalidate your efforts. Interoperability validation ensures that messages sent by your piece of software are well understood by the surroundings. IHE Profiles inform about the expected outcome of a particular transaction. The software tests must use this data to make sure the current products' iteration behaves as expected.

Furthermore, IHE Profiles rely heavily on proven standards and are generally interoperable. It may happen, however, that some specific features or options are not implemented in one particular IHE Profile implementation. Hence, it is necessary to perform interoperability validation before proceeding to the next phase. There exist some IHE testing frameworks, such as the NIST XDS Test Suite⁶ or the now famous IHE Connectathons⁷, where teams can test their IHE-enabled products against others and possibly receive compliance statements.

Security testing is a delicate matter, since in most small projects, there is no *true* security expert. Testing procedures must assess the security of the software module for scenarios designed in the planning phase. For safety critical systems, availability has also to be tested, by putting systems under heavy charges and inserting deficient nodes.

Usability is one aspect that can make or destroy your project. It is important to involve user judgment early in the development lifecycle. IHE Profiles are designed to integrate user's workflow without much disturbances.

3.3 The Game's Isn't over Until It's over⁸

The post-game phase represents the end of a release. The project enters in the post phase upon agreement that all environment variables are completed. The system is typically finished and release preparation tasks such as integration and delivery are done at this step.

⁶ <http://ihexds.nist.gov/>

⁷ <http://gazelle.ihe.net/content/ihe-europe-2011-connectathon>

⁸ Yogi Berra, American professional baseball player and manager.

Integration. Integration is the operation, which consists in inserting the current release of the product in the operating environment. For implementations of typical IHE use-cases, IHE simplifies this integration, by providing guides and illustrations. Part of the integration work is about configuring the environment to line up with product's final specifications. Even through everything was planned in the pre-game phase, unexpected conditions may arise and a return in another *sprint* cycle may be necessary. The integration phase also includes extensive testing and quality assurance.

Delivery. At this point the product is ready to be delivered to the customer. In this phase, an extensive documentation has to be created and information procedures set up, including user training.

Postmortem: Lessons Learned. In almost all developments, some processes may go straight and others just don't. Recording in a written and accessible form what went wrong and what went right allows making continuous process improvements. The recording usually happens in the *Sprint Retrospective* meeting. Knowing what can adversely affect an ongoing development effort and how to respond to it, helps making future developments more predictable and comfortable for the team.

4 Discussion

The methodology presented above builds on the success of agile methodologies and on the rigor and consistency of IHE Profiles. It simplifies the development of projects in small-to-medium scale regional hospitals and healthcare organizations, by providing a guide about subjects that require a special care. Furthermore, this methodology also provides links to helpful IHE documents facilitating the process of writing down good risk assessment plans, creating SOA architectures based on IHE Profiles or preparing the security of the product.

This methodology is meant to be used as a supplementary tool and not as a all-in-one guide. It was created as a response to problems we had in our previous developments. No standard evaluation was performed using existing frameworks, which may restrain the scope of usability of this methodology to non-safety critical and non-business critical projects. However, concepts presented here are taken from our past experiences and may still prove valuable in order to avoid the same mistakes again and again.

5 Conclusion

We presented here a new agile methodology centered on IHE Profiles that allows assisting small-to-medium scale healthcare development tasks based on IHE. The methodology presented here builds on the agility principles of the successful Scrum agile method and is primarily intended to be used by small-to-medium scale regional hospitals and healthcare organizations where teams are sufficiently small and proximity is good.

References

1. Laplante, P.A., Nell, C.J., Sangwan, R.S.: Healthcare Professionals' Perceptions of Medical Software and What to Do About it. *J. IEEE Computer* 39, 26–32 (2006)
2. Leveson, G.S., Turner, C.S.: An Investigation on the Therac-25 Accidents. *Computer*, 18–41 (1993)
3. Koppel, R., et al.: Role of Computerized Physician Order Entry Systems in Facilitating Medication Errors. *JAMA* 9, 1197–1202 (2005)
4. Kohn, L.T., Corrigan, J.M., Donaldson, M.S.: *To Err Is Human: Building a Safer Health System*, p. 80. Nat'l. Academies Press (2000)
5. Walker, J., Pan, E., Johnston, D., Adler-Milstein, J., Bates, D.W., Middleton, B.: The value of health care information exchange and interoperability. *J. Health Affairs*, hlthaff-w5 (2005)
6. IHE IT Infrastructure (ITI) Technical Framework. IHE International, vol. 1–3, Revision 7 (2010)
7. Schwaber, K., Beedle, M.: *Agile software development with Scrum*. Prentice Hall PTR (2001)
8. Alves, B., Müller, H., Schumacher, M., Godel, D., Abou Khaled, O.: Interoperability prototype between hospitals and general practitioners in Switzerland. In: *Medinfo 2010*, Cape Town, South Africa, pp. 366–370. IOS Press (2010)
9. Müller, H., Schumacher, M., Godel, D., Abou Khaled, O., Mooser, F., Ding, S.: *MediCoordination: A practical approach to interoperability in the Swiss health system*. In: *The Medical Informatics Europe Conference (MIE 2009)* (2009)
10. Alves, B., Schumacher, M., Godel D., Richard, P. Abu Khaled, O., Müller, H.: Prototype d'interoperabilit entre hpitaux et mdecins traitants. In: *Actes de GISEH 2010*, Conference Francophone "Gestion et Ingnieurie des Systmes Hospitaliers", Clermont-Ferrand, France (2010)
11. Alves, B., Schumacher, M.: The Open-Source Integrating the Healthcare Enterprise (IHE). In: *Proceedings of the IADIS e-Health Conference*, Rome, Italy (2011)
12. Takeuchi, H., Nonaka, I.: The new new product development game. *J. Harvard Business Review* 64, 137–146 (1986)
13. IHE Cookbook: Preparing the IHE Profile Security. IHE International (2008)
14. Caumanns, J., Kuhlisch, R., Pfaff, O., Rode, O.: IHE ITI Infrastructure White Paper: Access Control. IHE International (2009)
15. IHE IT Infrastructure Technical Committee White Paper: Template for XDS Affinity Domain Deployment Planning. IHE International, Version 15 (2008)

Obstetric Medical Record Processing and Information Retrieval

Miroslav Bursa¹, Lenka Lhotska¹, Vaclav Chudacek¹, Michal Huptych¹,
Jiri Spilka¹, Petr Janku², and Martin Huser²

¹ Department of Cybernetics, Faculty of Electrical Engineering,
Czech Technical University in Prague, Czech Republic

miroslav.bursa@fel.cvut.cz

<http://bio.felk.cvut.cz>

² Obstetrics and Gynaecology clinic,
University Hospital in Brno, Czech Republic

Abstract. This paper describes the process of mining information from loosely structured medical textual records with no apriori knowledge. In the paper we depict the process of mining a large dataset of ~50,000–120,000 records × 20 attributes in database tables, originating from the hospital information system (thanks go to the University Hospital in Brno, Czech Republic) recording over 10 years. This paper concerns only textual attributes with free text input, that means 613,000 text fields in 16 attributes. Each attribute item contains ~800–1,500 characters (diagnoses, medications, etc.). The output of this task is a set of ordered/nominal attributes suitable for rule discovery mining and automated processing.

Information mining from textual data becomes a very challenging task when the structure of the text record is very loose without any rules. The task becomes even more difficult when natural language is used and no apriori knowledge is available. The medical environment itself is also very specific: the natural language used in textual description varies with the personality creating the record (there are many personalized approaches), however it is restricted by terminology (i.e. medical terms, medical standards, etc.). Moreover, the typical patient record is filled with typographical errors, duplicates, ambiguities and many (nonstandard) abbreviations.

Note that this project is an ongoing process (and research) and new data are irregularly received from the medical facility, justifying the need for robust and fool-proof algorithms.

Keywords: Swarm Intelligence, Ant Colony, Textual Data Mining, Medical Record Processing, Hospital Information System.

1 Introduction

1.1 Motivation

In many industrial, business, healthcare and scientific areas we witness the boom of computers, computational appliances, personalized electronics, high-

speed networks, increasing storage capacity and data warehouses. Therefore a huge amount of various data is transferred and stored, often mixed from different sources, containing different data types, unusual coding schemes, and seldom come without any errors (or noise) and omissions. Massively parallel distributed storage systems are used nowadays to provide computational nodes with data in reasonable time.

1.2 Nature Inspired Methods

Nature inspired metaheuristics play an important role in the domain of artificial intelligence, offering fast and robust solutions in many fields (graph algorithms, feature selection, optimization, clustering, etc). Stochastic nature inspired metaheuristics have interesting properties that make them suitable to be used in data mining, data clustering and other application areas.

With the boom of high-speed networks and increasing storage capacity of database clusters and data warehouses, a huge amount of various data can be stored. *Knowledge discovery* and *Data mining* is not only an important scientific branch, but also an important tool in industry, business and healthcare. These techniques target the problematic of processing huge datasets in reasonable time – a task that is too complex for a human. Therefore computer-aided methods are investigated, optimized and applied, leading to the simplification of the processing of the data.

Ant Algorithms. Ant colonies inspired many researchers to develop a new branch of stochastic algorithms: *ant colony inspired algorithms*. Based on the ant metaphor, algorithms for both static and dynamic combinatorial optimization, continuous optimization and clustering have been proposed. They show many properties similar to the natural ant colonies, however, their advantage lies in incorporating the mechanisms, that allowed the whole colonies to effectively survive during the evolutionary process.

1.3 Knowledge Extraction

Several techniques to extract knowledge from raw data have been developed in the past. These techniques have various and multiple origins: some result from the statistical analysis of the data, the regressions, decision trees, etc.; some resulting from the artificial intelligence such as the expert systems, intelligent agents, fuzzy logic, etc.

Text Extraction. The accuracy for relation extraction in journal text is typically about 60 % [1]. A perfect accuracy in text mining is nearly impossible due to errors and duplications in the source text. Even when linguists are hired to label text for an automated extractor, the inter-linguist disparity is about 30 %. The best results are obtained via an automated processing supervised by a human [2].

Ontologies have become an important means for structuring knowledge and building knowledge-intensive systems. For this purpose, efforts have been made to facilitate the ontology engineering process, in particular the acquisition of ontologies from texts.

2 Input Dataset Overview

The dataset consists of a set of approx. 50 to 120 thousand records (structured in different relational DB tables; some of them are not input, therefore the range is mentioned) \times approx. 20 attributes. Each record in an attribute contains about 800 to 1500 characters of text (diagnoses, patient state, anamneses, medications, notes, references to medical stuff, etc.). For textual mining, 16 attributes are suitable.

The overview of one small (in field length) attribute is visualized in Fig. [2]. Only a subsample (about 5 %) of the dataset could be displayed in this paper, as the whole set would render into a uncomprehensible black stain. The vertices (literals) are represented as a green circle, the size reflects the literal (i.e. word) frequency. Edges represent transition states between literals (i.e. the sequence of 2 subsequent words in a sentence/record); edge stroke shows the transition rate (probability) of the edge. The same holds for all figures showing the transition graph, only a different visualization approach has been used.

It is clear, that human interpretation and analysis of the textual data is very fatiguing, therefore any computer aid is highly welcome.

3 Graph Explanation

In this paper we describe *transition graphs*. These are created for each attribute. An attribute consists of many records in form of a sentence. By *sentence* we hereby mean a sequence of literals, not a sentence in a linguistic form. The records are compressed – unnecessary words (such as verbs *is*, *are*) are omitted. In this paper, only the attribute `delivery_anesthetics` is visualized, as it is the simplest one.

Vertices of the transition graph represent the words (separated by spaces) in the records. For each word (single or multiple occurrence) a vertex is created and its potency (number of occurrences is noted). For example, the words *mesocaine*, *anesthetics*, *not*, *mL* form a vertex. Note that also words as *mesocain*, *mezokain* and other versions of the word *mesocaine* are present. For a number (i.e. sequence of digits) a special literal `_NUMBER_` is used.

Edges are created from single records (sentences entered). For example the sentence *mesocaine 10 mL* would add edges from vertex *mesocaine* to vertex `_NUMBER_` and from vertex `_NUMBER_` to the vertex *mL* (or the edge count is increased in case it exists). For all records, the count of the edges is also useful. It provides an overview on the inherent structure of the data – the most often word transitions.

4 Motivation

The task of this work is to provide the researchers with a quick automated or semi-automated view on the textual records. Textual data are not easy to visualize. The word frequency is inappropriate, although it is very simple. Therefore we decided to extract information in the form of a transition graph.

Using these graphs a set of rules for information retrieval is being created (defined). These rules serve for extraction of (boolean) attributes from the textual rules. These attributes are used in automated rule discovery and can be further used for recommendation.

5 Nature Inspired Techniques

Social insects, i. e. ant colonies, show many interesting behavioral aspects, such as self-organization, chain formation, brood sorting, dynamic and combinatorial optimization, etc. The coordination of an ant colony is of local nature, composed mainly of indirect communication through pheromone (also known as *stigmergy*). The high number of individuals and the decentralized approach to task coordination in the studied species means that ant colonies show a high degree of parallelism, self-organization and fault tolerance. In studying these paradigms, we have high chance to discover inspiration concepts for many successful meta-heuristics.

5.1 Ant Colony Optimization

Ant Colony Optimization (ACO) [3] is an optimization technique that is inspired by the foraging behavior of real ant colonies. Originally, the method was introduced for the application to discrete and combinatorial problems.

Ant Colony Methods for Clustering. Several species of ant workers have been reported to form piles of corpses (cemeteries) to clean up their nests. This aggregation phenomenon is caused by attraction between dead items mediated by the ant workers. The workers deposit (with higher probability) the items in the region with higher similarity (when more similar items are present within the range of perception). For example, the *Messor sancta* ants organize dead corpses into clusters; brood sorting has been studied in ant colony of *Leptothorax unifasciatus*.

This approach has been modeled in the work of Deneubourg et al. [4] and in the work of Lumer and Faieta [5] to perform a clustering of data.

Methods using pheromone also exist, namely A²CA [6]. Another approach can be found in the work of J. Handl in [7] (an ATTA algorithm), which introduces modified neighborhood function (penalizing high dissimilarities), short-term memory with lookahead (jumping ants), increasing radius of perception, time-dependent modulation of the neighborhood function.

ACO_DTree Method. The ACO_DTree method is a hybrid evolutionary approach for binary decision tree construction [8,9]. The tree is induced using the known data and it can be used for unsupervised clustering later: each leaf of the classification tree can be interpreted as a cluster. The algorithm uses a population of classification trees that is evolved using an evolutionary approach. Creation of the trees is driven by a pheromone matrix, which uses the ACO paradigm. The high number of individuals and the decentralized approach to task coordination in the studied species means that ant colonies show a high degree of parallelism, self-organization and fault tolerance.

This approach has been utilized (with improvements) to simplify the structure of the vast dataset by finding the most important state transitions between literals, producing a probabilistic transitional model. The output structure is presented to the analyst for further processing/iteration.

For clustering, the ACO_DTree method [8,10] and ACO inspired clustering [5] variations have been successfully used.

6 Automated Processing

Automated layout of transition graph is very comfortable for an expert, however the contents of the attribute is so complicated, that a human intervention is inevitable. Examples of automated layout can be seen in Fig. [1] and Fig. [2].

The figure Fig. [1] shows a transitional graph where only positioning based on the word distance from the sentence start is used. Although it might look correct, note that the same words are mispositioned in the horizontal axis.




Fig. 1. A fully automated transition graph showing the most important relations in one textual attribute. No clustering has been used. The layout is based on the word distance from the start of the sentence. Note the mis-alignment of the similar/same words. Refer to section [2].

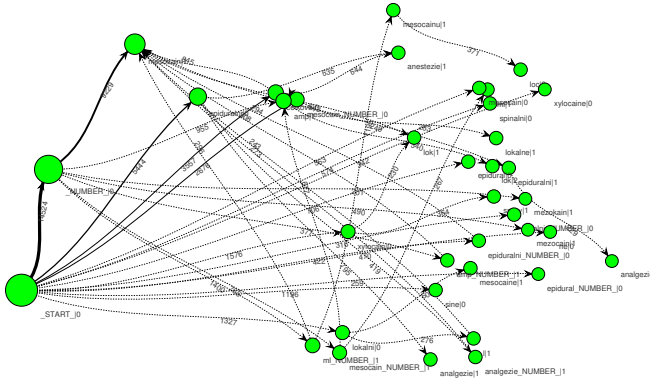


Fig. 2. A fully automated transition graph (sub-graph) showing the most important relations in one textual attribute. The ACO approach has been used to cluster the corresponding vertices. Refer to section [2].

7 Expert Intervention

A human intervention and supervision over the whole project is indiscutable. Therefore also human (expert) visualization of the transition graph has been studied.

The vertices in a human-only organization are (usually) organized depending on the position in the text (distance from the starting point) as they have the highest potency. Number literal (a wildcard) had the highest potency, as many quantitative measures are contained in the data (age, medication amount, etc.). Therefore it has been fixed to the following literal, spreading into the graph via multiple nodes (i.e. a sequence *mesocain 10 mL* become two vertices – *mesocain_NUMBER_* and *mL*). This allowed to organize the chart visualization in more logical manner. Time needed to organize such graph was about 5–10 minutes. The problem is that the transition graph contains loops, therefore the manual organization is not straightforward.

An aid of a human expert has been used in semi-automated approach (see Fig. [3] where the automated layout has been corrected by the expert. The correction time has been about 20–30 seconds only.

8 Results and Conclusion

The main advantage of the nature inspired concepts lies in automatic finding relevant literals and group of literals that can be adopted by the human analysts and furthermore improved and stated more precisely. The use of induced probabilistic models in such methods increased the speed of loosely structured textual attributes analysis and allowed the human analysts to develop lexical analysis grammar more efficiently in comparison to classical methods. The speedup (from

References

1. Freitag, D., McCallum, A.K.: Information extraction with hmms and shrinkage. In: Proceedings of the AAAI Workshop on Machine Learning for Information Extraction (1999)
2. Lafferty, J., McCallum, A., Pereira, F.: Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In: Proceedings of the ICML, pp. 282–289 (2001); Text processing: interobserver agreement among linguists at 70
3. Dorigo, M., Stutzle, T.: Ant Colony Optimization. MIT Press, Cambridge (2004)
4. Deneubourg, J.L., Goss, S., Franks, N., Sendova-Franks, A., Detrain, C., Chretien, L.: The dynamics of collective sorting robot-like ants and ant-like robots. In: Proceedings of the First International Conference on Simulation of Adaptive Behavior on From Animals to Animats, pp. 356–363. MIT Press, Cambridge (1990)
5. Lumer, E.D., Faieta, B.: Diversity and adaptation in populations of clustering ants. In: From Animals to Animats: Proceedings of the 3th International Conference on the Simulation of Adaptive Behaviour, vol. 3, pp. 501–508 (1994)
6. Vizine, A.L., de Castro, N.L., Hruschka, E.R., Gudwin, R.R.: Towards improving clustering ants: An adaptive ant clustering algorithm. *Informatica* 29, 143–154 (2005)
7. Handl, J., Knowles, J., Dorigo, M.: Ant-based clustering and topographic mapping. *Artificial Life* 12(1), 35–61 (2006)
8. Bursa, M., Lhotska, L., Macas, M.: Hybridized swarm metaheuristics for evolutionary random forest generation. In: Proceedings of the 7th International Conference on Hybrid Intelligent Systems 2007 (IEEE CSP), pp. 150–155 (2007)
9. Bursa, M., Lhotska, L.: Ant colony cooperative strategy in biomedical data clustering. *SCI*. Springer, Heidelberg (2007)
10. Bursa, M., Huptych, M., Lhotska, L.: Ant colony inspired metaheuristics in biological signal processing: Hybrid ant colony and evolutionary approach. In: *Biosignals 2008-II*, vol. 2, pp. 90–95. INSTICC Press, Setubal (2008)

Health Care Performance Monitoring Using an Inpatient Reference Process Model and Clinical KPIs

Eva Gattnar^{1,2}, Okan Ekinici¹, and Vesselin Detschew²

¹ Clinical Competence Center Cardiology, Siemens Healthcare
Allee am Röhelheimpark 3a, 91052 Erlangen, Germany
{eva.gattnar, okan.ekinci}@siemens.com

² Institute of Biomedical Engineering and Informatics, Ilmenau Technical University
POB 100565, 98684 Ilmenau, Germany
vesselin.detschew@tu-ilmenau.de

Abstract. Health care processes are nowadays heavily dependent on Information Technology (IT). On the other side, assessing a higher quality of health care to patients has become increasingly important. Effective health care processes are thereby heavily dependent on a comprehensive IT-support. Therefore, a thorough understanding about both the hospital IT-systems and clinical pathways is required, to identify quality problems. Improving quality of care through process performance measurement in hospitals and the identification of bottlenecks in performed workflows is thereby a promising concept. On this account we propose a novel approach based on an inpatient reference process model aiming a standardized and automated retrieval of quality and performance metrics along the clinical pathways of time-critical diseases in the context of various clinical standards and systems.

Keywords: performance monitoring, clinical key performance indicators, patient process, clinical workflow, clinical process modeling, events.

1 Introduction

With the enormous impact of Information Technology (IT), there is a major demand for standardization in health care nowadays. On the other side, assessing a higher quality of health care to patients has become increasingly important. Effective health care processes are thereby heavily dependent on a comprehensive IT-support. Therefore, a thorough understanding about both the hospital IT-systems and clinical pathways is required, to identify quality problems.

In recent years, hospital performance assessment projects and health care quality reporting have become therefore common worldwide [1], resulting in numerous initiatives aiming the development of clinical quality indicators to measure health care quality including hospital care performance measurement [2]. Improving quality of care through process performance measurement in hospitals and the identification of bottlenecks in performed workflows is thereby a promising concept. Process performance measures can assist hospitals in assessing their workflows and identifying areas for improvements in the field of diagnostics and treatment.

Several health care organizations recognized the importance of using performance metrics to monitor the clinical workflows. However, there is a lack of approaches aiming also the “real-time” end-to-end patient process monitoring, which can provide the possibility, to assist the clinicians in improving the quality of care already during the diagnostics and treatment and not only afterwards. Event-based approaches using workflow-events are a promising opportunity to measure the process performance and to monitor the clinical pathways, to improve the quality of care and benchmark the healthcare organization. On this account we propose a novel approach based on an inpatient reference process model aiming the standardized and automated retrieval of quality and performance metrics like Key Performance Indicators (KPIs) along the clinical pathways of time-critical diseases in the context of various clinical standards and systems. Beside the definition of the new model we develop the required clinical quality metrics systematically and analyze the hospital IT-systems, devices and communication standards regarding the process-based extraction of timestamps. In this paper we describe the conditions behind the approach and the results achieved.

2 Health Care Process Performance Monitoring

Processes are significant components within a health care organization. Process orientation is among others the most important method to support patient-oriented high quality of care. The economical benefit of recognized, well-defined, optimized, IT-supported and hence resulting mature processes is approvable too. Following the international Detecon study [3], a higher maturity of an organization’s processes results directly in an improved Return on Equity (ROE) as well as in an improved Return on Assets and Return on Investment. The resulting positive financial and quality-based effects are exactly measurable [4]. Consequently, to be successful and to provide an efficient and effective diagnostics and treatment, health care organizations have to apply optimization approaches aiming a higher process maturity and improved process performance [4][5].

Health care performance correlates often directly with process performance [5]. Therefore, processes should be applied as an important basis for health care quality improvement. Process models can help to analyze the process quality using structured process descriptions to allow process assessments. Secondly, they provide the basis for improvement steps to achieve a better process performance and in the following a higher quality of the processes within the health care organization. In general, process performance monitoring initiates and supports process improvement within an organization [5].

Thereby, hospitals need reliable and valid indicators for benchmarking and process improvement [6][7]. Following [6], health care process monitoring is impossible without using clinical indicators. Various types of such indicators can be either generically defined to support the measurement of several diseases or specifically defined to provide disease-specific measures [6]. Thereby, process-related indicators enable the performance measurement of the activities and tasks in patient’s “episodes of care” [8]. Based on the performance measurement theories and methodologies derived from industry, they are commonly called Key Performance Indicators (KPIs).

In particular, the development of reliable and valid clinical indicators requires a rigorous scientific design, definition and implementation [9]. Numerous organizations, such as the Institute of Medicine (IOM) [10], accrediting organizations (JCAHO, NCQA), government healthcare research and quality agencies (AHRQ, CMS) and various research groups devise and publish clinical indicator sets for measuring purposes in various domains of quality of care [11]. However, not all of these indicators are considered by physicians or hospitals to be valuable or relevant in their practice. Furthermore, health care systems are complex and therefore indicators may not reflect accurately all aspects of quality [12]. Another problem is that their development is frequently driven by feasible availability of data or ease of measurement, rather than by an objective definition of important clinical indicators [5]. Consequently it is essential to define indicators involving all stakeholders like physicians and IT-experts and additionally to consult approved guidelines, to ensure the indicator relevance, reliability, reproducibility and validity [11] and provide the possibility to develop both generic and disease-specific indicators.

We focus on clinical quality of care indicators and relate to the processes, which are performed in the care of time-critical diseases. The usage of time-based indicators for quality assessment in time-critical workflows represents a possibility for measurement, documenting and improving the quality of care [13][15]. Aiming the collection of time-based indicators in hospitals required IT-timestamps have to be retrieved. IT-systems act in our approach as enabler for performance monitoring, clinical quality improvement as well as risk and cost reduction [14][16]. The involved IT-systems are the Hospital Information System (HIS) and in the radiology department: the Radiology Information System (RIS) and imaging devices (Modalities). Outside the radiology department the Laboratory Information System (LIS) is used. Because several systems and communication standards (DICOM and HL7) (see Fig.1) are involved in the clinical workflow, we developed a new inpatient process model to enable performance monitoring across several clinical systems [16].

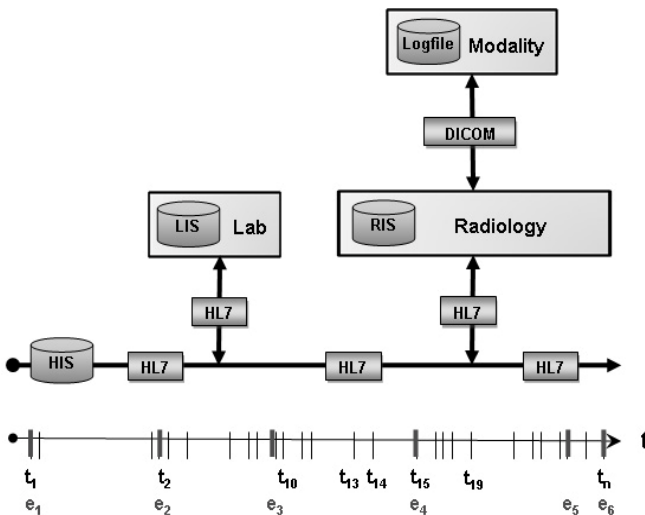


Fig. 1. IT-systems and clinical standards involved in the diagnostic workflow in hospitals

Fig.1 shows the involved IT-systems and the timeline during the diagnostics of an inpatient (i.e. a patient who stays at hospital for diagnostics and treatment). Additionally, timestamps (t) and events (e) are exemplary marked in the figure. Event (e1) means the patient's admission at hospital. The two following events mark the beginning and the end of the laboratory procedure, i.e. the KPI "order to lab test". The fourth and fifth events allow the imaging procedure monitoring at a modality (i.e. imaging device). Here the KPI "order to imaging" can be determined. Thereby, the timestamps, these events are based on, are extracted from the logfiles stored at the modalities. Using modality logfiles for monitoring purposes is a novel approach we introduce to enable process monitoring at the lowest process level. Finally, the last event indicates the finishing of the diagnostic process flow. The difference between the first and last event timestamps provides the KPI "door to diagnosis", which is very important in time critical diseases like heart attack and stroke [17].

3 Results and Validation

We developed a new generic inpatient reference process model to provide the possibility of standardized and comprehensive performance process monitoring in hospitals in the field of time-critical diseases. In order to monitor the process cycle times, health care organizations are up to date faced with the problem of extracting IT-timestamps from several not well-connected IT-systems, which are not related to a specific inpatient process. Consequently, it is not possible to monitor the whole end-to-end inpatient process-flow starting with the patient's admission at hospital until the patient's discharge. Secondly, valid and reliable performance monitoring requires standardized, meaningful and well-defined timestamps and events, which have to be monitored in order to measure the process cycle times. We provide these possibilities and fulfill both requirements bridging the gap between the IT and the clinical context proposing our model and KPIs [16]. As the definition of such events and the indicators based on them is very time-consuming and difficult [13], a predefined generic approach - as we present it - is desirable by the stakeholders. It can be used already in an early phase during the assessment in hospital to analyze the pathways and to define the measure points based on the IT-system architecture [14].

The first model development phase comprised the construction of our newly clinical reference process model using events, functions, connectors, control flows and process path elements. We selected the relevant process tasks based on interviews we performed with sophisticated clinicians. This group was built up of physicians like cardiologists, neurologists and radiologists who have huge and long-time experience in hospitals in heart attack and stroke treatment. Individually observations of the workflows during diagnostics and treatment of heart attack and stroke patients in several hospitals built up the second information source. After finishing the modeling tasks the physicians were consulted again in order to approve the model, the acceptance of the chosen modeling technique as well as the clinical correctness. In further steps we enhanced the model analyzing clinical pathways, as they are used in best practice hospitals. "Best practice" means in this context, that the considered hospitals are best in treating heart attack resp. stroke regarding quality of care.

In the following, we conducted interviews with involved physicians again, to derive relevant, valid, reliable and feasible clinical indicators for process monitoring. Additionally we consulted clinical guidelines for stroke and heart attack treatment. We derived sets of time-based indicators following [18],[19], which we approved again with experienced clinicians. Afterwards we generalized the indicators to use them in our generic reference process model. Finally we investigated the IT-infrastructures, which are nowadays commonly installed in hospitals.

To provide process-based and integrated performance measurement support we selected the event-based process chain (EPC)-method for modeling. EPCs have become widely-used for business process modeling in Europe, in particular in countries where SAP acts as a leading Enterprise Resource Planning (ERP) system [20]. EPCs are inspired from petri nets, incorporate role concepts and data models like ER-models or UML class diagrams [20]. According to [21], the EPC belongs to the most advanced business process modeling languages beside the UML and is one of the most used languages in the measurement researching community [22].

We evaluated and verified our approach with several clinicians and medical specialists applying it in the hospital environment. Therefore we simulated the clinical pathways of several hospitals using our EPC-model in the ARIS-Toolset. ARIS provides the possibility to test and to simulate the modelled processes already at design time, i.e. before an implementation in the clinical IT-environment is started or performed. An integrated “Process Performance Manager“ (ARIS PPM) can be used for process cycle times analysis purposes. ARIS PPM extracts information from logfiles and shows them graphically including process bottlenecks. The second evaluation step involved the mapping of the IT-systems, which were used in the hospitals in question to the modelled events. Thereby, the events, which can be automatically extracted, were identified. Finally the derived clinical indicators were evaluated again in discussion with clinicians as well as through a comprehensive literature study.

The objective of this practical validation was first to demonstrate the applicability of the approach in the hospital environment. This included the approval of the appropriateness of the modelling approach. First, the used EPC-methodology was approved as understandable for clinicians and as practicable for process monitoring in hospitals. As we modelled and use separate EPC-modules in our approach, which are loosely connected using process interfaces, we assure the flexibility that is required in such a highly flexible clinical environment. This means that several diagnostic and treatment procedures as well as performance measures can be repeated and performed in parallel if it is necessary. Finally, our modelled generic indicators were approved as *relevant, valid, reliable* and *feasible*. They are *relevant* because they relate to clinically important, evidence-based aspects of care over which clinicians have significant control. Furthermore they are *valid* as they correlate strongly with the measurement goals. Very important was the approval of the indicator *reliability* as the indicators are mapped to selected model events and therefore valid for comparison and standardized recording. Finally, the indicators are *feasible* as they are sustainable over time. In future work we will formalize our approach stronger and consider a graph valuated mechanism therefore. As we focus on time-critical diseases, a

statistical evaluation would require a comprehensive study involving several hospitals. In near future, statistical results will be provided therefore using simulation.

4 Related Work

So far, there is no process model published that describes the integration of clinical performance metrics for acute diseases taking in particular into account the possibilities for their automatic extraction using clinical information systems during a patient's stay in hospital and additionally considering the characteristics of the clinical environment, such as flexible workflows and modular process structure. Nevertheless, focusing on classical workflow analysis, event-based process monitoring and mining approaches with the objective to optimize the workflows are presented in [23], [24], [25] and [26]. However, they were not designed for health care purposes and focus on more standardized processes as they are usually performed in clinical domain. An identification approach aiming the clinical context identification and categorization is proposed in [27].

A KPI-based framework for process-based benchmarking of hospital information systems is presented in [28]. In contrast to our approach, documentation processes and outcome criteria related to the process flow and underlying structures in the area of the Hospital Information System (HIS) aiming benchmarking possibilities are there focused on. An approach using the derivation of hospital-specific clinical guidelines is described in [29]. Thereby, lifecycle support for medical guidelines and pathways is the objective.

Further clinical (reference) process models are presented in [30], [31], [32], [33], [34]. In [35] the authors present a modeling approach, which introduces principles of process modeling in healthcare using EPCs and employ these principles to existing medical information systems by implementing them using a Workflow Management System (WFMS). However, the possibility for event-based clinical performance measurement in time-critical diseases, as we introduce it in our approach, is not provided. Additionally, the model-based generic and disease-specific performance monitoring is also not supported.

5 Conclusion and Outlook

Especially in the field of time-critical diseases, performance monitoring and measurement is important regarding process optimization and quality improvement in hospitals. Our approach comprehends performance monitoring using standardized time-based performance indicators based on clinical IT-systems and standards as well as the developed inpatient clinical process model. In future work we will investigate the usage of our approach in clinical workflow engines to provide the possibility to manage and improve the IT-based workflow execution and to enable process mining.

References

1. Groene, O., Skau, J., Frolich, A.: An international review of projects on hospital performance assessment. *International Journal for Quality in Health Care* 20(3), 162–171 (2008)
2. Copnell, B., et al.: Measuring the quality of hospital care: an inventory of indicators. *Internal Medicine Journal* 39(6), 352–360 (2009)
3. n.e.: Der Einfluss der Prozessreife auf den Erfolg von Telekommunikationsunternehmen, Detecon study (2009), <http://www.detecon.com/bpm>
4. Gattnar, E., Ekinci, O.: Process Assessment and Improvement Using Capability Maturity Models. In: *Podium der Wirtschaft, Recht & Soziales*, vol. 22, pp. 121–148. VHK, Altdorf b. Nbg (2011)
5. Gattnar, E., Ekinci, O.: Qualitätsmanagement im Krankenhaus – Konzepte, Rahmenbedingungen und Standards. In: *Podium der Wirtschaft, Recht & Soziales*, vol. 21, pp. 67–130. VHK, Altdorf b. Nbg (2011)
6. Mainz, J.: Developing evidence-based clinical indicators: a state of the art methods primer. *International Journal for Quality in Health Care* 15(supp. 1), i5–i11 (2003)
7. de Vos, M., et al.: Quality measurement at intensive care units: which indicators should we use? *Journal of Critical Care* 22, 267–274 (2007)
8. Rubin, H.R., Pronovost, P., Diette, G.B.: The advantages and disadvantages of process-based measures of health care quality. *International Journal for Quality in Health Care* 13(6), 469–474 (2001)
9. Rubin, H.R., Pronovost, P., Diette, G.B.: From a process of care to a measure: the development and testing of a quality indicator. *International Journal for Quality in Health Care* 13(6), 489–496 (2001)
10. *Crossing the Quality Chasm: A new Health System for the 21st Century*. Committee on Quality of Health Care in America, Institute of Medicine (IOM), 4th edn. National Academy Press, Washington (2003)
11. McGlynn, E.A.: Identifying, Categorizing, and Evaluating Health Care Efficiency Measures. Final Report (prepared by the Southern California Evidence-based Practice Center—RAND Corporation, under Contract No. 282-00-0005-21). AHRQ Publication No. 08-0030. Agency for Healthcare Research and Quality, Rockville, MD (2008)
12. Pencheon, D.: *The good indicators guide: understanding how to use and choose indicators*. NHS Institute for Innovation and Improvement, Association of Public Health Observatories (APHO) (2007)
13. Gattnar, E., Ekinci, O., Detschew, V.: Event-based Workflow Analysis in Healthcare. In: *Proceedings of the 9th International Workshop on Modeling, Simulation, Verification and Validation of Enterprise Information Systems (MSVVEIS 2011)*, in conjunction with the 13th International Conference on Enterprise Information Systems (ICEIS 2011), pp. 61–70. SciTePress (2011)
14. Gattnar, E., Ekinci, O., Detschew, V.: Clinical Process Modeling and Performance Measurement in Hospitals. In: *Proceedings of the 5th International Workshop on Advances in Quality of Service Management (AQuSerM 2011)*, in conjunction with 15th IEEE Enterprise Computing Conference (EDOC 2011), pp. 132–140. IEEE Computer Society Press, Los Alamitos (2011)
15. Gattnar, E., Ekinci, O., Detschew, V.: Interoperable Process Monitoring using Clinical IT-Standards and Healthcare Information Systems. In: *Proceedings of the IADIS International Conferences Informatics 2011, Wireless Applications and Computing 2011, Telecommunications, Networks and Systems 2011*, pp. 59–66. IADIS Press (2011)

16. Gattnar, E., Ekinci, O., Detschew, V.: An Integrated Approach to Standardized and Automated Extraction of Clinical Process Cycle Times. In: Proceedings of the IADIS International Conferences eHealth 2011, pp. 128–136. IADIS Press (2011)
17. Gattnar, E., Ekinci, O., Detschew, V.: A Novel Generic Clinical Reference Process Model for Event-Based Process Times Measurement. In: Abramowicz, W., Maciaszek, L., Węcel, K. (eds.) BIS 2011 Workshops. LNBIP, vol. 97, pp. 65–76. Springer, Heidelberg (2011)
18. Wollersheim, H., et al.: Clinical indicators: development and applications. *The Netherlands Journal of Medicine*, Van Zuiden Communications B.V. 65(1), 15–22 (2007)
19. Scott, I., Phelps, G.: Measurement for improvement: getting one to follow the other. *International Medicine Journal*, Royal Australasian College of Physicians 39, 347–351 (2009)
20. Korherr, B., List, B.: Extending the EPC and the BPMN with Business Process Goals and Performance Measures. In: Proceedings of the 9th International Conference on Enterprise Information Systems (ICEIS 2007) (2007)
21. List, B., Korherr, B.: An Evaluation of Conceptual Business Process Modelling Languages. In: Proceedings of the 21st ACM Symposium on Applied Computing (SAC 2006). ACM Press (2006)
22. Sánchez Gonzáles, L., Rubio, F.G.: Measurement in business processes: a systematic review. *Business Process Management Journal* 16(1), 114–134 (2010)
23. zur Mühlen, M., Rosemann, M.: Workflow-based process monitoring and controlling—technical and organizational issues. In: Sprague, R. (ed.) Proceedings of the 33rd Hawaii International Conference on System Science (HICSS-33), pp. 1–10. IEEE Computer Society Press, Los Alamitos (2000)
24. van der Aalst, W.M.P.: Business Alignment: Using Process Mining as a Tool for Delta Analysis. *Requirements Engineering* 10(3), 198–211 (2005)
25. van der Aalst, W.M.P., et al.: Workflow mining: A survey of issues and approaches. *Data & Knowledge Engineering* 47(2), 237–267 (2003)
26. van der Aalst, W.M.P., ter Hofstede, A.H.M., Weske, M.: Business Process Management: A Survey. In: van der Aalst, W.M.P., ter Hofstede, A.H.M., Weske, M. (eds.) BPM 2003. LNCS, vol. 2678, pp. 1–12. Springer, Heidelberg (2003)
27. Ghattas, J., Peleg, M., Soffer, P., Denekamp, Y.: Learning the Context of a Clinical Process. In: Rinderle-Ma, S., Sadiq, S., Leymann, F. (eds.) BPM 2009. LNBIP, vol. 43, pp. 545–556. Springer, Heidelberg (2010)
28. Jahn, F., Winter, A.: A KPI Framework for Process-based Benchmarking of Hospital Information Systems. In: Proceedings of the 23th International Conference of the European Federation for Medical Informatics (MIE 2011), pp. 542–546. IOS Press (2011)
29. Lenz, R., Reichert, M.: IT support for healthcare processes – premises, challenges, perspectives. *Data & Knowledge Engineering* 61, 39–58 (2007)
30. Eisentraut, K., et al.: Abbildung klinischer Behandlungspfade als Modelle in der Unified Modeling Language. 53. Jahrestagung der Deutschen Gesellschaft für Medizinische Informatik, Biometrie und Epidemiologie (gmds) (2008)
31. Mauro, C., Happle, T., Sunyaev, A., Krcmar, H., Leitmeister, J.M.: From Medical Processes to Workflows – Modeling of Clinical Pathways with the Unified Modeling Language. In: Proceedings of 3rd International Conference on Health Informatics (HealthInf 2010), pp. 388–391. INSTICC Press (2010)
32. McGregor, C., et al.: A Structured Approach to Requirements Gathering Creation Using PaJMa Models. In: Proceedings of the International Conference of IEEE Engineering in Medicine and Biology Society, Vancouver, Canada, pp. 1506–1509 (2008)

33. Winter, A., Haux, R.: A Three Level graph-based Model for the Management of Computer-Supported Hospital Information Systems. *Methods of Information in Medicine* 34(4), 378–396 (1995)
34. Buchauer, A., et al.: 3LGM: Method and Tool to support the management of heterogeneous hospital information systems. *Computers in Medicine, Polish Society of Medical Informatics*, 72–82 (1997)
35. Becker, J., Fischer, R., Janiesch, C.: Optimizing U.S. Health Care Processes - A Case Study in Business Process Management. In: *Proceedings of the 13th Americas Conference on Information Systems (AMCIS 2007)* (2007)

A Model for a Motivational System Grounded on Value Based Abstract Argumentation Frameworks

Eugenio Di Tullio¹ and Floriana Grasso²

¹ Dipartimento di Informatica, University of Bari, Italy
ditullioeugenio@gmail.com

² Department of Computer Science, University of Liverpool, UK
floriana@liverpool.ac.uk

Abstract. Digital interventions for promoting behavioural changes have become more and more prevalent, due to the ubiquitous role played by the Web 2.0. Not many of these programmes are however grounded on well established theories, both from a psychological and from a dialogue perspective. We present in this paper a model for incorporating a general framework for abstract argumentation, into a motivational intervention based on the Transtheoretical Model of Behaviour Change. A preliminary implementation of the model as a proof of concept has been carried out on the domain of healthy eating.

Keywords: Motivation, health promotion, argumentation systems.

1 Introduction

Digital interventions to promote healthy lifestyles are more and more ubiquitous. [13] offers a detailed overview of Internet based interventions, and conclude that, as their effect can vary substantially, it is important to identify the parameters that can contribute to their success. Their analysis seems to suggest that more successful interventions are strongly grounded on theory, especially if the theory deals with planned behaviour, and that using a variety of techniques which could impact different stages of the changing process is crucial, as well as the use of different styles of interactions. In this paper we show how insights from a relatively new research field, Argument and Computation [10], can be utilised to augment the design of a motivational system, based on the Transtheoretical Model of Change. A small prototype has been implemented, in the domain of healthy eating. This paper, after introducing some relevant background literature, will discuss the design of the system, and show a walked-through example to demonstrate the way we envisage this system to work.

2 Background

The Transtheoretical model of Change [9] is a widely accepted theory that attempts to model how people modify their behaviour. The model suggests that

people progress through “stages of change”, from a first *precontemplation* stage, when people see no problem with their current behaviour, up to a *maintenance* stage, where the new behavior is continued on a regular basis. In each of the stages, an advisor can use various strategies to foster movement to the next stage. In particular, when people have to move from precontemplation to *contemplation*, a stage where it is clear that there is a problem to address, information alone is not sufficient, and it has been argued that appeal to persuasion or argumentation could be beneficial [3,11]. Therefore it makes sense to explore what research fields like persuasive technologies or argument and computation can contribute to the problem specification. In particular, we were interested in those theories which explicitly model extra-rational factors in the argumentation process, such as perspectives and values. A recent work [12] in particular attempts a systematic definition of the concept of “value” and the way this can be used in pragmatic argumentation to reason about goals, augmenting the “Value-Based Argumentation Framework” as defined in [2]. The system presented in this paper is based on this work, and uses the definitions of values and perspectives as described in the paper. This section gives a brief overview of these notions, but we refer to the paper for a more extensive description.

Perspectives

A state of affair can be evaluated from different points of view, or perspectives. This allows to express concepts like: from the “health” perspective t is preferable to s , while from a “travel comfort” perspective, the opposite is true. A **perspective** is defined by [12] as a pre-order on states in a certain domain, and is denoted with \leq_p . We use p, q, r to denote perspectives. When $s \leq_p t$ (where s and t are states), we say that t is at least as preferred as s from perspective p .

The same state of affair can be evaluated from different perspectives. In some cases it is not known which of two states is preferred from a given perspective, but one can assume that one perspective is influenced by another. For example one can assume that perspective $p = \text{“Being successful”}$ **positive influences** perspective $q = \text{“Happiness”}$ thus if a state is preferred from p , it will be preferred from q too. We denote that a perspective p positively influences the perspective q with notation $p \uparrow q$ and similarly we use $p \downarrow q$ for a negative influence.

Perspectives influencing each other may create **influence chains**. These chains can be represented with a directed graph (Fig. 1) where the dashed arrows indicate a negative influence, while the solid ones a positive influence. The left graph denotes how perspective p positively influences q , which in turn negatively influences r . It follows that p negatively influences r (right graph).




Fig. 1. Perspective Influence Chain

Values

Perspectives are subjective points of view and are not necessarily shared by every agent in the dialogue. The way in which an agent reasons with perspectives is based on the agent's set of values and preferences. An **agent preference** is a pre-order on states denoted with $<_{\alpha}$. When $s <_{\alpha} t$ we say that agent α prefers state t to state s . Therefore agent preferences are perspectives which an agent considers valid. If one or more preferences maximise a particular perspective, this perspective is considered a value for that agent. A **value** is therefore a perspective maximised from one or more agent's preferences. Starting from an agent's values, it is possible to infer other agent's preferences, kept implicit in the user model, by finding those which maximise his values.

Values can also be in an order relationship, to allow for situations in which, for example, a user may prefer the value Health to Happiness, and thus will choose a state transition that maximises Health over one that maximises Happiness. A **Value System** is a pre-order on values, and is denoted with \prec_{α} . When $W \prec_{\alpha} V$ we say that agent α prefers to promote the value set V over the value set W . Transitions may promote or demote a particular value set:

- **pro:** function $pro : S \times S \rightarrow V$ determines the values promoted by a transition from a state to another. This function is defined as: $pro(s, t) = \{v \in V_{\alpha} \mid t <_v s\}$. When $pro(s, t) = V$ we say that the transition from s to t promotes V .
- **dem:** function $dem : S \times S \rightarrow V$ determines the values demoted by a transition from a state to another. This function is defined as: $dem(s, t) = pro(t, s) = \{v \in V_{\alpha} \mid t <_v s\}$. When $dem(s, t) = V$ we say that the transition from s to t demotes V .
- **Neutral transition:** when a transition from s to t neither promotes nor demotes a value ($v \notin (dem(s, t) \cup pro(s, t))$) we say that that transition is neutral for that value.

By means of the value ordering and the functions pro and dem , it is possible to derive α preferences.

3 System Description

The system's architecture is shown in Fig. 2. The jVS dialogue manages the interaction with the user. It is build on top of an ASPIC type dialogue [1], which is an implementation of a dialogue game based on an abstract argumentation

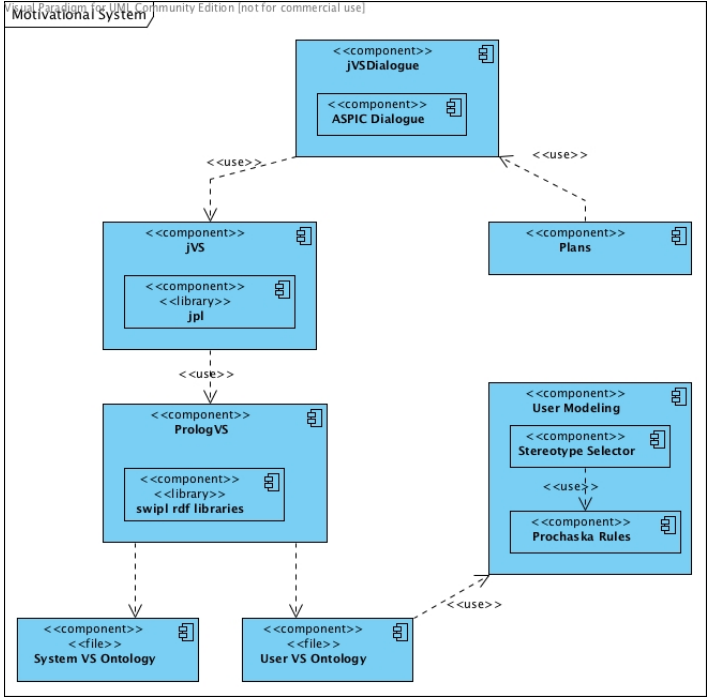


Fig. 2. System Architecture

framework. The jVS interfaces the dialogue component with the reasoning engine, a set of prolog libraries which reason about the value systems according to the theory at [12]. The libraries use information from a set of ontologies describing the stereotypical model of a user in each of the 6 stages of change, according to the Transtheoretical model, as well as the set of values and perspectives which can be applied to a state of the world. The system maintains one belief model, in form of an ontology, for each user of the system, as well as the system's own. Aside from domain related notions, the model of values and preferences in the ontology is shown in Fig. 3. The User Modelling component chooses the appropriate stereotype for the user, from information retrieved during the interactions, on the basis of a widely used Transtheoretical model questionnaire. The Plans component manages the plans used by the system to build the interaction with the user in each stage.

The interaction is driven by discourse plans describing the evolution of the motivational strategy. A plan consists of domain knowledge, eliciting questions to build the user model of their values, suggestions to the user depending on which stage the user is in, and motivational expressions to operate some positive reinforcement when the user succeeds in achieving a goal. The plan content depends on the stage of change of focus at any one time, as to each stage corresponds a different set of goals. For instance in the precontemplation stage one goal might be raising awareness of the problem the user is facing.




Fig. 3. Value/Preference Ontology

The main drive for deciding what to say is the user's stage of change. The discourse plan, on the basis of the different stages of change, identifies appropriate argumentation paths in the VS. One of the paths is chosen on the basis of the user model (taking into account what the user already said and whether the user had already 'accepted' some argumentations beforehand). The use of a VS path allows to detect inconsistencies in the user's system of values. In this paper we show the value systems using a very simple, preliminary, formalisation, not comprising the stage of change user profile, that has mainly the purpose of better explain how the system manipulates values and how they are connected together. In the remainder of the section, an example will help understand this mechanism.

3.1 A Walked through Example

Let us suppose the system contains the states: (i) Eating junk food less than 4 times a month; (ii) Eating junk food between 4 and 8 times a month; (iii) Eating junk food more than 8 times a month, while the perspectives to evaluate the states are Healthy eating; Fitness; Health; Social life. Let us also suppose that the following pre-order applies in, respectively, the system's set of beliefs (left) and the user's set of beliefs (right):

System's VS	User's VS
Healthy Eating \uparrow Fitness	-
Fitness \uparrow Social life	Fitness \uparrow Social life
Fitness \uparrow Health	-
Social life \prec Health	Social life \prec Health

At the start of the interaction, the user makes no connection between healthy eating and health, therefore the user justifies some behaviours, like eating too much junk food, without loss of coherence in his system. Also, the user values

health and social life, with the former preferred over the latter, therefore if the system succeeds in showing that some behaviour is negatively affecting health, the user might agree that it is a behaviour that should be modified. Let us suppose the dialogue is concerned with the transition:

A = Eating junk food more than 8 times a month is preferred to Eating junk food less than 4 times a month

The user has a preference for the transition, while the system considers A a transition the user should avoid, as the perspectives that A demotes are more than those which A promotes. But in the user’s model, the user is not “aware” of all A’s negative impacts. The situation is therefore:

System's VS	User's VS
Perspectives that the transition A promotes	
Social Life	Social Life
	User's preferences
Perspectives that the transition A demotes	
Healthy Eating	
Fitness	
Health	
Social Life	

The argumentation the system may attempt with the user is therefore:

Eating junk food more than 8 times a week, instead of less than 4 times a week, is not advisable because your Health would be penalised, because A penalises healthy eating, which in turn promotes fitness, which in turn promotes health.

The system hopes that by integrating these new connections, the user would be persuaded that the transition is negative. The VS focuses on the health perspective, because the user has indicated this is preferred over other perspectives. If the user accepts the line of reasoning above, the image of the user’s VS will be modified as in Fig. 4. A screenshot of the system corresponding to one phase in this dialogue is shown in Fig. 5.




Fig. 4. User and System’s VS after the interaction in the example

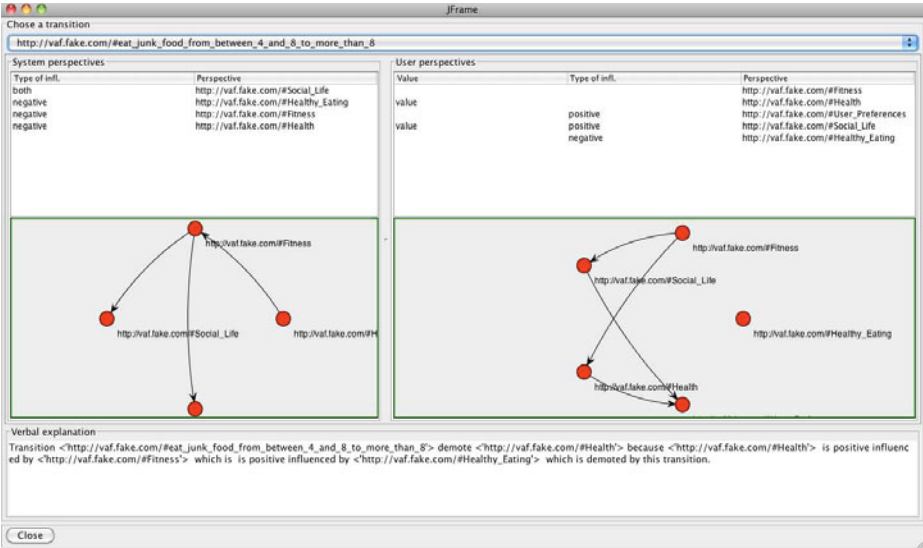


Fig. 5. Screenshot of the system for the example

4 Conclusion

This paper provides a practical, although prototypical, demonstration of a motivational system that uses the Transtheoretical Model [9] and the Value System [12] to adapt the interaction on the basis of the user's state of change and on his personal values.

Other work has been done on digital intervention for motivation, using similar theories and ideas. [5] presents a motivational system to stop smoking that interacts with the user via phone text. The system texts users periodically with messages that encourage to persevere, and that focus on their success so far. Users can also text the system in case they crave for cigarettes, receiving messages that help them to carry on with the quitting attempt. The paper proves that this way to interact with the user considerably increases the rate of success of the smoke cessation intervention. In [7] the Transtheoretical Model and various motivational techniques are used for motivating people to save energy. They successfully apply well known theories to new contexts, paying much attention to the system's capability of understanding which changing stage the user is in. The system is however still at a theoretical stage. [6] presents a software aimed at behavioural researchers, where they can build platforms to test their interventions on real users and process the results. Researchers are provided with an end-user development environment where they can build their own interfaces and facilities that will be used by their test subjects, giving them tailored advice. Similar tools are intended to be used for the system in this paper.

The system presented in this paper has been implemented as a prototype. The argumentation component is fully implemented, while the strategic component

is still at a design stage. The role of natural dialogue for interacting with the system has been left for the moment outside the scope of this work. Testing of the system showed that it is able to provide very simple argumentation, and it is envisaged that with more complex ontologies and with an interface with a NL dialogue system it will be possible to obtain more complex and convincing natural dialogue with the user. A small evaluation is being designed as a Wizard of Oz experiment [8], and plans are under way to adapt the system so that it can be interfaced with an embodied conversational agent developed at the first author's research lab [4].

References

1. Amgoud, L., et al.: Project N 002307 ASPIC, Deliverable D2.6 - Final review and report on formal argumentation system, available as Technical Report ULCS-07-005, Department of Computer Science, University of Liverpool, <http://www.csc.liv.ac.uk/research/techreports/tr2007/tr07005abs.html>
2. Bench-Capon, T.J.M.: Persuasion in Practical Argument Using Value Based Argumentation Frameworks. *Journal of Logic and Computation* 13(3), 429–448 (2003)
3. Cawsey, A., Grasso, F.: Goals and Attitude Change in Generation: a Case Study in Health Education. In: Jokinen, K., Maybury, M., Zock, M., Zukerman, I. (eds.) *Proceedings of the ECAI 1996 Workshop on: Gaps and Bridges: New directions in Planning and NLG*, pp. 19–23 (1996)
4. de Rosi, F., Novielli, N., Carofiglio, V., Cavalluzzi, A., De Carolis, B.: User Modeling And Adaptation In Health Promotion Dialogs With An Animated Character. *International Journal of Biomedical Informatics*, 514–531 (2006)
5. Free, C., Knight, R., Robertson, S., et al.: Smoking cessation support delivered via mobile phone text messaging (txt2stop): a single-blind, randomised trial. *Lancet*
6. Hare, J., Osmond, A., Yang, Y., Wills, G., Weal, M., De Roure, D., Joseph, J., Yardley, L.: LifeGuide: A platform for performing web-based behavioural interventions. In: *WebSci 2009: Society On-Line*, Athens, Greece, March 18–20 (2009)
7. He, H.A., Greenberg, S., Huang, E.M.: One size does not fit all: applying the transtheoretical model to energy feedback technology design. In: *Proceedings of the 28th International Conference on Human Factors in Computing Systems (CHI 2010)*, pp. 927–936. ACM, New York (2010)
8. Kelley, J.F.: An iterative design methodology for user-friendly natural language office information applications. *ACM Transactions on Office Information Systems* 2(1), 26–41 (1984)
9. Prochaska, J., Clemente, C.D.: Stages of Change in the Modification of Problem Behavior. In: Hersen, M., Eisler, R., Miller, P. (eds.) *Progress in Behavior Modification*, vol. 28. Sycamore Publishing Company, Sycamore (1992)
10. Rahwan, I., Simari, G.R.: *Argumentation in Artificial Intelligence*, 1st edn. Springer Publishing Company, Incorporated (2009)
11. Reiter, E., Robertson, R., Osman, L.: Types of Knowledge Required to Personalize Smoking Cessation Letters. In: [25], pp. 389–399
12. van der Weide, T., Dignum, F., Meyer, J., Prakken, H., Vreeswijk, G.: Practical Reasoning Using Values. In: McBurney, P., Rahwan, I., Parsons, S., Maudet, N. (eds.) *ArgMAS 2009. LNCS (LNAI)*, vol. 6057, pp. 79–93. Springer, Heidelberg (2010)
13. Webb, T.L., Joseph, J., Yardley, L., Michie, S.: Using the Internet to promote health behavior change: A systematic review and meta-analysis of the impact of theoretical basis, use of behavior change techniques, and mode of delivery on efficacy. *Journal of Medical Internet Research* 12(1) (2010)

An Assessment of the Potential for Personalization in Patient Decision Aids

Øystein Eiring and Laura Slaughter

Norwegian Electronic Health Library, Postboks 7004, 0130 Oslo. Oslo University Hospital
oye@helsebiblioteket.no, lauras@idi.ntnu.no

Abstract. An increasing number of patient decision aids are being developed to assist patients in making personalized choices among health care options, but little is known about the current use of and potential for personalization of web-based decision aids. The purpose of this study is to estimate the potential for personalization of patient decision aids. We developed a coding scheme for personalization and analyzed web-based patient decision aids from all relevant developers according to the scheme. The most relevant subgroups of users and the current representation of the groups in the tools were identified. We then identified system behavior relevant to web personalization, and instances of adaptive system behaviors. The decision aids included in this study exhibit four out of five classes of system behavior eligible for personalization. With few exceptions, the tools do not contain automatic, adaptive behavior. Patient decision aids hold potential for web personalization. Relevant techniques are largely unexplored.

1 Background

Patient decision aids are evidence-based tools designed to help people participate in making specific, deliberated, personalized choices among health care options, in ways they prefer. According to a systematic review on the effectiveness on patient decision aids, the tools «differ from usual health education materials because of their detailed, specific, and personalized focus on options and outcomes for the purpose of preparing people for decision making»[1].

The term «personalized» is often included in the definitions and declared purposes of patient decision aids. Personalization can pertain to several aspects of a decision aid, such as the structure and content of the tool, the decision-making process, and the resulting choice. Within the domain of information and communication technology, personalization implies the use of technology to accommodate the differences between individuals. Technological personalization of patient decision aids could potentially tailor healthcare and healthcare communication by adjusting to the different biological, psychosocial and contextual idiosyncrasies of patients.

Web personalization is the employment of user features in web systems that adapt their behavior to the user. The overall aim is to meet the needs, goals and preferences of a variety of people. The adapted content can be variations regarding information, products,

people, services and activities. An exploration of the techniques in the domain of decision aids could hold potential for more informed, satisfactory, effective and personalized decisions. A large inventory of techniques for adaptively selecting, structuring and presenting content in web systems to user features (attributes and interaction data) have evolved during the last two decades [2].

Despite the fact that decision aids are intended to support the personalization of care, little is known about the current use of, and potential for web personalization inherent in the tools.

2 Design and Methods

The objective of this mixed-methods study is to estimate the potential for web personalization in web-based patient decision aids.

Based on a comprehensive anthology [2] we developed a coding scheme for current web personalization techniques. The scheme includes a checklist of possible user features and a catalog of adaptive system behaviors.

We identified developers of web-based decision aids by hand-searching the quality-assessed patient decision aids in the Ottawa Inventory[3]. Developers of decision aids only available in PDF format were excluded. One decision aid from each developer was included for further study. When a developer had produced more than one decision aid, we selected the decision aid that included the richest functionality present in the developer's portfolio of tools.

Applying the coding scheme, we identified explicit and implicit subgroups with comparable user features in the decision aids. Subgroups existent in the tools that were not present in the generic coding scheme, were added to the scheme. We identified and described the linguistic representation of subgroups in the decision aids.

We then mapped the system behavior in the decision aids to the classes and prerequisites of adaptable system behavior present in the coding scheme. Finally, we systematically identified user-adaptive behavior present in the decision aids.

3 Results

259 decision aids developed by 22 producers were found in the Ottawa Inventory. The 10 producers that met the inclusion criteria were responsible for 223 of the decision aids included in the inventory.

3.1 A Coding Scheme for Personalization of Patient Decision Aids

According to the coding scheme, the basic components of personalization are the media content, user features, user model construction and representation, and adaptive system behavior. We analyzed the decision aids according to the four classes. User features can broadly be classified into the user's knowledge level, interests, preferences, goals/tasks, background, individual traits and context. Adaptive system behaviors include adaptive

navigation support, adaptive selection, organization and presentation of content, adaptive search, adaptive collaboration and personalized recommendations. We adjusted and updated the coding scheme to the features identified in the decision aids. In example, we included somatic parameters as a subclass of individual traits.




Fig. 1. Screenshot of the coding scheme. The subclass *individual traits* is expanded.

3.2 Content Types

8 out of 10 decision aids included at least two media types. 8 decision aids could be classified as hypermedia, by definition present when hyperlinks and at least 2 media types are integrated in a web system.

3.3 Explicit and Implicit Subgroups

The majority of the implicit and explicit subgroups targeted and represented in the decision aids could be classified in accordance with the user features subclasses in the coding scheme. The most frequently found subgroups were coping styles, emotional reactions, cognitive skills, user's beliefs, experiences of users, literacy level and somatic parameters.

Most of the content that could be differentiated according to subgroups pertained to the somatic parameters of individual patients. The most frequently used somatic parameters that defined subgroups of users were risk factors, factors relevant for the eligibility for treatment, incidences, prevalences, and probabilities, outcomes of decisions, etiology, lab test results and prediction of recovery.

One of the decision aids included an *evidence summary* that summarized the research underlying the decision aid. 53 % of the sentences in the evidence summary

contained information that described different subgroups. This information was, with few exceptions, not found in the decision aid.

3.4 Representation of Subgroups

The decision aids embedded a number of different strategies to address user groups with different features. The most frequently applied representations of subgroups identified in the decision aids are the following:

- Listing several subgroups and making specific statements true for each subgroup one by one
- Making statements that are untrue or irrelevant to at least one subgroup
- Alluding to subgroups without specifying the attributes of the subgroups
- Giving an average for all subgroups combined
- Suggesting that a patient belongs to one, particular subgroup
- Listing only some subgroups
- Not acknowledging the existence of relevant subgroups
- Asking user to determine the relevant subgroup her-/himself
- Helping the patient determine the relevant subgroup e.g. through an interactive tool
- Describing how health personnel should determine the relevant subgroup
- Giving general information but acknowledging that subgroups do exist
- All decision aids in the study included information that was true and/or relevant only to subsets of users with particular user features.

3.5 User Model Construction and Representation

Direct or indirect evidence of the application of user models was not found.

3.6 System Behavior and Adaptation

All decision aids presented selected and organized content. 6 of the decision aids included a search field, but only 1 included a search field that searched specifically for content included in the tool only. The search fields in the remaining 5 decision aids searched the content in all of the website. 1 developer included a tailored search field to selected websites. With 2 exceptions, the decision aids exhibited static navigation. None of the producers enabled user collaboration. 1 included a user forum directly related to the tool. 1 decision aid included automatic recommendation of content based on implicit interaction data. Users could manually enter personal data in 9 of the 10 decision aids, the most frequent feature being the possibility to adjust text size.

Table 1. Decision aids included in the study

Clinical condition	Developer	Hyper media	Adaptive behavior	Corpus
Treatment options for early breast cancer	Cardiff University and others, UK	Yes	Basic	Closed
Stroke prevention in atrial fibrillation	Healthwise Inc, US	Yes	Basic	Closed
Statins for prevention of cardiovascular disease	Mayo Clinic, US	No	Basic	Closed
H1N1 vaccination	CHEO, Canada	Yes	Basic	Closed
Treatment options for early breast cancer	Baylor College of Medicine, US	Yes	Medium	Closed
PSA testing	Health Dialog Inc, US	Yes	Basic	Closed
Long term feeding tube placement in elderly patients	Ottawa Patient Decision Aid Research Group, Canada	Yes	Basic	Closed
Mammography screening	Public Health Agency of Canada	No	Basic	Closed
Premixed insulin for type 2 diabetes	AHRQ, US	Yes	Basic	Closed
Treatment options for early prostate cancer	National Cancer Institute, US	Yes	Basic	Closed

4 Discussion

Whereas paper-based and linear media can be personalized to a very limited degree, web-based hypermedia systems can adapt their content and presentation to individual users employing an array of techniques. The web-based decision aids included in this study display 4 out of 5 classes of system behavior eligible for personalization. None of the systems contain advanced adaptive behavior. All the selected decision aids are closed corpus systems.

The web-based decision aids analyzed in this study target and address several subgroups within what is mainly a one-size-fits-all format. A number of strategies are applied to relate the content to different subgroups, of which some are potentially misleading, inappropriate, and might entail increased cognitive burden and unnecessary uncertainty. A relatively large amount of content could have been differentiated according to different patient's somatic parameters.

In the case where the evidence supporting the decision aid was provided, our analysis indicates that most of the information could be differentiated to individual patients. This information was only to a limited degree reflected in the decision aid.

As decision aids are developed to support patient's personalized choices, the general absence of personalization in current decision aids is a paradox. According to our findings, a significant amount of content in both the decision aids and their underlying

evidence base could be tailored to the somatic parameters or other features of individual subgroups.

In addition to personalized *selection* and *organization* of content, system behavior could adapt to users in a variety of ways. To give an idea of the potential, personalization techniques could include:

- *Presentation* of the content most relevant to the user by priority-on-context techniques such as coloring or scaling
- Supported *navigation* by augmenting links with annotations that give visual cues, for instance of progress
- *Recommendation* of peers with similar features that is treated at the same hospital, to the patient
- Personalized web *search* adapted to the health literacy of the user
- Personalized *collaboration* where users e.g. build a shared list of questions to ask their caregiver or physician

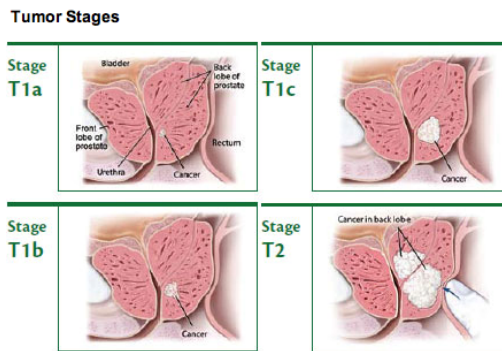


Fig. 2. In the patient decision aid *Treatment choices for men with early prostate cancer*, information about different tumor stages is presented identically to all users. The presentation could be personalized utilizing techniques that highlight the tumor stage relevant to the individual patient. Content fragments could be dimmed, colored, scaled or sorted. Reprinted with permission from the National Cancer Institute.

5 Conclusion

Existing web personalization techniques constitute a rich resource of developmental possibilities for various aspects of decision aids that is largely unexplored. By adapting system behavior already present in current web-based decision aids, the content and functionality could potentially be tailored to a range of user groups. The time and effort in using the techniques must be weighted against the possible benefits.

Advanced personalization techniques require the representation of user features in user models, which is unlikely to be worth the cost when decision aids are stand-alone applications. User models could potentially be built utilizing input from the electronic patient journal. Embedding decision aids in patient portals holds promise of enabling

adaptive behavior of decision aids. The advance of semantic web technologies such as RDF and OWL, and the use of an open corpus knowledgebase, could bring added possibilities.

References

1. O'Connor, A.M., Bennett, C.L., Stacey, D., Barry, M., Col, N.F., Eden, K.B., Entwistle, V.A., Fiset, V., Holmes-Rovner, M., Khangura, S., Llewellyn-Thomas, H., Rovner, D.: Decision aids for people facing health treatment or screening decision. *Cochrane Database Syst. Rev.* (3), CD001431 (2009)
2. Brusilovsky, P.: Adaptive Navigation Support. In: Brusilovsky, P., Kobsa, A., Nejd, W. (eds.) *Adaptive Web 2007*. LNCS, vol. 4321, pp. 263–290. Springer, Heidelberg (2007)
3. <http://decisionaid.ohri.ca/AZinvent.php> (accessed July 20, 2011)

An Analysis of Twitter Messages in the 2011 Tohoku Earthquake

Son Doan, Bao-Khanh Ho Vo, and Nigel Collier

National Institute of Informatics
Chiyoda-ku, Hitotsubashi, Tokyo, Japan
{doan,khanhvo,collier}@nii.ac.jp

Abstract. Social media such as Facebook and Twitter have proven to be a useful resource to understand public opinion towards real world events. In this paper, we investigate over 1.5 million Twitter messages (*tweets*) for the period 9th March 2011 to 31st May 2011 in order to track awareness and anxiety levels in the Tokyo metropolitan district to the 2011 Tohoku Earthquake and subsequent tsunami and nuclear emergencies. These three events were tracked using both English and Japanese tweets. Preliminary results indicated: 1) close correspondence between Twitter data and earthquake events, 2) strong correlation between English and Japanese tweets on the same events, 3) tweets in the native language play an important roles in early warning, 4) tweets showed how quickly Japanese people's anxiety returned to normal levels after the earthquake event. Several distinctions between English and Japanese tweets on earthquake events are also discussed. The results suggest that Twitter data can be used as a useful resource for tracking the public mood of populations affected by natural disasters as well as an early warning system.

Keywords: Twitter, social media, earthquake, surveillance, natural language processing.

1 Introduction

Social media such as Facebook and Twitter have proven to be useful resources for understanding public opinion towards natural disaster events. Such resources can be used to detect general events in politics, e.g., elections [10], and finance, e.g., stock market changes [2,10] and oil price changes [10], as well as in alerting disasters such as earthquakes and typhoons [12]. Other social data such as search queries have been successfully used in public health to build bio-surveillance systems for early warning of influenza-like illness [7,8,11], showing high correlations with Centers for Disease Control and Prevention (CDC) reports. Within the wider Web, the BioCaster project has worked on detecting and tracking infectious diseases using newswire reports [5]. Twitter, the largest micro-blogging service with about 200 million users as of March 2011 [1], can generate 200 million *tweets* a day. Tweets are short but condensed personal messages with a 140 character limit designed for rapid reporting from mobile devices. Several applications using Twitter messages in biosurveillance

systems have been developed. For example, Flu Detector used Twitter messages to detect the ILI rate in the United Kingdom [9], and DIZIE, which is part of the BioCaster project, is an experimental syndromic surveillance system [6].

The great Tohoku earthquake happened on 11th March 2011 was the most powerful known earthquake to have hit Japan, and one of the five most powerful earthquakes in the world overall since modern record-keeping began in 1900 [3]. The earthquake triggered a tsunami, causing massive loss of life and destruction of infrastructure, and in turn lead to a number of nuclear accidents in Fukushima prefecture, affecting hundreds of thousand of residents. This was described as “the toughest and the most difficult crisis for Japan” by the Japanese Prime Minister [4]. Greater understanding of social responses during such disaster periods should help metropolitan governments and public health agencies to gain greater insights for preparedness and response. Twitter data, being real time and large-scale, offers a unique insight into public opinions as the disaster develops.

In this paper, we analyzed over 1.5 million Twitter messages for the period starting 9th March 2011 until 31st May 2011 – the time when the main crisis happened - in order to review social attitudes during the time when the earthquake occurred. We focused on tracking keywords related to three main topics: earthquake/tsunami, radiation and public anxiety for the Twitter user population in the metropolitan Tokyo area; an area that experienced severe tremors, social anxiety and mild radiation but no major loss of life. To gain greater insights into differing attitudes between local and foreign residents we explicitly differentiated English and Japanese tweets. Our results show high correlations between Twitter data and real world events as well as how quickly Japanese people’s anxiety returned to a stable level after the disasters. To the best of our knowledge, this is the first such study on Twitter data during the 2011 Tohoku earthquake.

2 Methods

2.1 Twitter Corpus

We collected Twitter data for three months, starting from March 9th 2011 to May 31st 2011 using Twitter API (<http://dev.twitter.com/>) with the geolocation feature set to track messages originating within Tokyo. The resulting corpus had a total of 48,870 tweets in English and 1,611,753 tweets in Japanese. The details of tweets by dates in both English and Japanese are depicted in Figure 1. In order to analyze the data, we wrote a simple parser in Python to parse the text and used relevant keywords for filtering.

2.2 Earthquake Events and Relevant Keywords

Our empirical analysis focuses on the events during the 2011 Tohoku earthquake. Within the stream of Twitter messages we studied three indicators of public response: 1) earthquake and tsunami, 2) radiation caused by the Fukushima Daiichi plant’s

meltdown, and 3) public anxiety. The first two types of indicators are aimed at showing people’s awareness of the earthquake, tsunami and radiation and the last indicators looks at how people in Tokyo are anxious about these events. Essentially, the events happened as the sequences as follows: The first is the earthquake occurring at 05:46:23 UTC on Friday, 11th March 2011. The second is the tsunami which happened after the earthquake a few minutes. The third is the nuclear explosion at the Fukushima Daiichi plant which the first explosion at reactor 1 happened at 6:36 UTC on 12th March 2011.




Fig. 1. Tweet numbers by dates in English (left) and in Japanese (right)

We manually investigated both English and Japanese tweets and constructed lists of key terms that are relevant to the events. The lists of English and Japanese terms indicating earthquake and tsunami, radiation and public anxiety are shown in Table 1. Although the number of terms might not be comprehensive we believed that they are good enough for our investigation. The reason we included earthquake and tsunami as a single event was because we found there were few English tweets about tsunami in our corpus.

2.3 Data Analysis

Keyword filtering has shown to be a simple, but effective way to filter tweets for relevant topics [7,8,9,10,11]. In our study, we filtered tweets by event keywords in Table 1. We normalized tweets per day by dividing the total number of filtered tweets to the total number of tweets per day, i.e.,

$$f(\text{event}) = \frac{\# \text{filtered tweets per day}}{\# \text{tweets per day}},$$

where $f(\text{event})$ is the relative frequency of events per day.

Table 1. List of relevant keywords for the Earthquake and Tsunami, Radiation, and Anxiety events

English terms	Japanese terms
<i>Earthquake and Tsunami event</i>	
earthquake, quake, quaking, post-quake, shake, shaking, shock, aftershock, temblor, tremor, movement, sway, landslide seismic, seismography, seismometer, seismology, epicenter, tsunami, wave	大地震 (major earthquake), 大震災 (great earthquake), 震災 (earthquake disaster), 地震 (earthquake), 余震 (aftershock), 揺れ (quake/tremor), 震度 (seismic intensity), 震源 (epicenter), マグニチュード (magnitude) 津波 (tsunami)
<i>Radiation event</i>	
radiation, radioactivity, radioactive, nuclear, power plant, reactor, iodine, TEPCO, meltdown, sievert, micro sievert, iodine, isodine, explosion, caesium, strontium, plutonium, uranium	放射 (radiation), 放射線 (radiation ray), 放射能 (radioactivity), 放射性物質 (radioactive material), 原発 (nuclear power plant), 東京電力 (TEPCO), メルトダウン (meltdown), マイクロシーベルト (micro sievert), ヨウ素 (iodine), イソジン (isodine), ヨウ化カリウム (potassium iodide), 炉心溶融 (core meltdown), 爆発 (explosion)
<i>Anxiety event</i>	
die, death, risky, scary, scared, incredible, freaked out, chaos, evacuate, help, unable to contact, bad, worrying, worried, anxious, annoying	死亡 (death), 死ぬ (die), やばい, やばかった, ヤバイ, やばっ, やべ (risky; dangerous), 怖い, 怖かった, 怖っ, 恐れ (scary, scared), すごい, すげえ, すげー, すっげー (incredible), びびる, びびった (freaked out), 混乱 (chaos), 避難 (evacuation), 助けて (help), 連絡とれない (unable to contact), 大変 (bad; oh, my God), 心配 (worrying), 船酔い (seasick)

3 Results and Discussions

Figure 1 supports Sakaki *et al.*'s [12] observations that the number of earthquake tweets increases significantly directly after a major earthquake. The data indicates that Twitter users would like to broadcast their experience immediately.

Event 1: Earthquake & tsunami event

Earthquake and tsunami keyword frequencies for both English and Japanese are shown in Figure 2. We noticed that there is a sharp and sudden rise in the number of tweets immediately preceding the first major tremor. Note that the earthquake happened at 14:46:23 JST (05:46:23 UTC) on Friday, 11th March 2011 with 9.0 magnitude earthquake near the east coast of Honshu, Japan which was 373 km NE of Tokyo¹.

First, we considered how quickly Twitter users responded to the earthquake. It is unknown when the first public report about the earthquake was in Tokyo but the first tweet on the topic originating in Tokyo occurred at 05:48:08 UTC, 1 minute and 25 second right after the earthquake happened at the epicenter. It is unsurprising but noteworthy that the first tweet was in Japanese.

¹ <http://earthquake.usgs.gov/earthquakes/recenteqsww/Quakes/usc0001xgp.php>

Within our corpus the first English tweets on the earthquake are given below, with the first two tweets sent from an iPhone:

```
11-03-2011T05:48:54 Huge earthquake in TK we are affected!
11-03-2011T05:49:01 BIG EARTHQUAKE!!!
11-03-2011T05:50:00 Massive quake in Tokyo
```

The first Japanese tweets on the earthquake are as follows.

```
11-03-2011T05:48:08 "地震!" [Earthquake!]
11-03-2011T05:48:08 "地震だー縦揺れ!" [Earthquake ~ vertical shake!]
11-03-2011T05:48:14 "地震!!!!!" [Earthquake!!!!!]
```

We can easily see that first Twitter users responded very quickly, with the first English and Japanese tweets occurring about two minutes right after the earthquake happened. Japanese tweets preceded the English tweets by about 47 seconds. This might be because the numbers of Japanese language users are far greater than English language users in Tokyo. We also note that when the earthquake occurred, because of network outage there was no contact by cell phones but people could still access the Internet through 3G services with smartphones such as the iPhone.

We note that the first tweet from a Tokyo resident about a tsunami in Tohoku was a re-tweet at 06:02:35 UTC, 12 minutes after the first tsunami was reported. The first tweet about a tsunami was an eye witness tweet at 2011-03-11T 05:52:23 UTC, 6 minutes after the earthquake occurred at its epicentre.

```
11-03-2011T 05:52:23 "オレ、津波の様子見てくるわ!!!!" [I can see the tsunami coming!!!!]
```




Fig. 2. Keyword frequencies for the earthquake event over time for English (left) and Japanese tweets (right)

Let us consider the details of the aftershocks which were described in Wikipedia².

“Japan experienced over 900 aftershocks since the earthquake, with about 60 registering over magnitude 6.0 Mw and at least three over 7.0 Mw. A magnitude 7.7 Mw and a 7.9 Mw quake occurred on March 11 and *the third one struck offshore on 7 April with a disputed magnitude*”

² http://en.wikipedia.org/wiki/2011_Tōhoku_earthquake_and_tsunami

“Four days later *on April 11, another strong magnitude 6.6 Mw aftershock struck Fukushima*, causing additional damage and killing a total of three people”

As reported above, there are two other significant earthquake events: **7th April** and **11th April**. Both English and Japanese tweets in Figure 1 show how significant they are since peaks occur on both of those date with frequencies of 0.12 (English) and 0.07 (Japanese) at 7th April and 0.18 (English) and 0.09 (Japanese) at 11th April, respectively. It seems from our observations that Japanese language speakers were more concerned on 11th April than 7th April.

The below tweets were the first to show concern about nuclear plants right after the earthquake.

11-03-2011T05:57:53 "原発大丈夫かな?" [Is the nuclear power plant okay?]

11-03-2011-03-11T08:43:36 "詳しくRT @u_tips: 福島第2原発が穏やかじゃないですね。" [In detail RT @u_tips: Fukushima Dai-ni power plant is on alert.]

11-03-2011T09:50:49 "福島原発ヤバい状況らしい。。。政府が国民を欺かないことを願います" [The Fukushima plant is in a really bad situation... I hope that the government won't deceive the public.]

From drill down analysis we noticed that many people reported the situation happening in Tokyo from their own personal experiences such as a lack of food in convenience store on 11th and 12th March.

11-03-2011T11:27:03 People r suggested to prepare an "emergency kit" consist of blanket, water, canned food, flashlight, aid kit, clothes #bigearthquakeinjapan

12-03-2011T01:00:18 Wow. I've never seen a convenience store depleted of food before, even during the Great Handgun quake. At least I got toilet paper.

It is easy to see that such information could be automatically harvested for timely planning in future disasters. From Figure 2, we can see that both English and Japanese tweets correlate closely, reflecting the fact that public concern in both English and Japanese are the same during the earthquake events.

Event 2: Radiation event

As reported in many public newswires, radiation was one of the main concerns after the earthquake and tsunami severely damaged the Fukushima Daiichi nuclear plant, causing three of its reactors to experience a meltdown although this fact was not confirmed until several months later. The radiation keyword frequencies for both English and Japanese tweets are depicted in Figure 3.

Following the timeline of the Fukushima Daiichi nuclear disaster³, there were significant events about explosions of reactors as follows:

12th March, 15:36JST: Massive hydrogen explosion on the outer structure of the unit 1.

14th March, 11:01JST: The unit 3 reactor building explodes, injuring six workers.

15th March, 20:00JST: A major part of the fuel in reactor 2 drops to the bottom of the reactor pressure vessel. Radiation levels at the plant rise significantly but subsequently fall back.

³ http://en.wikipedia.org/wiki/Timeline_of_the_Fukushima_I_nuclear_accidents

Radiation equivalent dose rates of 400 millisieverts per hour (400 mSv/h) are observed at one location in the vicinity of unit 3.

Figure 3 shows that after the March 11th earthquake, Japanese tweets showed further peaks on the 12th and 15th March whilst English tweets reached peaks one or two days later on the 13th and 17th March, respectively. Although the cause is not clear this indicates that Japanese people in Tokyo were concerned about radiation earlier than foreign residents in Tokyo. Once again the results indicate the important role that aggregated tweets in the native languages play in early warnings.

When the earthquake hit the Fukushima nuclear plant on 11th April, both English and Japanese tweets reached their peak a day later, on 12th April. This indicates that the event is of major concern to both Japanese people and foreigner residents in Tokyo.




Fig. 3. Keyword frequencies for the radiation event over time for English (left) and Japanese tweets (right)

Below are some examples of early English tweets at the peak on 11th and 12th April.

11-04-2011T23:21:09 Earthquakes, tsunamis, radiation what else you got? Exploding Fuji? Let me get out my folding umbrella... #fb #quaketrashtalk #backchannel

11-04-2011T23:51:09 This is the "New Normal" life after 3/11. Always fearing how the situation of nuclear plant is #prayforjapan

12-04-2011T00:08:01 Nuclear Agency Japan has increased the level of nuclear disaster to level 7 as the worst, which equal to Chernobyl #Fukushimasradiation

12-04-2011T02:41:40 @<username> 4 of 6 reactors in meltdown. Spent fuel rods melting. Two explosions, containment breached. Already a 7. #fukushima #niisa

Typically, several tweets focus on concerns about radiation in relation to tap water in Tokyo since 13th March.

13-03-2011T07:06:37 People are asked to close window, door; not to use AC; use mask & not to drink tap water #Fukushimasradiation

23-03-2011T06:18:50 210 becquerel iodine (normal 100 becquerel) discovered in Tokyo tap water. Infants are urged to avoid drinking it #Fukushimasradiation

Event 3: Anxiety event

Using keywords with Twitter data can track the public mood such as in elections or oil price changes [10]. It is interesting to see if they can be used to follow the anxiety of Tokyo residents during the earthquake. Figure 4 shows anxiety frequencies for English and Japanese tweets. Since the frequency of English tweets are relatively low since we had few tweets about anxiety, we included it here as a reference.

Figure 4 shows that the highest frequencies of the anxiety event were at the starting point -11th March and then decreased over a two week period and kept stable after that. There was a slight rise in concern on 11th April but the upward change was not so high when comparing to the earthquake and nuclear events in Japanese. This perhaps underlines the well reported response about how Japanese people kept calm during the disasters. Despite the relatively small number of English language tweets, aggregated message counts from foreign residents show many peaks. However, when comparing the peaks we can see similar trends with Japanese tweets, especially from 20th April to 11th May.




Fig. 4. Frequencies of the anxiety event by dates of Japanese tweets

Below are some examples for anxiety of people in Tokyo in English tweets,

```
11-03-2011T06:58:49      nerves frayed out on the streets. everyone
emptied out of buildings, on their phones, worried. some inspecting
damage to buildings
11-03-2011T07:09:56      @<username> I'm okay, thx. Worried about
others...
11-03-2011T07:14:52      @<username> thanks ! Now Me and my family are
ok. But worrying about my men...
```

4 Conclusions

In this paper, we provided an empirical investigation that showed high correlations between aggregated tweets and disasters during earthquake events, i.e., earthquake and tsunami, radiation. It appears that there is strong potential for tracking public anxiety in resident populations affected by the disaster. The results reveal that tweets in native

language play an important role in early warning in terms of their volume and timeliness. Strong correlation between Twitter and public health events leads us to believe that Twitter data can be a useful resource in an early warning surveillance systems as well as a tool for analyzing public anxiety and needs during times of disaster.

In the future, we plan to extend our work on analysis to other aspects of the earthquake using publicly available metrics for evaluation. Automated methods to find relevant terms for tracking during disasters will also be investigated.

Acknowledgements. Authors would like to thank to Twitter Inc. for providing API functions to access the Twitter data.

References

1. BBC (2001), <http://www.bbc.co.uk/news/business-12889048> (retrieved March 28, 2011)
2. Bollen, J., Mao, H., Zeng, X.: Twitter mood predicts the stock market. *Journal of Computational Science* 2(1), 1–8 (2011)
3. CBS News. New USGS number puts Japan quake at 4th largest, March 14. Associated Press (2011), Archived from the original on (April 5, 2011), <http://www.webcitation.org/5xgjFTgf4> (retrieved March 15, 2011)
4. CNN. Japanese PM: 'Toughest' crisis since World War II (March 13, 2011), Archived from the original on (April 12, 2011), <http://edition.cnn.com/2011/WORLD/asiapcf/03/13/japan.quake/index.html?iref=NS1> (retrieved March 13, 2011)
5. Collier, N., Doan, S., Kawazoe, A., Goodwin, R., Conway, M., Tateno, Y., Ngo, Q.-H., Dien, D., Kawtrakul, A., Takeuchi, K., Shigematsu, M., Taniguchi, K.: BioCaster: detecting public health rumors with a Web-based text mining system. *Bioinformatics* (2008), doi:10.1093/bioinformatics/btn534
6. Collier, N., Doan, S.: Syndromic Classification of Twitter Messages. In: Kostkova, P., Szomszor, M., Fowler, D. (eds.) *eHealth 2011*. LNICST, vol. 91, pp. 186–195. Springer, Heidelberg (2012)
7. Eysenbach, G.: Infodemiology: Tracking Flu-Related Searches on the Web for Syndromic Surveillance. In: *AMIA Annu. Symp. Proc. 2006*, pp. 244–248 (2006)
8. Ginsberg, J., Mohebbi, M., Patel, R., Brammer, L., Smolinski, M., Brilliant, L.: Detecting influenza epidemics using search engine query data. *Nature* 457(7232), 2012–2014 (2009)
9. Lampos, V., Cristianini, N.: Tracking the flu pandemic by monitoring the social web. In: *Proc. of the 2nd International Workshop on Cognitive Information Processing (CIP)*, pp. 411–416 (2010)
10. O'Connor, B., Balasubramanian, R., Routledge, B.R., Smith, N.A.: From tweets to polls: Linking text sentiment to public opinion time series. In: Cohen, W.W., Gosling, S. (eds.) *Proceedings of the 4th International AAAI Conference on Weblogs and Social Media*, pp. 122–129. AAAI Press (2010)
11. Polgreen, P.M., Chen, Y., Pennock, D.M., Nelson, Weinstein, R.: Using internet searches for influenza surveillance. *Clinical Infectious Diseases* 47(11), 1443–1448 (2008)
12. Sakaki, T., Okazaki, M., Matsuo, Y.: Earthquake shakes twitter users: realtime event detection by social sensors. In: *Proc. of the 19th International World Wide Web Conference, Raleigh, NC, USA*, pp. 851–860 (2010)

Web-Based Stress Management System Goes Mobile: Considerations of the Design of a Mobile Phone App

Åsa Smedberg¹ and H el ene Sandmark²

¹ Department of Computer and Systems Sciences, Stockholm University, Forum 100,
164 40 Kista, Sweden

asamed@dsv.su.se

² School of Health, Care and Social Welfare, Division of Public Health Sciences,
M alardalen University, 721 23 V aster as, Sweden

helene.sandmark@mdh.se

Abstract. High levels of stress at work, great demands and lack of balance between work and family; these are examples of components in the daily lives of many people in modern society. Interventions of different kinds are needed. However, changing patterns of thinking and behaviors is not an easy task, and it demands for continuous support and learning efforts. This paper presents a web-based stress management system and the way it was transferred to a mobile phone app. The solution and considerations are presented. The design principles of the web-based system were used to examine the use of the mobile phone app. The results showed that the app to a large extent meets the design principles, and some of the principles seemed even more applicable in the mobile phone setting due to its mobility.

Keywords: Mobile phone app, stress management, web-based, design principles, functions, stress areas.

1 Introduction

Today many people are on sick leave as a result of high levels of stress [1]-[2]. Studies show that high or unclear demands, low control and low social support at work increase the risk of "job stress". If this state of high levels of stress is prolonged, the risk of impaired health gets high. There may also be a disturbance in the balance between work and family when the demands of one sphere do not comply with the requirements of the other, or when the individual's efforts to fulfill its role in the work are affected negatively by the demands from family [3]-[4]. Social support at work or in private life has a mediating effect of perceived stress, and can to some extent, prevent negative stress.

Prevention and health promotion measures aim at creating attitude and behavioral changes in people who suffer from stress symptoms. Empowerment is a central concept for increasing control in life. An important question is how to build on and reinforce authentic participation, a sense of value and mastery in decision making.

As for individual efforts in stress management to deal with negative stress exposure, Jacobson was first to design a program of progressive muscle relaxation to guide people to achieve relaxation [5]. This can prevent stress reactions. It has been concluded that the effectiveness of interventions varies according to health-outcome measures being used [6]. However, to combine different techniques for muscle relaxation and cognitive-behavioral skills have seemed to be more effective across a variety of health-outcome measures than using a single technique [6].

When dealing with stress-related health problems the Internet can play an important role, with its facilities for exchanging information and support, as well as opportunity to have continuous interactions. Through web-based systems, people with high levels of stress can more easily come into contact with health professionals and peers for support and get support on a continuous basis. Self-help groups on the Internet have evolved in different health areas; stress and coping with stress constitutes one. There are also various types of web-based intervention systems that use different types of models for dealing with people's stress. However, available web-based systems today tend to focus on either information or communication, and to separate guidance from health experts from social support among peers.

In an attempt to close this gap, we have previously developed a web-based system that takes into account the benefits of having different actors and types of support coming together [7]-[9]. The web-based system has now also been transferred to a mobile phone setting. Since a mobile phone has certain characteristics that differ from a computer, new conditions follow. The advantage of being available to the user in most social situations need to be exploited, but also poorer properties such as a smaller screen and less performance need to be taken into account [10]. This paper presents considerations as well as preliminary results of the mobile phone app development.

In next sections, the web-based stress management system and its design principles are introduced. Thereafter, the mobile phone app is presented. The purpose was to let the mobile phone app be an extension of the web-based stress management system, designed in accordance with the principles used for the web-based system (presented in [7], e.g.). Therefore, we also compare functionality of the mobile phone app and the web-based system, and examine the mobile phone app based on the design principles used for the web-based system. The paper ends with some concluding remarks.

2 The Web-Based Stress Management System

To manage stressful life situations, to develop and maintain a healthy lifestyle and inner balance, can be difficult. E-health communities can contribute to this process through its ability to foster continuous interaction between members. Previous studies have shown that knowledge from health experts and experiences of peers can create a good synergy; the use of web-based ask-the-expert functions, between citizens and experts, and conversations among peers in e-health communities has shown to offer different and complementary support [11]-[12].

2.1 Design Principles

The web-based stress management system is based on eleven design principles generated through research. This section provides an introduction to these design principles.

1. Social and Learning Activities. Well-functioning web-based communities foster social and empathical communication [13]-[14]. At the same time, they need to support the learning processes regarding stress and how to manage stress. Both social and learning activities must therefore coexist.

2. Integration of Knowledge and Actors. Health experts and peers possess different kinds of knowledge and experiences and contribute differently to the system. If diversity is allowed in the system instead of having actors and knowledge kept separate, new understanding is more likely produced. To have direct communication with peers and also with health professionals can create a turning point for people trying to manage a stressful life situation.

3. Extended Relations. New materials and new ideas can be introduced when different health experts are invited to online conversations. Also, people with previous symptoms of stress, who have managed to recover from difficult situations in life, can contribute with their insights. By this, both strong-tie and weak-tie relationships are made possible [15].

4. Learning Conversations on Stress. Through questions and conversations about practical situations, experienced and anticipated difficulties and issues of motivation, users can get help in their learning processes and become more aware of them [12].

5. Tagged Conversational Topics and Acts. In order to allow users to navigate easily among conversations, topics and answers could be tagged by content and type of response. This could also help users to reflect on stress-related issues and to gain increased awareness of learning processes.

6. Variety Management in Stressful Situation. The system should help the users to boost their resources and strategies in order to better manage and prevent stress. This can be done in various ways, through a battery of actions, such as advice on how to approach a specific stressful situation, and how to filter the environmental signals and experienced demands, as well as guidance on how to perform physical exercises to relieve tension.

7. System Diversity and Variety. Individual differences in terms of personal stressors and reactions [16]-[17] must be handled by the system. Various stress-related areas can enable people with different needs to navigate between perceived stress situations, information sources, exercises and discussions.

8. Continuous Learning. Continuous efforts are needed for people who want to change unhealthy behaviors and patterns of thinking. Therefore, it is important that the system has mechanisms for ongoing support in everyday situations.

9. Viable Platform. The system in itself should be learning based on the users and user patterns. By identifying repeated conversation topics and requests from the users, feedback on what goes on in the users' lives is available, and the system can be

adjusted when necessary. For example, exercises can be adjusted to better fit the stress symptoms among the users.

10. Multimedia for Virtual and Real Life Integration. Communication between the system's medical staff and the ones with stress symptoms can be supported by multimedia applications. Breathing exercises can, for example, be easily demonstrated through online video clips, and also other physical exercises and instructions can benefit from using multimedia. The recorded exercises can be performed by the users whenever they choose. Other types of tools can be used to support the dialogue and make them richer.

11. Embedded Feedback Mechanisms. The system should include stress measurements such as physiological markers and self-reporting since they help measure the users' learning progress. Results from these tests could also be useful to identify additional needs that the system should meet.

2.2 The Web-Based System Design

The design principles presented in the previous section constitute the basis of our web-based stress management system. The system is a self-help system for stress management outgoing from the idea that both stress experts and peers should contribute with their different skills and experiences. The system provides the user with information in the form of research and real life stories, practical exercises (both text and image-based) and opportunities to have conversations with different groups of participants, with both stress experts (through ask-the-expert function and group counseling) and peers (through forums and group counseling). As the different active participants and the various types of support are combined in one system, we call the system a multiple-help system for stress management [7]-[9]. The stress management system is also divided into four different stress areas: Sleep, Work and study, Balance in life (work-study-life balance) and Physical well-being, based on previous research [3]-[4], [18], [19]. Variety among the users with stress symptoms has to be met by variety of the system.

3 Transfer to a Mobile Phone App

We now move on to describe the design of the recently developed prototype of a mobile phone app for stress management. The work involved how to use the basic ideas and functions in a mobile setting. The technical platform used is Android.

3.1 The Mobile Phone App Functions

When the mobile phone app is started, it displays the five different types of help or functions of the system together with explanatory text. There is also a button in the lower corner that provides brief information of the app and the research project. This home page is seen on next page (Fig. 1). The five types of help, or functions, available in the system are (from the top): Book counseling session, Research and real life stories, Ask-the-expert, Exercises and Forum.



Fig. 1. Home page of the mobile phone app for stress management (in Swedish)

The functions work as follows.

Book Counseling Session. The function displays the group counseling sessions related to the stress area selected via a spin component (android technique). If the user clicks on one of the counseling sessions, a dialog with information of the chosen counseling session is shown, and also a button to be used for signing up and another one for cancelling a session. Another dialog box will then appear in which the user can save the time and date of the session in the calendar on the phone. The user cannot see registered counseling sessions in the app but only in the phone calendar. The user is not able to attend the group sessions via the app; instead, he or she needs to be online in the web-based system to participate in the registered sessions.

Research and Real Life Stories. The user can choose to either read or listen to the documented research results and stories told that are related to the stress area selected. Depending on what options are selected in the tab sliding drawer, data from the web-based system are filtered and presented accordingly. The app is prepared to present information as text or played as audio sound. If sound is the option, the audio file is streamed from the web-based system.

Ask-the-Expert. The user can see the frequently asked questions and answers for each stress area and also post own questions. To submit questions, the user clicks on a button for "Ask question" presented in a sliding drawer at the bottom of the page. In a questionnaire, the user can specify the stress area in question, add a title to the question and type the question itself. At the moment, the app solution does not allow the expert's answer to be presented in the app; instead, the answer is handled through the web-based system and sent as a personal e-mail to the sender of the question. Eventually, it is also presented among the frequently asked questions and answers.

Exercises. The user can read, see or listen to the exercises related to the stress area being selected. As in "Research and real life stories", the app is prepared to support the formats selected by the user via buttons located in a sliding drawer. If the user clicks on an item in the list, text is displayed or video and/or sound streamed from the web-based system.

Forum. The user can actively participate in conversations with peers. As the web-based system uses a pre-programmed standard forum system, the phone app adds a link to the forum (via WebView) instead of retrieving the separate messages. Since the standard forum system, VanillaForums, supports mobile phone apps, the interface is automatically adjusted. The way it is displayed can be seen below (in Fig.2).

3.2 The Functions in Relation to Stress Areas

When one of the functions has been activated (“Forum”, e.g.), the system will display a new page with actions available, and also with all the five functions visible through a tab system (as presented in Fig.2). The five functions have been given the following tab names: *Råd* (=Advice): Booking of counseling sessions, *FoB* (=R&S): Research and Stories, *FAQ*: Ask-the-expert, *Övn* (=Exc): Exercises and practical tools, *Forum*: Forum for conversations.

The figure below shows also how the four different stress areas implemented in the web-based system are displayed. The stress areas are (from the top): Work and study, Physical well-being, Balance in life and Sleep. By clicking on one of the areas, conversations on topics related to it can be read and taken part in.




Fig. 2. Page showing the five different functions at the top. “Forum” is activated and forum discussions in the four stress areas are displayed via WebView (in Swedish).

The forum function of the system is treated differently since it uses a pre-programmed standardized forum system. The other four functions, however, have been able to be programmed from scratch and adjusted to the overall phone app interface. On next page is an illustration of the tab for *booking counseling sessions* (Fig.3). The chosen stress area is *balance in life*. In the text box, information about the scheduled counseling session is given (title, date, time and maximum participants).

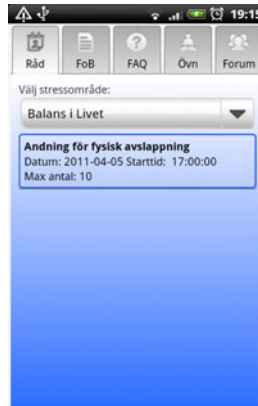


Fig. 3. Page for booking counseling sessions. It illustrates how a session is displayed in the area dealing with issues related to balance in life (in Swedish).

3.3 Comparison of Functionality

In order to find out about similarities and differences in functionality between the two systems – the web-based system and the mobile phone app - the main five functions of the systems were compared (see Table 1).

Table 1. Comparison between the web-based system and the mobile phone app regarding the main five functions

Function	Web-based system	Mobile phone app
Research & real life stories	Presentation of stories and research as both text and sound.	Ability to choose between reading or listening to selected stories and research (via 'sliding drawer').
Practical tools, exercises	Presentation as text, sound and video clips.	Presentation as for Research & real life stories. Video clips are also available.
Ask-the-expert	Frequently questions and answers are presented. The user can post new questions (and receive personal answers through e-mail).	Frequently questions and answers are presented. The user can post new questions (and read personal answers through e-mail).
Group counseling	Information on upcoming counseling sessions are displayed, and the user can sign up for one or more sessions. Information about registered sessions are displayed. The session takes place in a certain chat area in the system.	Information on upcoming counseling sessions are displayed, and the user can sign up for one or more sessions. The user can also add the session to his/her phone calendar.
Forum	The forum system offers a participation and displays questions and answers between peers.	As the web-based system uses a pre-programmed standard forum system, the phone app offers links to the forum system (via WebView).

3.4 Evaluation Based on the Design Principles

In addition to comparing functionality between the two systems, the design principles of the web-based system were used to conduct a minor evaluation of the mobile phone app. This is a summary of the evaluation results.

Social and learning activities take place through the app. Linking the app to the web-based forums and ask-the-expert function sessions allows for communication among peers, and also with experts. To allow for **integration of knowledge and actors**, the app communicates and displays information, exercises, forum discussions, etc., to and from the web-based system. Thereby, the two systems together let both health experts and peers contribute in different ways. Expert answers to own posted questions still need to be worked on technically in order to have them displayed in the mobile phone. **Extended relations** are possible in both environments. Through the app, **learning conversations on stress** among peers (on practical situations, difficulties, motivation, etc.) are possible to follow and participate in. The question of **tagged conversational topics and acts** is to be dealt with by the web-based system (standard forum system). A battery of actions are available that allows the user to be better equipped for dealing with stressful situations, and thereby increasing his or her **variety management in stressful situations**. This is also related to empowerment among the users. Also, individual differences are handled by the system through its stress areas (the same as in the web-based system) and its different types of help, allowing for **system diversity and variety**. Furthermore, **continuous learning** is something even more emphasized and evident in the mobile phone setting, where ongoing support in everyday situations is possible. The app is strongly related to the web-based system and should be adjusted to changes in the web-based setting, thereby allowing for a **viable platform**. The app is designed to be prepared for addition of stress areas, e.g. The app uses **multimedia for virtual and real life integration** when demonstrating physical exercises, such as breathing techniques through online video clips. The mobile phone setting allows the user to view exercises independent of both time and place. So far, **embedded feedback mechanisms**, in the form of stress measurements, are to be handled by the web-based system and will eventually be available in the phone app.

4 Final Remarks

In this paper, we have presented a prototype of a mobile phone app for stress management associated with a web-based system. So far, the mobile phone app has shown promising results when evaluated by criteria from research in the area of stress management and online support. This preliminary work will be followed by user-based evaluations to identify needs to be addressed in the next iteration of design.

References

1. Henderson, M., Glozier, N., Holland, E.K.: Long Term Sickness Absence is Caused by Common Conditions and Needs Managing. *BMJ*. 330, 802–803 (2005)
2. Dekkers-Sánchez, P.M., Hoving, J.L., Sluiter, J.K., Frings-Dresen, M.H.: Factors Associated with Long-Term Sick Leave in Sick-Listed Employees: A Systematic Review. *Occup. Environ. Med.* 65, 153–157 (2008)

3. Sandmark, H.: Work and Family: Associations with Long Term Sick-Listing in Swedish Women. *BMC Public Health* 7, 287 (2007)
4. Sandmark, H.: Job Mismatching, Unequal Opportunities and Long-Term Sickness Absence in Female White Collar Workers in Sweden. *Scand. J. Public Health*. 37, 43–49 (2009)
5. Jacobson, E.: *Progressive Relaxation*, 2nd edn. University of Chicago, Chicago (1938)
6. Murphy, L.R.: Stress Management in Work Settings: A Critical Review of the Health Effects. *American Journal of Health Promotion* 11, 112–135 (1996)
7. Smedberg, Å., Sandmark, H.: Multiple Help Online: An Integrated E-Health System for Stress Management. In: *Proceedings of IADIS International Conference E-Health*, Freiburg, Germany, July 29-31, pp. 151–158 (2010) ISBN: 978-972-8939-16-8
8. Smedberg, Å., Sandmark, H.: Stress Intervention Online - Designing for Self-Help through Multiple Help. In: *Proceedings of the Third International Conference on eHealth, Telemedicine, and Social Medicine*, Guadeloupe, France, February 23-28, pp. 120–125 (2011) ISBN: 978-1-61208-003-1
9. Smedberg, Å., Sandmark, H.: Dynamic Stress Management: Self-Help through Holistic System Design. In: Smedberg, Å. (ed.) *E-Health Communities and Online Self-Help Groups: Applications and Usage*, pp. 136–154. IGI Global (2012)
10. Jones, M., Marsden, G.: *Mobile Interaction Design*, 1st edn. Wiley-VCH, Chichester (2005)
11. Smedberg, Å.: To Design Holistic Health Service Systems on the Internet. *Proceedings of World Academy of Science, Engineering and Technology*, 311–317 (November 2007)
12. Smedberg, Å.: *Online Communities and Learning for Health - The Use of Online Health Communities and Online Expertise for People with Established Bad Habits*. Doctoral Thesis. Department of Computer and Systems Sciences, Stockholm University (2008) ISBN: 978-91-7155-689-9
13. Preece, J.: Empathic Communities: Balancing Emotional and Factual Communication. *Interacting with Computers* 12, 63–77 (1999)
14. Preece, J.: *Online Communities – Designing Usability, Supporting Sociability*. Wiley & Sons (2000)
15. Haythornthwaite, C.: Social Networks and Online Community. In: Joinson, A., McKenna, K., Postmes, T., Reips, U.-D. (eds.) *The Oxford Handbook of Internet Psychology*, pp. 121–137. Oxford University Press, United Kingdom (2007)
16. McEwen, B.S.: The Neurobiology of Stress: From Serendipity to Clinical Relevance. *Brain Res.* 886, 172–189 (2000)
17. Levi, L.: Working Life and Mental Health - A Challenge to Psychiatry? *World Psych.* 4(1), 53–57 (2005)
18. Siegrist, J., Marmot, M.: Health inequalities and the psychosocial environment—two scientific challenges. *Soc. Sci. Med.* 58, 1463–1473 (2004)
19. Preckel, D., von Känel, R., Kudielka, B.M., Fischer, J.E.: Over Commitment to Work is Associated with Vital Exhaustion. *Int. Arch. Occup. Environ. Health* 78, 117–122 (2005)

Use of SMS for Tsunami Early Warnings at a Table Top Exercise^{*}

Vincenzo Lagani, Vasilis Kontogiannis,
Panagiotis Argyropaidas, and Catherine Chronaki

Foundation for Research and Technology – Hellas (FORTH)
N. Plastira 100 Vassilika Vouton, GR-700 13 Heraklion, Crete, Greece
{vlagani, vasilik, parg, chronaki}@ics.forth.gr

Abstract. Tsunamis are one of the most dangerous and destructive natural disasters: countries that experience a tsunami event are likely to undergo, after the immediate destruction of the regions nearby the coast, several secondary effects, for example epidemic outbreaks. Thus, early warning systems able to timely advise the authorities and the population of an imminent tsunami are valuable tools that should be implemented in any coastal region with high risk of seismic events. In this paper we present the experience that we gained during a Table Top exercise aimed at testing the effectiveness of an SMS – based Tsunami Early Warning System (TEWS). SMS showed to be a valuable additional support channel for propagating the alarm, even though some drawbacks were identified e.g. lack of reliability, formal procedures and follow-through training that must be carefully taken in account. Experience gained will be used in an operational exercise scheduled October 2011 and potentially pave the way for including SMS early warning in disaster management.

Keywords: Early Warning System, SMS, Tsunami, Disaster Management, Emergency Response.

1 Introduction

Recently, two extraordinary Tsunami events (Indian Ocean, 2004 and Pacific Coast of Japan, 2011) remind the entire world of the potential destructiveness of natural disasters. Tsunami consequences do not run out in the immediate term; on the contrary, long term effects of Tsunamis are even more devastating than the brute destruction spread by the giant waves. In particular, Tsunamis often create unsanitary health conditions that are likely to breed serious diseases and epidemics. Currently, the only countermeasures for contrasting Tsunami events consist in (a) developing emergency programs for population evacuation and (b) creating Tsunami Early Warning Systems (TEWS) that are able to timely detect the arrival of a Tsunami for

^{*} This work was supported by the European Commission, Humanitarian Aid and Civil Protection Directorate under contract 070401/2009/534360/SUB/A3: “POSEIDON: Earthquake followed by Tsunami in the Mediterranean Sea”.

advising both the authorities and the civil population to take action in view of the imminent risk.

Real world, operating TEWS have been deployed in several countries after 2004, for example in India [1] and Indonesia [2]. Interestingly, most recent TEWS started to employ SMS services for spreading alert messages; however, using SMS for broadcasting disaster warnings is still a controversial issue. Thanks to the wide diffusion of mobile phones, SMS can easily reach a large number of recipients. Recent studies have noted the possibilities of using SMS and MMS to provide personalized, media-rich, context-aware advice [3,4] and enhance the understandability of the alert for visual/hearing impaired citizens [5]. However, SMS technology is intrinsically unreliable, and sending large numbers of SMS immediately after a disaster may contribute to overloading the communication network [6].

In this framework, our research group along with National Civil Protection Authorities, the Red Cross, and the regional Government, recently engaged in a new initiative, the POSEIDON project [7], with the aim of conducting a European civil protection exercise to train in the European Civil Protection Mechanism that includes use of Information and Communication technology in population evacuation in case of Tsunami. The operational field exercise will be held in October 2011, in the island of Crete. For our simulation we plan to employ SMS as supplementary communication channel for advising both competent authorities and the general population about the imminent Tsunami. One of the preliminary activities of the project consisted in performing a Table Top exercise (May 2011), with the objectives of both strengthening the connections among the actors involved and refining the organization of the simulation. We sent, through SMS, a Tsunami alert to all the participants of the Table Top exercise that accepted to be involved in this first experimentation. After sending the SMS, we collected the impression and the comments of the recipients with a questionnaire prepared in advance. Our scope was to exploit this first experience in planning the SMS alert system for the field exercise.

2 Methodology

Setting of the TEWS Simulation during the Table Top Exercise: The tabletop exercise was hosted in the main building of the Decentralized Administration of Crete, on May 30, 2011. Close to one hundred (100) Representatives of Health Authorities, Civil Protection, Firefighters, Police, Port Authorities and Cretan Municipalities were present. Among the Table Top participants, eighty-four (84) accepted to participate in the simulation: the participants' role consisted in receiving two alert SMS on their mobile phones and subsequently answering a questionnaire. The first message contained a Tsunami alert for the Municipality of Heraklion, while the second one for the Municipality of Chania. Both messages contained a location – dependent hyperlink pointing at a map of the coastal area at risk. Messages were sent either in English or Greek, depending by the nationality of the recipient. The prototype .NET application for sending the messages, consisting in a Graphical User Interface connected to the Skype software via an application programming interface that allowed maintaining the list of contacts, their role, and language preferences.

Evaluating the Results from the Table Top Exercise: We measured the performance of the SMS service under two different aspects: (a) *technical effectiveness*, consisting in evaluating whether the SMS service succeeded in timely delivering all planned messages, and quantified as the percentage of sent SMS over the number of planned messages (%SENT) and the percentage of participants that actually received the SMS over the number of intended recipients (%DEL); (b) *informativeness*, much more complex to measure, since it involves the interaction between the reader and the message itself. For evaluating SMS informativeness we devised an ad-hoc questionnaire subdivided in the following sections (dimensions): *understandability*, *credibility*, *usability* and *usefulness*. Questions are structured as five points items, ordered from “totally negative” (score: 1) to “totally positive” (score: 5). The questionnaire is available online at: <http://kwiksurveys.com/?u=poseidon>

3 Results

Out of one-hundred sixty eight (168) planned messages (eighty-four (84) recipients, two messages each), we managed to send one-hundred nine SMS (%SENT = 64.9%) to sixty two (62) distinct recipients. At the end of the table-top exercise, we collected twenty eight (28) filled and four (4) blank questionnaires (thirty two (32) in total). We estimated the percentage of delivered messages as $\%DEL = 28/32 = 87.5\%$. After the questionnaires collection, we evaluated the results as follows: for each dimension D we calculated a normalized score as the average score of the questions associated to D. We then calculate the mean of the normalized scores across all questionnaires. This procedure produced a single, mean normalized score for each dimension, ranging from 1 to 5, where “1” indicates that the participants’ average evaluation of the message was completely negative, while “5” corresponds to a totally positive assessment. In particular, the score for each dimension are the following: Understandability, 3.66; Credibility, 3.58; Usability, 3.67; Usefulness, 2.26.

4 Discussion and Conclusion

The scope of our first experience in using SMS for TEWS during a Table-Top exercise consisted in detecting possible technical and organizational problems to be solved before the wider and more complex POSEIDON field exercise. From this perspective, our simulation was a success: the outcomes of the exercise pointed out a number of unexpected problems that otherwise would have never been identified. It was confirmed that exercises are the best way to prepare for the unexpected.

The first difficulty that we experienced was a malfunctioning of the Skype software; due to this blatant technical problem, our simulation plan was substantially subverted. This inconvenient taught us that SMS technology in general is not reliable, even for sending a few hundreds of messages, especially if an external application (Skype) is used as SMS gateway. However, SMS unreliability does not only reside into the application employed for sending the messages; any part of the technical infrastructure used for transmitting the SMS can fail. Despite SMS low reliability, the results of our simulation also underlined that SMS can still be a useful media for diffusing Tsunami alerts. The participants generally rank as “sufficient” to “good” the

Understandability, Credibility and Usefulness of the message. This result suggests that, once the message is received, it can be an effective tool for informing the recipient about the imminent risk of Tsunami. The only dimension that did not pass the threshold of sufficiency was “Usefulness”, the questionnaire dimension measuring the helpfulness of the accessory information provided by the hyperlinks; we hypothesize that the difficulties in accessing the hyperlinks may be due to a lack of training in processing and acting upon the information in the SMS message. Thus, the results of the Table Top exercise suggested us the following actions:

- the reliability of the SMS service must be improved, at least as far as concern the SMS gateway;
- the content of the message must be improved, especially regarding the accessibility of the multimedia information;
- participants (i.e. civil workers) and the public (should the service be expanded) need to be prepared and trained to receive and process SMS information;
- Alternative means of communicating the information in the SMS message need to be explored in tandem. SMS messages alone are not adequate for public awareness and informed response to disasters.

In synthesis, this experience confirmed that SMS messages seem to be a valuable tool for diffusing Tsunami alerts. However, the low reliability of SMS technology indicates that SMS can only be employed as a *complementary* medium along with other communication and awareness methods. Civil protection workers and the public need practical training on action upon SMS receipt through customized profiles that are included in preparedness exercises, to cultivate trust and improve effectiveness of response and follow-through based on the individual role and skills.

References

1. Kumar, T.S., Kumar, C.P., Nayak, S.: Performance of the Indian Tsunami Early Warning System. In: International Archives of the Photogrammetry, Remote Sensing and Spatial Information Science, Kyoto, Japan, vol. XXXVIII, part 8 (2010)
2. Lauterjung, J.: Installation of a tsunami early warning system in the Indian Ocean. In: Proceedings of the 13th Annual ACM International Workshop on Geographic Information Systems (GIS 2005), p. 1. ACM, New York (2005)
3. Meissen, U., Voisard, A.: Increasing the Effectiveness of Early Warning via Context-aware Alerting. In: Proceedings of ISCRAM Conference, Washington, DC, USA (May 2008)
4. Malizia, A., et al.: CAP-ONES: An Emergency Notification System for all. In: Proceedings of the 6th International ISCRAM Conference, Gothenburg, Sweden (May 2009)
5. Mitchell, H., Johnson, J., La Force, S.: Wireless Emergency Alerts: An Accessibility Study. In: Proceedings of the 7th International ISCRAM Conference, Seattle, USA (May 2010)
6. Pau, L.F., Simonsen, P.: Emergency messaging to general public via public wireless networks. In: Proceedings of ISCRAM Conference, Washington, DC, USA (2008)
7. Evangelia, T., et al.: Risk Assessment Study based of the 365AD Earthquake to Drive a Civil Protection Exercise. ERCIM News (April 2010), <http://ercim-news.ercim.eu/en81/special/risk-assessment-study> (visited on July 11, 2011)

Seamless Evaluation of Interactive Digital Storytelling Games: Edugames4All

Patty Kostkova

City eHealth Research Centre
City University London
patty@soi.city.ac.uk

Abstract. Interactive Digital Storytelling (IDS) is a growing hybrid discipline bringing together computer games and cinematic storytelling enabling entertaining education. CeRC developed IDS game “Edugame4All” to improve children understanding of hygiene and responsible antibiotic use. This short paper presents an evaluation results assessing knowledge gain against the project prescribed learning objectives and proposes a novel framework called “seamless evaluation” inserting fine-grain questions into the story narrative and introducing a debrief session to evaluate without decreasing players’ immersion.

Keywords: educational web games, game evaluation, interactive storytelling.

1 Introduction

Serious games are increasingly becoming established as a new method of education [1]. IDS is a growing hybrid discipline bringing together computer games and cinematic storytelling enabling a creation of unique learning applications [2]. The positive use of games to aid learning has been established [3-4]. The Edugames4All project (www.edugames4all.org) based on the DG SANCO-funded e-Bug project (www.e-bug.eu) developed a story-telling game to teach children basic principles of hygiene and antibiotic resistance [5]. As a traditional ‘role-playing game’ it relies heavily on narrative [2]. There were 3 missions implemented, each teaching a set of learning objectives (LOs) prescribed by the project. In this paper we discuss the evaluation results and propose a seamless evaluation utilizing the storytelling nature of the game and engaging the player in a dialogue with the characters allowing to fine-tune the understanding of the player while not disturbing the immersion.

2 Interactive Story Telling and Edugames4All

Unlike linear digital storytelling, an interactive storytelling makes the player “feel” he/she is changing the course of the game. As defined by O’Hagan, the diagram in Figure 1 illustrates the 4 levels of the game.

2.1 Edugames4All: Case Study

The IDS mechanics was chosen based on focus groups conducted with 13-15 years old children in the UK [6]. There are three puzzles to solve, each testing several LOs.

Here we describe the first puzzle based around hand hygiene related LOs (the 2nd and 3rd teach appropriate use of antibiotics and the issue of antibiotic resistance). Referring to standard IDS levels, defined in [1], illustrated on Figure 1, in the “animation base”, the player is presented with a scene, chooses narrative to talk to characters, collects evidence, investigates evidence in a laboratory and presents an answer to “puzzle” – the cause of infection or reason for an illness – to the boss of e-Bug (“Big C” character, Figure 2). At the second “dialog base”, there is the story. For example, the story in the first puzzle is as follows: at a BBQ party, a famous actor gets by infection transmitted as a result of poor hand hygiene at a bathroom but poisoning food and insufficient hygiene in the kitchen must be eliminated by investigating a Chef and by collecting and testing evidence samples. Investigate dialogue in the kitchen scene is illustrated on Figure 3.




Fig. 1. 4 levels of interactive story-telling (taken from [1])



Fig. 2. Big C Fig. 3. Chef investigation narrative Fig. 4. “Micro-vision”

The game mechanics “rules” require the player to test the samples and prove the hypothesis. This includes so called “micro-vision” (illustrated on Figure 4) turning the room into a “microbe-vision light” allowing users to see microbes on the scene and collect samples for testing in a microbiological laboratory. The evidence is tested and hypothesis about the cause of the infection either proved or disproved. Finally, on the 4th level - “plot outline” - the story “Big C at the Investigators HQ assigns player a task to investigate the reason of poisoning and approves the correct investigation outcome: the Actor’s sickness was caused by a contamination of the salad at the BBQ by microbes transmitted due to poor hand hygiene at a toilet.

3 Evaluation Results

The games were evaluated during the development phase in schools in the UK in terms of knowledge change and usability. Final game was evaluated in summer 2010

Table 1. Summary of player responses for questionnaires for each Mission, including the specific questions provided at each level corresponding to e-Bug pack learning outcomes. The pre and post responses are written via symbols between the ‘&’ where C=correct, W=wrong, ?=other (either answer not provided or both True and False selected – wrong in either case). For example, ‘C&C’ = correct pre and correct post answers; ‘C&(W|?)’ = correct pre-answer and either wrong post-answer or other response selected. ‘Remaining’ refers to the tallied player responses for the following cases: ‘?&?’ , ‘?&W’ , and ‘W&?’.

Mission	Question	# Kids Started	# Kids Finished	C&C	W&W	(W ?)&C	C&(W ?)	Remaining	McNemar chi-square	p value
Mission 1	1. Viruses and bacteria are two types of micro-organisms.	22	10	16	1	4	1	27	1.25	0.26
	2. If you cannot see a microbe in plain daylight, it is not there.	22	10	18	0	3	1	27	0.56	0.45
	3. Fungi are microbes.	22	10	17	2	3	0	27	2.08	0.15
	4. Virus microbes can be spread by sneezing or coughing.	22	10	18	0	2	2	27	0.06	0.80
	5. Virus microbes can be spread by sneezing or coughing.	22	10	19	0	0	3	27	4.08	0.04
	6. Virus microbes can spread if people don't wash their hands after using the loo.	22	10	20	0	1	1	27	0.13	0.72
	7. Washing vegetables and cooking meat removes bad microbes from our food.	22	10	19	1	1	1	27	0.13	0.72
	8. People should wash their hands before eating.	22	10	21	0	0	1	27	2.25	0.13
	9. Before making a meal, people don't need to wash their hands.	22	10	20	1	0	1	27	2.25	0.13
	10. Microbes are found on most surfaces like your skin, or food that you eat.	22	10	19	0	2	1	27	0.08	0.77
Mission 2	1. Antibiotics kill bacteria.	14	10	10	1	1	2	35	0.75	0.39
	2. Microbes cannot be spread by sneezing or coughing.	14	10	6	0	2	6	35	2.53	0.11
	3. Vaccines help protect people against some microbes.	14	10	11	0	0	3	35	4.08	0.04
	4. Bacteria and viruses are the same.	14	10	11	0	0	1	37	2.25	0.13
	5. Antibiotics kill viruses.	14	10	3	4	5	1	36	2.04	0.15
	6. Most coughs and colds get better without antibiotics.	14	10	9	1	3	1	35	0.56	0.45
	7. If antibiotics are overused, bacteria can mutate to become resistant to the drugs.	14	10	13	1	0	0	35	N/A	N/A
	8. You should only use antibiotics with your doctor's permission.	14	10	12	2	0	0	35	N/A	N/A
	9. Overuse and abuse of antibiotics leads to antibiotic resistance.	14	10	13	1	0	0	35	N/A	N/A
	10. Antibiotics are needed to treat bacterial illnesses and to make your feel better.	14	10	11	1	2	0	35	1.13	0.29
Mission 3	1. Microbes are contagious.	22	13	17	0	3	2	27	0.05	0.82
	2. If you cannot see a microbe in plain daylight, it is not there.	22	13	20	0	1	1	27	0.13	0.72
	3. Bad microbes can spread when you touch something or someone that is sick.	22	13	18	1	3	0	27	2.08	0.15
	4. If people wash their hands, they are less likely to get ill.	22	13	19	0	3	0	27	2.08	0.15
	5. You should only use antibiotics with your doctor's permission.	22	13	19	0	0	3	27	4.08	0.04
	6. You can stop taking antibiotics when you are feeling better.	22	13	7	5	7	3	27	1.23	0.27
	7. Finishing your course of antibiotics ensures that all bacteria is out of your system.	22	13	7	6	6	3	27	0.69	0.40
	8. You should never give your antibiotics to a friend to use.	22	13	22	0	0	0	27	N/A	N/A
	9. People should take antibiotics to prevent them from getting sick (like a vaccine).	22	13	12	3	3	4	27	0.32	0.57
	10. Antibiotics kill our good bacteria in our gut.	22	13	10	2	7	3	27	1.23	0.27

by 49 students attending “Technology Camp,” a summer school held at Elstree School, near Reading, Berkshire. Their age range was 10-13 year olds including both girls and boys. The study took two 1.5 hour timeslots provided for each of the two sets of children (divided at the convenience of the summer school counselors). The results of the evaluation using pre and post questionnaire, administered on paper before and after children played the game, were analyzed using t-pair test and McNemar test, as illustrated in Table 1. We observed a large dropout rate (which is not unusual and was observed at the junior game evaluation [6]). Mission 1 – 22 started, 10 finished Mission 2 – 14 started, 10 finished, Mission 3 – 22 started, 13 finished. Further Several students played more than one mission – hence the higher starting numbers for Missions 1 and 3. On average (49 students divided into thirds), approximately 16 students should have played each Mission. The prescribed Learning Outcomes were mapped to individual Missions as requested by the project, however, as they were the least common denominator across European science curricula these proved to be too simple and did not provide a sufficient level of granularity. Table 1 illustrates most kids at this science camp knew correct answers beforehand, thus, we could measure little significant knowledge change. A negative result was achieved - $P=0.04$ for three LOs; however, the knowledge change was from correct to incorrect. The problem was that the majority of students knew the answers before game playing as the level of biology and science knowledge at the camp was well above average. Secondly, the LO set to allow unification of the project aims cross Europe proved too simple to test a deeper understanding but provided a baseline for a novel finer-grained seamless evaluation.

4 Work in Progress: Seamless Evaluation

In addition to designing the LOs at the right level of depth to better evaluate the educational impact of games, the administration of tests needs to be "seamless" by inserting questions in the game mechanics in order to decrease an exposure suspicion bias while maintaining user immersion and enjoyment. The need for innovative non-intrusive approach to evaluation has been highlighted [8], however, to actually design a seamless evaluation framework an appropriate design methods matching the mechanics of evaluating game are required. In our previous research, a seamless evaluation of the Edugames4All platform game was achieved by inserting a “How to be a Millionaire” scenes (player was asked questions in a competitive quiz format) before and after playing a bacteria fighting session in human body [5]. IDS games are providing a nature format for seamless evaluation. Further, to enhance the leaving effect a “debriefing” phase is to be inserted providing reflections on one’s decisions by observing the game from critical distance [3]. Alex, the “investigator” characters in Edugames4All, has to return to the Mission HQ and report each decision back to Big C. These two implementations of seamless evaluation are currently in tested on a Global Handwashing Day Game [7] and preliminary results will be presented at the conference.

5 Conclusion

Computer games have been established as new media for education. We described the design of the IDS game Edugames4All teaching hygiene and antibiotic resistance and the results of an evaluation using pre and post questionnaires methodology. A work in progress method, the seamless evaluation, integrating fine-grained questions into the game dialogue and introducing a “debrief” session evaluates knowledge gain without decreasing players’ immersion.

Acknowledgements. Thanks to DG SANCO, David Farrell, L Lazareck, D Weerasinghe, e-Bug team lead by C McNulty, J Weinberg, D Lecky; and artists N Lai, S Beveridge.

References

- [1] O’Hagan, M.: Video games as a new domain for translation research, *Revista Tradumàtica – Traducció I Tecnologies de la Informació I la Comunicació 05: Localització de videojocs*, ISSN 1578-7559, <http://www.fti.uab.cat/tradumatica/revista>
- [2] Spieerling, U.: Interactive Digital Storytelling: Towards a Hybrid Conceptual Approach. In: *Proceedings of DiGRA 2005 Conference: Changing Views – Worlds in Play (2005)*
- [3] Klabbers, J.: The gaming landscape: A taxonomy for classifying games and simulations. In: Copier, Raessens (eds.) *Level up: Digital Games Research Conference*, Utrecht Uni. (2003)
- [4] Robertson, J., Collins, T., Kelly, A.: *Tales of Adventure Author: Learning through Making Games. Narrative and Interactive Learning Environments*, Edinburgh (2008)
- [5] Farrell, D., Kostkova, P., Lazareck, L., Weerasinghe, D.: Developing web games to teach microbiology. *The Journal of Antimicrobial Chemotherapy* 66(suppl. 5), v33–v38
- [6] Farrell, D., Kostkova, P., Weinberg, J., Lazareck, L., Weerasinghe, D., McNulty, C.A.M.: Computer games to teach hygiene: An evaluation of the e-Bug junior game. *The Journal of Antimicrobial Chemotherapy* 66(suppl. 5), v39–v44
- [7] <http://www.edugames4all.org/IntegratedCRD.nsf/a2e0dec55e90eaac80257914005ad890/c7b7205052d61c578025791400661778?OpenDocument>
- [8] Vavoula, G., Meek, J., Sharples, M., Lonsdale, P., Rudman, P.: A Lifecycle Approach to Evaluating MyArtSpace. In: *Proc. 4th Int. Workshop of WMUTE 2006*, pp. 18–22 (2006)

Patient-Centered Care as a Learning Process

Jan Aidemark and Linda Askenäs

Linneaus University
351 95 VÄXJÖ
Sweden

{Jan.aidemark,Linda.askenas}@Lnu.se

Abstract. The focus of this paper consists of investigations into a strategic planning framework for information systems in support of patient-centered care (PCC) processes. The planning perspective that underpins the research includes learning theories, organizational learning and knowledge management in general. A brief review of current PCC goals and perspectives is used as a starting point for an investigation of PCC activities and support system. The examples of PCC activities are organized as a learning process and presented in a framework.

Keywords: Patient learning, Information systems, Planning framework, E-health.

1 Introduction

Patient-centered care (PCC) is an emerging approach to health care, on which a great deal of hope has been pinned as it indicates lower costs and better care quality (Laine and Davidoff, 1996). PCC is a break in a traditional disease-oriented model of care (Epstein, 2007). However, a coherent definition of PCC is still to appear, one recent effort being given by Robinson et al. (2008, p. 600) according to which “fundamental characteristics of PCC” include “a) patient involvement in care and b) the individualization of patient care”. In the center of this definition lies the interaction between patient and caregiver in the hope of a better quality of care, increased patient satisfaction and higher patient adherence (Robinson, et al. 2008, p. 601). This makes PCC an approach for how to design and perform health care which contains implications for every aspect of the health care process. In this paper we focus on how information systems and information technology (IS/IT) could be used to support a PCC strategy towards health care. Our approach is to build on a learning tradition within IS/IT and to apply this to the task of planning and designing IS/IT support for PCC activities and processes. The contribution of this paper is a literature review that focuses on the connection between pressing PCC issues and current theoretical frameworks, concept and models. The study evolves into an outline of knowledge- and learning-oriented planning and design framework for PCC support.

2 Background: Patient-Centered Care (PCC)

A number of possible aspects of the concept of patient-centered care (PCC) have been identified. A fairly early expression of patient-centered care was made by Sherer et al. (1993), who defined it as a “[d]esign of patient care wherein institutional resources and personnel are organized around patients rather than around specialized departments”. This can be seen as an organizational approach to achieving patient-centered care. From economic aspects a number of benefits would be accrued from a patient-centered approach. Charnel and Frampton (2008) argue that there is a solid business case for PCC, with great possibilities for lower costs as well as for an increased market share. Robinson et al. (2008) adopt four PCC perspectives on: 1) The public policy perspective, focusing on the partnership among practitioners, patients and their relatives, and on the care decisions that are made, so that these decisions follow the needs and wants of the patient and enable the patient to take part in these decisions in an informed and knowledgeable way. 2) The economic perspective, which focuses on patients as consumers of health care and on their competence to make informed decisions with respect to various quality characteristics of the care. Robinson et al. 2008 mention as one example the possibility of direct-to-consumer marketing of pharmaceuticals, and how this connects to the patient as decision-maker involved in care decisions. 3) A clinical practice perspective, focusing on the view of the caregiver and on the changes in how, for example, practitioners understand, interact with and treat their patients. The PCC approach affects how the caregiver understands patient needs and tailors the care to be given as well as generally increasing the humaneness of the care. 4) The patient’s perspective, finally, focuses on patient priorities, which might include factors like respect, courtesy, competence, efficiency, decision-making, time for care, availability, accessibility, wants/needs of care and communication/information. To handle the complexity of PCC models have been developed of what could be included in a PCC care process. Steward et al. (1995) propose a model of six interacting components: 1) exploration of the illness/disease and the patient’s experience of it, 2) understanding the patient as a whole person, 3) planning care, goals of treatment and areas of responsibilities, 4) work with prevention and health promotion, 5) creating in the patient-doctor relation a working therapeutic relationship including power sharing, and 6) a realistic and effective use of time and money. Building on this model, Aita et al. (2005) suggest a somewhat generalized and more concise one for PCC. This model contains four components: 1) the physician perspective, including style, value or medicine philosophy, 2) the patient perspective, including priorities, values and the philosophy of health, 3) practice organization, including priorities and culture, and 4) the community aspect, including culture and priorities.

3 Investigations of a Patient-Centered Care Model

3.1 The Condition for Understanding - What Is Knowledge

The focus in support of care processes is often on fact, information and messages that are directed towards the patient. The move towards learning/knowledge puts the focus

on what happens with the facts when they have reached the patient, and on what it takes for facts to become actions. In the core of this we find the concept of knowledge and the processes that surround it, like creation, storage, transfer, use and so on, in short, what a patient understands from the facts presented. This is a process going on in all the aspects, including measures, planning, choices, actions and evaluations. Forbat et al. (2009) discuss the nature of PCC and its key feature of patient involvement in the care process, which means that the patient is part of planning, decision-making, delivery and evaluation. However, this involvement is not enough, but the patient must really be engaged in the care process. This requires that the patient gains a personal understanding of the sickness, and this does not mean just an information-seeking process. Forbat et al. (2009) argue that patient and care personnel must be coached into being able to be engaged. The method which Forbat et al. (2009) suggest builds on study groups of patients and staff that work according to a Plan Do Study Act (PDSA) method as a support for learning and building personal understanding.

3.2 Measuring – Gathering Facts

Gathering facts by taking measurements on the patient is central for making plans and decisions about the continued care process. In PCC this has been translated into a need of personalized measurements that capture the individual perspective, or, how the patient perceives his or her situation. The outcome of a measurement might be different whether a patient or a physician makes the assessment (for example, Stephens, 1997). In the care of chronic illness special measures for understanding the long term progress of the individual have been developed. The typical example of this is the “quality of life” (QoL) concept. Davidson et al. (2004), researchers on heart failure, argue for the task of assessing patient needs as being central to tailoring the care for the individual patient. The approach builds on structured measures to capture this information, including for example patient satisfaction, quality of life and utility. But the key is how to individualize these measures, by involving the patient and making him or her take active part in the measurements, and turning it into a long-term process and not a one-time happening. In their paper the authors argue for a general framework and for developing specific and quantitative needs assessment questionnaires within each area. They also point towards a number of problems with the QoL measure. First, this concept is not properly defined, which makes it somewhat less reliable even though it often includes factors like physical functions, psychological processes, social and economic concerns, and spiritual/existential aspects.

3.3 Planning – Exploring Options for Care

A great deal of what precedes the actual decision and execution of individual care could of course be seen as care planning. From a decision theory aspect the planning phase is about generating options which are to be decided on later. An example of this could be found in a study of a goal-setting support for diabetes patients (Langford et

al., 2007). In this case a framework for a self-management goal cycle has been set as support for the patients. The central activities include: 1) visits, both personal and group meetings, also including scheduled or ad hoc telephone calls; 2) goal lists, i.e. lists of things the patient should achieve; 3) measurements, including one measure for the likelihood of reaching a goal and one for measuring the ability of the caregiver to solve problems associated with the goal; 4) a checklist for the caregiver as support when gathering facts about the situation; and 5) problem-solving activities when a problem arises about the goal.

3.4 Decision-Making – Making Care Choices

The decision on what care a patient needs is the centerpiece of the care process. A patient who is well informed and is thus empowered to participate actively in the decision plays an important part in a patient-centered care model. The preceding activities build up to this moment, the patient-doctor meeting, facts and measures, the planning and the learning processes all supporting an empowered patient. Great advantages are expected from a patient-centered decision including a focus on issues that are important to the patient, decisions made in line with patient values, the improvement of patient compliance with and commitment to the care. In the long run this would also mean a lower total cost for care. Shared decision-making is a concept advocated in a brief case study (Walker, 2008) as part of a patient-centered care model. Walker (2008) defines shared decision-making as “[t]he collaboration between patients and caregivers to choose treatment options in line with patient health plans”. The patient in the case becomes informed by support like counseling for good information gathering and understanding, videos with treatment choices and facts in form of a treatment handbook. Gustavsson et al. (1999) present a study of a support system that was provided to a group of patients for home use. The patients were provided with support like information, decision support, and connections to experts and other patients. The impact of the system was measured with self-reports of their quality of life and of the frequency and duration of their use of medical services. The benefits included, for instance, patients spending less time during care visits, communicating with health providers by phone and experiencing fewer and shorter hospitalizations.

3.5 Action and Learning: Connecting the Cause and Effect of Self Care

Taking the action that is the actual care is of course the goal of the care process. The actions can be taken by patients more or less on their own, or be performed by someone else, maybe a caregiver from the professional organization, or by someone close to the patient. The focus of PCC is to empower the patient to be able to perform self-care. Robinson et al. (2008) study PCC from an adherence perspective, defining the concept as the patient’s efforts to follow health care advice. Adherence is related to the concept of compliance, defined in Robinson et al. (2008) as a when a person’s behavior coincides with clinical advice. To use the term adherence reflects more of PCC, i.e. individual care and patient involvement. Self-management is one approach

to creating patient-centered care. Coleman and Newton (2005) discuss self-management in the context of chronic care, including a number of aspects: 1) Patient education and application in real time situations. 2) Providing the patient with problem-solving skills. 3) Helping the patient to act with efficacy, using the most appropriate alternatives. 4) Training the patient in recognizing barriers to action and how to remove these. 5) Performing motivational interviews to support patients to commit themselves to the care. However, self-management does not work on its own; it requires that the other aspects of the care process are in line. The patient must be able to understand the sickness with its social as well as physical consequences, do measurements, obtain help with goal-setting and perform follow-up and evaluation. The care process must be integrated with a focus on the patient's ability to enact the care. Instructions like drug warning labels form one possible type of support for self-care, but making instructions into effective tools for learning and acting is not simple. Webb et al. (2008) give an example of a warning text for prescription drugs. The findings from a discussion group show that a majority of the participants found that the texts contained difficult language or were confusing. The participants in the study requested more actionable texts in the most simple and concise manner possible. Webb et al. (2008) argue that the level of misunderstandings among patients with regard to warning texts is high, and that such texts should be developed towards adopting a patient perspective.

3.6 Keeping Records: Documentation and Evaluation

Evaluation is needed to ensure that PPC approaches promote better health for patients. Building the databases needed as sources of facts is closely related to record-keeping. Both evaluation and records are ongoing activities that are performed in each phase of the care process, as discussed throughout this section.

Evaluations are necessary for making a good case for PCC. For example, Stewart et al. (2000) made a study of the differences in outcome between cases using and those not using a patient-centered approach. The outcome of this study showed that patient-centered communication influences patients' health because they perceived that their visit was patient-centered. Stewart et al. (2000) therefore drew the conclusion that patient-centered practice improved health status and increased the efficiency of care by reducing diagnostic rests and referrals. For the patient to find common ground with the physician was especially important. Cassivi et al. (2008) discuss the problems with measuring by giving an example of a framework in the thoracic surgical area. The main problem is identified as what is called a void in quality measures. The problem is one of a duality in the use of measures. The measures are both used for quality improvement or for economic aspects of care and, together with these, as means for supporting the right care choices. This is also reflected in the type of measures, whether they are outcome-oriented or process-oriented. To fill this void in measures, frameworks for measuring and measures should be developed on the basis of a patient-centered view. Cassivi et al. (2008), also point toward the need of building databases of records from the measures and the importance of these records for future quality improvements. The financial outcome is

also connected with the evaluation and measurement of PCC and its importance for care institutions. Charmel and Frampton (2008) point to number of economic reasons for PCC, including reduced length of stay, a lower cost per case, fewer adverse events, higher employee retention rates, a decrease in malpractice claims and increased market shares.

3.7 Interactions: The Patient-Doctor Meeting

In the communication processes the exchange between patient and doctor is the key piece, which has received a fair amount of attention from several perspectives on the PCC area. This activity is not entirely comparable to the other six above. The patient-doctor contact is a meeting that in practice is the time and place where many of the other activities take place. However, what we here focus on is the nature and condition of the interactions between patient and doctor as such. The nature of this interaction is a key to patient learning in the sense that it is in the presence and interaction with another person (i.e. the doctor) that a fact becomes trustworthy and real to the patient. In this sense it is a central aspect of the process of patient learning. Culture-competent communication has been discussed as one aspect of the patient-doctor meeting. One example of a frame for these processes is the cultural awareness model, a model suggested by Teal and Street (2009) consisting of four elements: communication repertoire, situational awareness, adaptability, and core cultural issues. These four elements indicate a set of skills that the caregiver should possess, and which could form an area for learning and KM activities. It would probably be even more to the point to have support systems for the caregivers in the current situation and as support for learning over time. Another example comes from cancer care, where Epstein and Street (2007) present a model for patient-centered communication. This model consists of six interconnected functions: fostering healing relationships, exchanging information, responding to emotions, making decisions, managing uncertainty, and enabling patient self-management. The book argues for the connection between good communication and the outcome of the care given.

4 Conclusion: PCC as Learning Process

The aim of this paper has been to investigate how to understand and model a health care process as a patient learning process for the use as a IS/IT planning approach. Starting with a tentative model of PCC, we have investigated current practices of PCC as reported in the literature, and found the model a useful way of understanding the PCC process. All the seven types of activity are in use today, albeit not in a coherent or systemic way. The focus of most of the activities is not on the learning aspect as such, however important part it is. Most of the practices can be seen to deal with more than one type, but they usually have a focus on one of them. There are clear needs of development in all these learning processes, and the way they interact over time, all directed towards an integrated and systemic understanding of this problem area. For the care organization learning process there has been progress in directions like care

improvement research, but more generally it would be expected that lessons could be learnt from organizational learning areas. When it comes to community learning, which takes place on a person-to-person level theories within the social learning area could be expected to contribute. The picture of the patient learning process seems to be a very information-oriented and rational one. This is the picture that emerges both in the practice as we have studied it and in the different situations that we have reviewed in this paper. This sits well in the planning perspective and a rational world view of information systems in general. However, this may not be the best way of understanding how a patient really experiences the process.

Acknowledgement. The project Bridging the Gap 2, Futurum - the academy for healthcare, Jönköping county council, Sweden.

References

1. Aita, V., McIlvain, H., Backer, E., McVea, K., Crabtree, B.: Patient-centered care and communication in primary care practice: what is involved? *Patient Education and Counseling* 58, 296–304 (2005)
2. Cassivi, S.D., Allen, M.S., Vanderwaerdt, G.D., Ewoldt, L.L., Cordes, M.E., Wigle, D.A., Nichols, F.C., Pairolo, P.C., Deschamps, C.: Patient-centered quality indicators for pulmonary resection. *Ann. Thorac Surg.* 86, 927–933 (2008)
3. Charnel, P.A., Frampton, S.B.: Building the business case for patient-centered care. *Health Care Financial Management* (March 2008)
4. Coleman, M.T., Newton, K.S.: Supporting self-management in patients with chronic illness. *American Family Physician* 72(8), 1503–1510 (2005)
5. Davidson, P., Cockburn, J., Daly, J., Sanson Fisher, R.: Patient-centered needs assessment – Rationale for a psychometric measure for assessing needs in heart failure. *Journal of Cardiovascular Nursing* 19(3), 164–171 (2004)
6. Dubé, L.: What's missing from patient-centered care? *Marketing Health Services* 23(1), 30–36 (2003)
7. Epstein, R.M., Street Jr., R.L.: Patient-centered communication in cancer care: promoting healing and reducing suffering. National Cancer Institute, NIH Publication No. 07-6225, Bethesda, MD (2007)
8. Forbat, L., Cayless, S., Knighting, K., Cornwell, J., Kearney, N.: Engaging patients in health care: An empirical study of the role of engagement on attitudes and action. *Patient Education and Counseling* 74, 84–90 (2009)
9. Gustafsson, D.H., Hawkins, R., Boberg, E., Pingree, S., Serlin, R.E., Graziano, F., Chan, C.L.: Impact of a patient-centered, computer-based health information/support system - reliability and validity in a patient population. *American Journal of Preventive Medicine* (16,1), 1–9 (1999)
10. Heisig, P., Iske, P.: European knowledge management framework. In: *European Guide to Good Practice in Knowledge Management* (Work in progress, Draft Version 3.0). CEN/ISSS, Berlin, Amsterdam (2003)
11. Laine, C., Davidoff, F.: Patient-centered medicine. A professional evolution. *JAMA* 275(2), 152–156 (1996)

12. Langford, A.T., Sawyer, D.R., Giomio, S., Brownson, C.A., O'Toole, M.L.: Patient-centered goal setting as a tool to improve diabetes self-management. *The Diabetes Educator* 33(suppl. 6), 139s–144s (2007)
13. Robinson, J.H., Callister, L.C., Berry, J.A., Dering, K.A.: Patient-centered care and adherence: Definitions and applications to improve outcomes. *Journal of the American Academy of Nurse Practitioners* 20, 600–607 (2008)
14. Sherer, J.L., Anderson, H.J., Lumsdon, K.: Putting patients first. Hospitals work to define patient-centered care. *Hospital* 67(3), 18–23 (1993)
15. Stephens, R.J., Hopwood, P., Girling, D.J., Machin, D.: Randomized trials of quality of life endpoints: are doctors ratings of patient symptoms interchangeable with patients self-ratings?! *Qual. Life Res.* 6, 225–236 (1997)
16. Steward, M., Brown, J.B., Weston, W.W., McWhinney, I.R., McWilliam, C.L., Freeman, T.R.: *Patient-Centered medicine: transforming the clinical method*, p. 25. Sage, Thousand Oaks (1995)
17. Stewart, M., Brown, J.B., Donner, A., McWhinney, I.R., Oates, J., Weston, W.W., Jordan, J.: The Impact of Patient-Centered Care on Outcomes. *Journal of Family Practice* 49(9), 796–804 (2000)
18. Teal, C.R., Street, R.L.: Critical elements of culturally competent communication in the medical encounter: a review and a model. *Social Science and Medicine* 68, 533–543 (2009)
19. Walker, T.: Shared decision making gains recognition as patient-centric care model. *Managed Health Care Executive* (April 2008)
20. Webb, J., Davis, T.C., Bernadella, P., Clayman, M.L., Parker, R.M., Adler, D., Wolf, M.S.: Patient-centered approach for improving, prescription drug warning labels. *Patient Education and Counseling* 72, 443–449 (2008)

Wireless Computer Games and Applications in the Medical Education Curriculum: Adventures in Pedagogy

Juanita Fernando and Nicolette Peters

Medicine, Nursing & Health Sciences, Monash University
Wellington Road, Monash, Victoria 3800, Australia
{juanita.fernando,nicolette.peters}@monash.edu

Abstract. This manuscript reports work on the first of several related projects in a series entitled “Adventures in Pedagogy”. Serious Computer Games and Applications on wireless devices connected to the Internet are commonly embedded into everyday clinical practice. Yet a review of the literature indicates a scarcity of such curricula in undergraduate medical education. Graduates are evidently not well prepared for wireless e-health practice during University study. Consequently, we have introduced a selective for first year medical students called “Computer Games and Applications for Health and Wellbeing”. The selective is designed to support the clinical application of the wireless tools in an ethical and practical manner while embedding fundamental IT concepts to help prepare graduates for new practice horizons.

Keywords: medical education, health informatics education, professional education, wireless technology, telemedicine, telehealth.

1 Introduction

Most clinicians agree that mobile access to Serious Computer Games and Applications (SCG&A) using wireless devices over the Internet has opened new practice horizons for health. The SCG&A are designed to facilitate improved health and wellness outcomes, epidemic intelligence and public health event detection [1]. Many SCG&A tools are already widely used in Australian, if not international, clinical practice.

Designers of medical education curricula have largely overlooked the SCG&A health practice milieu [2]. Anecdotal evidence suggests the enormous amount of information students must learn over their five year undergraduate qualification militates against explicitly introducing informatics into the medical curriculum [3]. Educators would be well advised to develop a dialectic curriculum that resolves the current mismatch between ubiquitous SCG&A tools for health and busy medical education models. We have begun to address the mismatch by introducing a Selectives unit, “Computer Games and Applications for Health and Wellbeing”, into the first year undergraduate medical education.

2 Study Design

We rely on a qualitative grounded theory study design to facilitate a research approach encouraging reflection on the continuous interplay between data collection and analysis.

2.1 Aim

The aim of this work, the first of several related projects in a series entitled “Adventures in Pedagogy”, reflects on the process of embedding health informatics into MBBS curriculum to support graduates for future wireless practice.

2.2 Methods

We use an action research method. Action research is well suited to exploring and sustaining change processes to established curriculum.

2.3 The Site and Milieu

The Selectives program at Monash University (Victoria, Australia) offers first year Bachelor of Medicine Bachelor of Surgery (MBBS) students’ non-clinical units to study during their second semester. The program provides students with opportunities to develop or acquire new or existing skills outside traditional areas of medical education. The Selectives program ran for eleven weeks for two hours between July and October 2011.

2.4 Participants

We are the only participants at this stage of the research program, which has received University human ethics authorization.

2.5 Evaluation

Students will evaluate the unit in October 2011 using a tool we have developed and the Student Evaluation of Teaching and Units tool administered by Monash University (4). Both evaluations will be triangulated with data collected from our reflections. These data will be analysed interpretively so we may analyze the full complexity of Selective outcomes.

3 Learning Objectives

Unit learning objectives were informed by several meetings with MBBS candidates of varying experience over many years, daily interactions with colleagues, familiarity

with the relevant educational and health informatics literature and professional expertise [5].

3.1 Course Content

The course content was designed so that the medical students might “dip their toes” into mobile e-health tools for education, epidemic intelligence and everyday practice. It incorporates technology skills assessment and training as required, followed by classes supporting students to devise a suitable evaluation tool. The tool will guide the students’ reviews and analyses of serious computer games, telehealth/telemedicine portals, smart phone and tablet applications, social media and 3D applications. The cohort will be organized into groups divided according to free applications using tablets or smartphones connected to the Internet from the classroom. At the Selective conclusion, each group will present their reviews. We hope to publish the student work collaboratively in the future.

4 Support and Logistics

Multidisciplinary University support for the unit included access to devices and other resources drawn from Faculty colleagues and those from IT Support and E-Learning divisions. Expertise from the entire Monash community underpinned much of the syllabus design.

The multinational Australasian College of Health Informatics (ACHI) membership made several suggestions for meaningful syllabus inclusions too. ACHI members often used their own professional networks to support the selective [6]. The authors drew on a generous range of international expertise to inform the classes on topics such as the application of telehealth and telemedicine (public health) to African and transient communities, 3D worlds hosted locally, in New Zealand and the United Kingdom. This pedagogical adventure would not have been able to proceed without such direct and collegiate support.

Finally, the Selectives’ focus on personal and professional development provided a ready-made vehicle for us to offer a new topic, simplifying logistics management. The unit did not require integration into an already busy curriculum [3]. The challenge of embedding health and medical informatics syllabus into medical education more generally is among the subjects of our current research interests.

4.1 Marketing the Selective to Medical Students

The published literature and anecdotal evidence based on research experience from the biomedical and health informatics arena informed the Selectives title. While many clinical students may be technologically savvy this does not necessarily translate to comfort using wireless devices for health in real life [2]. Clinical students often believe new informatics pathways in medical education are too hard to learn as part of their degree program (7). Thus the friendly, if a little misleading, term “computer

games”, a Selective component, is embedded in the unit title. To ameliorate this concern the entire syllabus overview was published to students on the University’s internal e-learning system.

4.2 Syllabus Overview

Key learning goals for each class were framed to ensure Selective learning objectives were largely achieved. However we hoped detailed content would largely be driven by student enrolees. Delivering the Selective was likely to call on the complete depth and breadth of our informatics and research experience.

5 Conclusion

The introduction of medical education curriculum methodically examining mobile access to innovative wireless SCG&A tools on the Internet seems a logical response to new practice horizons for health. We are gratified to be part of an international vanguard of medical informatics educators involved in enriching health professional education.

References

1. Health Informatics Society of Australia: Providing Leadership in E-Health (2011), <http://www.hisa.org.au/>
2. Otto, A., Kushniruk, A.: Incorporation of Medical Informatics and Information Technology as Core Components of Undergraduate Medical Education – Time for Change! *Stud. Health Technol. Inform.* 143, 62–67 (2009)
3. Confederation of Postgraduate Medical Councils (CPMC): Australian Curriculum Framework for Junior Doctors (2011), <http://curriculum.cpmec.org.au/>
4. Monash University. Student Evaluation of Teaching and Units (2011), <http://opq.monash.edu.au/us/surveys/setu/index.html>
5. Fernando, J.: Personal web page (2011), <http://users.monash.edu.au/~juanitaf/>
6. Australasian College of Health Informatics (ACHI) (2011), <http://www.achi.org.au>
7. Berglunda, M., Nilssona, C., Révaya, P., Petersson, G., Nilsson, G.: Nurses’ and Nurse Students’ Demands of Functions and Usability in a PDA. *Int. J. Med. Inform.* 76(7), 530–537 (2007)

The Process of Policy Authoring of Patient-Controlled Privacy Preferences^{*}

Thomas Trojer¹, Basel Katt¹, Thomas Schabetsberger²,
Richard Mair², and Ruth Breu¹

¹ Research Group Quality Engineering, University of Innsbruck, Austria
{thomas.trojer,basel.katt,ruth.breu}@uibk.ac.at

² ITH-icoserve GmbH, Innsbruck, Austria
{thomas.schabetsberger,richard.mair}@ith-icoserve.com

Abstract. Discussions about appropriate security controls to protect medical records led to the understanding that the patient her-/himself plays a crucial role in networked electronic health-care. Patients have individual privacy concerns and may want to execute their personal right of self-determination on access and usage of their medical records. The ability for patients to have control over their personal medical data is the essence of patient-centric networked electronic health-care, but poses challenges regarding its tool support. Since patients can be generally treated as non-security experts as well as non-health-care domain experts, usability-supporting factors of authoring tools for privacy preferences have to receive major attention by implementers. Additionally, domain characteristics have to influence the design of such authoring applications. Finally expressed privacy preferences have to be analysed to inform the patient-author and guide her/him in the policy authoring process. In this paper we discuss the process of authorization policy authoring for shared electronic health records which we use to implement patient-controlled access control authoring tools. Further a use-case in the context of a specific health-care infrastructure is presented.

Keywords: Privacy, Patient privacy policy, Access control, Authoring tools, Information self-determination, Integrating the Healthcare Enterprises (IHE).

1 Introduction

Discussions (e.g., in [9]) about appropriate security controls to protect personal medical records led to the understanding that the patient her-/himself plays a crucial role in networked electronic health-care. This is due to the fact that a patient constitutes the identified individual within a health-record and therefore processing of such medical data is bound to common data privacy regulations of

^{*} This work was partially supported by the Austrian Federal Ministry of Economy as part of the Laura-Bassi – Living Models for Open Systems – project FFG 822740/QE LaB, see <http://lab.q-e.at/>

a country (see e.g., [1]). Further patients have individual privacy concerns and may want to execute their personal right of self-determination on access and usage of their medical records. Legal regulations on data privacy and therein especially information self-determination represent the underpinning motive to let patients express personal privacy concerns. To meet these regulations, customized applications and IT-infrastructure have to be built to make electronic health-records controllable.

Therein we see a major requirement being usable authoring tools supporting and guiding patient-authors of privacy preferences during the authoring process. An important usability-supporting factor during this authoring process is the analysis of privacy policies. Analysis results are used to inform the patient-author about quality and effects her/his privacy settings imply on the overall functionality of the health-care information system. Besides policy analysis, the integration of health-care domain characteristics and domain knowledge also support usability. Such integrations try to answer questions like, *"who are the typical stakeholders within the health-care domain?"*, *"which data is involved and how sensitive is it?"* or *"what data is required by those stakeholders?"*. In this paper we present the process of authorization policy authoring for shared electronic health records and discuss usability-supporting factors involved. The importance of developing highly usable authoring tools comes from the fact that patients are not considered security experts and are not necessarily familiar with working processes of the health-care domain. Therefore patients have to be supported when trying to express their individual conception of privacy towards corresponding enforceable privacy policies.

2 Problem Statement

Privacy can be ensured when a consent or agreement on a purpose of use is stated by the identified individual which gets enforced by the security infrastructure. Further access restrictions limit usage of data in order to prevent potential damage and misuse. Here we want to specifically emphasize on the necessity that the explicit source for setting privacy preferences is the corresponding patient.

A problem gets visible when changing the perspective, asking how privacy policies can be declared in a way so that they match patients' individual conception of privacy. Since it is not feasible to put in place a trusted party who manages policies for each patient, a patient by her-/himself should be allowed to act as the author of privacy settings. An initial requirement to successfully empower patients to do so, is to consider usability-supporting factors of the policy authoring tools. These factors leads to a change in the traditional authoring process, which allows only security experts to define security artifacts. Furthermore, as privacy policy authoring requires health-care domain information, aspects of integration to an established health-care infrastructure have to be covered at the same time.

3 Related Work

Our ongoing work related to patient-controlled access control is based on proposals published by ELGA¹, which is the working group driving the Austrian e-health initiative [2].

There has been general work published in the field of usable security, e.g., [7]. The authors in this work state that when employing usable applications, guiding the user in a privacy policy authoring process will lower the risk of inappropriate use of personal information. In their work they conducted an empirical study to evaluate the use of tools guiding and not guiding users through the authoring process. Significant advantages of the employment of guided tools are shown in their study. This also justifies our effort on implementing patient-controlled access control policy authoring. Still, our work differs by the use-case within the health-care domain together with the domain-aware analysis of patient privacy policies to support the user. Further the authors in [13] evaluated the SPARKLE policy workbench, an enterprise privacy policy authoring application in order to gain information on usability challenges. We are able to develop our usability requirements based on parts of their work, although related to characteristics of a networked health-care landscape.

A core part of the process of policy authoring we propose in this work is authorization policy analysis. An analysis component therefore analyses patient privacy settings and provides feedback to the patient-author. Policy analysis, similar to what we implemented is covered in [10]. Still, in our work we dynamically retrieve health-care domain characteristics, required to enable domain-aware analysis.

Katt et al. [8] propose an architecture for enforcing access control in *Integrating the Healthcare Enterprises* (IHE) based systems. IHE is also the basis for our work regarding the retrieval of domain characteristics. Their work can be used to implement the actual enforcement of patient privacy policies.

4 Authoring of Privacy Preferences

In this section we define the authoring process and discuss usability-supporting factors of authoring of privacy preferences performed by non-security experts. In order to allow those privacy preferences to be machine-interpretable and enforceable by a security infrastructure we translate them to access control policies.

Fig. 1 describes the activities and artifacts involved in our policy authoring process. Based on the type or expertise of a user, a set of templates for declaring privacy settings is gathered and made accessible to the user. After setting privacy options based on domain information a policy analysis component decides whether the privacy policy is enforceable or needs to undergo further editing. From this process we extract factors which play a part in supporting the usability of policy authoring tools. These usability-supporting factors are shown in

¹ electronic health-record (German, "Elektronische Gesundheitsakte")

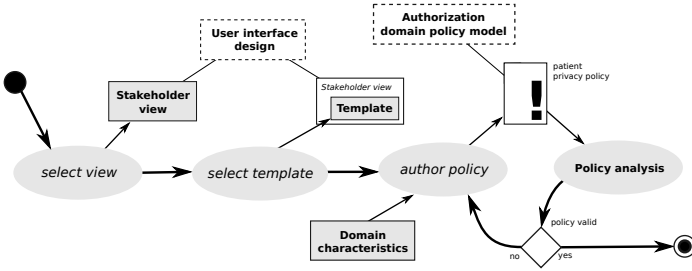


Fig. 1. Authorization policy authoring process represented as activity diagram

Fig. 2 as part of our authorization policy authoring framework. A central building block of this framework is the *authorization policy authoring model* which provides entities required to be presented by *policy authoring user interfaces* as well as to be translated into *patient privacy policies*. Entities of this model correspond to the health-care domain and define authorization aspects. These authorization aspects are introduced to the model via extending the *authorization domain policy model* (cf. our previous work [14]). The authorization domain policy model defines e.g., an *access target* describing the requesting *subject* and an *action* to be executed on an information system *resource*. Further *permission* or *restriction* entities cover those access targets and further limit access by additionally providing *conditions* or *obligations* to be fulfilled. Based on these definitions enforceable access control policies can be generated. Fig. 2 indicates this via the vertical connections between the policy authoring user interfaces and the patient privacy policy through the authorization policy authoring model.

Below we discuss the usability-supporting factors related to authorization policy authoring. *User interface design* and *stakeholder views* are only conceptually mentioned here, since the actual design of authoring tools as well as studies on user groups and detailed use-case scenario analysis, respectively, are out of scope and considered future work.

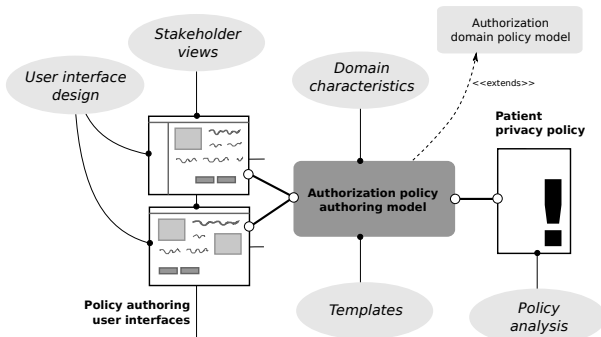


Fig. 2. Usability-supporting factors of access control policy authoring

Stakeholder Views. A stakeholder view defines the authoring functionality which is available for a specific type of user, e.g., a patient or medical professional. By analysing different use-case scenarios, where access to a shared electronic health record of a patient is involved, all required stakeholder views can be identified. A typical use-case scenario regarding a patient is the viewing of her/his own health record or the definition of trusted medical staff gaining extended access rights to health records. On the other hand, medical practitioners want to access the health record of a patient during or after a medical treatment or want to have the option to protect very critical health records from patient access (e.g. if laboratory markers suggest to diagnose a severe disease).

Having different views allows us to define multiple customized applications related to one authorization policy authoring model, expressing the needs of different stakeholders. Further, when developing views according to this model, we can guarantee that the user interface reflects domain and access control aspects appropriately.

User Interface (UI) Design. When defining graphical user interfaces human cognition as well as user behavior during task execution has to be considered. A usable UI is designed in a way e.g., to show interface elements in a well-placed (grouped) manner, describes necessary steps to reach a certain goal, gives a user a history of previous actions taken, lets a user abort (and maybe continue) at any time and keeps a user informed about the application state. In our case the state of the application links back to *policy analysis* as it will be described later in this section.

Domain Characteristics. Policies employed by a security infrastructure may be diverse and complex. A patient's ability to define privacy policies may be influenced by multiple information sources from within the electronic health-care network. Auxiliary information have to be used by the policy authoring application in order to enable a domain-aware policy authoring process. We identify the following types of domain-related information and associated attributes to be integrated:

- *Patient*, i.e. unique patient identifiers, corresponding health records, related medical practitioners, etc.
- *Medical data*, i.e. record identifiers, record types, related stakeholders (e.g., creator or identified patient)
- *Health-care provider*, i.e. working roles, unique identifiers
- *Health-care work processes*, i.e. record type – working role mappings, types and purposes of data processing, patient – practitioner relationships (via medical treatments, referrals or maintained health records), practitioner – resource *needs-to-know* [4] relationships, etc.

Domain characteristics represent the core factor to actually implement usable policy authoring tools. Integration of domain characteristics in the context of an IHE-based health-care infrastructure is discussed in Section 5.

Templates. Templates are partial instances of the authorization domain policy model. A template covers a common concept of a health-care working environment which implies an authorization rule supporting it. Templates are defined via a *label* which associates entities of the authorization domain policy model (used as placeholders) and the actual domain data specified by domain characteristics. The total amount of available templates defines the instantiation possibilities of the underlying authorization domain policy model. Therefore templates have to be based on evaluated use-case scenarios in order to be meaningful and to provide guidance and an overview of functionality to the user. Finally for each template different *stakeholder views* and *user-interface* designs can be considered.

E.g., the template with the label *family practitioner* associates an authorization domain policy stating that the *selected medical practitioner* (i.e. a placeholder for domain data) is *permitted* (i.e. an instance of the *permission* entity) to access *all health records* (i.e. the actual domain data of type health record). As another example, a template labeled *referral* associates a policy which allows a *selected medical practitioner* (i.e. the target of the referral) to access (i.e. a stated *permission*) a basic set of health records (e.g., defined via *record types* or via the (derived) *record sensitivity*) of a *patient*. Such templates, as they correspond to the describing authorization domain policy model can be easily transformed to enforceable access control policies, e.g., expressed by the *Extensible Access Control Markup Language* (XACML) [12].

Policy Analysis. Policy analysis is in general considered a recurring task. By analyzing policy artifacts, feedback can be provided to inexperienced or non-expert users. In our context, a patient expressing privacy preferences can be informed e.g., if certain access rules are in conflict. Further a patient has to be warned if her/his settings would lead to privacy at risk or interfere with working processes in the health-care domain, e.g., by restricting access to medical data where access is required.

By assuming syntactically valid authorization policies, we see two different types of analysis to be performed:

- *Conflict* or *redundancy detection* between policies (see e.g., [11]), i.e. access control is undecidable or one policy dominates another policy, respectively
- *Constraint checking* regarding health-care work processes, i.e. evaluating if e.g., privacy is potentially at risk as no inter-personal relationship between a patient and a practitioner can be derived

Policy analysis is bridging the gap between patients and security and domain experts. On one side patients request privacy protection which conforms to their conception of privacy and on the other side security experts are able to actually express such preferences in a machine-readable way. Further domain experts can contribute knowledge about properties for a functioning health-care domain. A policy analysis component functions herein as a kind of advisory system to a patient. By reporting issues regarding domain characteristics together with the authorization policies themselves, policy analysis positively contributes to the usability of an authoring application.

5 Use-Case: Authoring Tools for IHE IT Infrastructure

In order to establish authorization policy authoring tools in the context of a health-care network infrastructure, we have to identify the required interfaces for retrieving domain characteristics and build upon a specific access control framework. The Austrian ELGA proposes the use of IHE-based systems together with authorization enforcement via an XACML infrastructure. Fig. 3 shows the IHE-profiles which the authoring application integrates, as well as XACML security components it relates to. Detailed descriptions of the security components, their interaction and access control mechanisms are found in [8,12,6] and out of scope of this specific work.

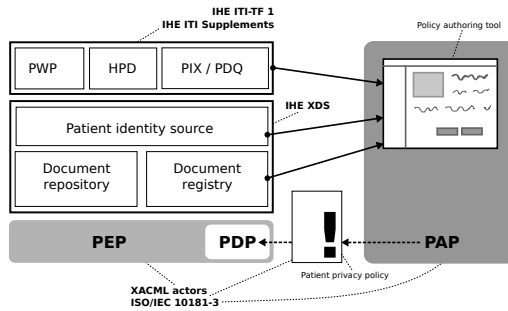


Fig. 3. Integration of policy authoring tools, their relation to the XACML actors and their dependencies to an IHE-based infrastructure

In the following we identify and list different IHE-profiles the authoring application needs to incorporate:

- *Cross-enterprise document sharing* (XDS) [3] *document registry*, in order to support retrieval of *medical data* and *health-care work processes* domain characteristics (cf. Section 4 – *Domain characteristics*)
- XDS *patient identity source*, in order to get local patient information (e.g., from within a specific hospital) as part of the *patient* domain characteristics
- *Patient identifier cross-referencing* (PIX) and *patient data query* (PDQ) [3], to support a unified patient identifier and the retrieval of extensional patient information as part of *patient* domain characteristics, respectively
- *Healthcare provider directory* (HPD) [5] to fetch provider meta-data like work role or credentials, as part of *health-care provider* domain characteristics
- *Personnel white pages* (PWP) [3], to gain extensional provider information as part of *health-care provider* domain characteristics

Employing these profiles within a health-care network maintaining electronic health records enables highly domain-aware policy authoring. The Austrian ELGA and our business partner *ITH-icoserve* correspond to IHE profiles, allowing our work to be put in a practical context.

6 Conclusion and Future Work

In this paper we described usability-supporting factors for developing effective authorization policy authoring tools. The authoring application enables patient-controlled privacy preferences which get transformed to enforceable access control policies. These policies protect access to shared medical data within a networked electronic health-care system. We elucidated factors related to the design of such application which are related to the authorization policy model, domain characteristics influencing the specification of policies and the analysis of policies to provide feedback to the patient-author of policies. Finally we described a use-case of an IHE-based health-care infrastructure and a set of IHE profiles to be integrated within an authorization policy authoring tool. These profiles define the required interfaces to retrieve domain characteristics.

Future work discusses three major work packages: (i) Usability studies guiding the design of user-interfaces and stakeholder views, (ii) a more formal discussion of policy analysis regarding policy conflicts and domain-related constraints and (iii) experiences from our efforts on integrating our authoring prototype to the national health-care network in conformance to all national regulations.

References

1. European Commission. Directive 95/46/EC, Data Protection Directive (1995)
2. IBM Austria. Feasibility Study for implementing the electronic health record (ELGA) in the Austrian health system (2006)
3. IHE. IT Infrastructure (ITI) Technical Framework. Integration Profiles, vol. 1
4. IHE. IT Infrastructure Access Control (White Paper) (2009)
5. IHE. IT Infrastructure (ITI) Technical Framework, Supplement, Healthcare provider directory (HPD) (2010)
6. ISO. ISO/IEC 10181-3:1996 Information technology – Open Systems Interconnection – Security frameworks for open systems: Access control framework (1996)
7. Karat, C., Karat, J., Brodie, C., Feng, J.: Evaluating interfaces for privacy policy rule authoring. In: CHI 2006. ACM (2006)
8. Katt, B., Breu, R., Hafner, M., Schabetsberger, T., Mair, R., Wozak, F.: Privacy and Access Control for IHE-Based Systems. In: Weerasinghe, D. (ed.) eHealth 2008. LNCSIT, vol. 1, pp. 145–153. Springer, Heidelberg (2009)
9. Kotschy, W.: STRING ELGA Datenschutzrechtliche Analyse (German, Electronic health record – Data privacy aspects). Austrian Federal Ministry of Health (2005)
10. LeMay, M., Fatemeh, O., Gunter, C.A.: PolicyMorph: Interactive Policy Transformations for a Logical Attribute-Based Access Control Framework. In: SACMAT 2007. ACM (2007)
11. Moffett, J.D., Sloman, M.S.: Policy conflict analysis in distributed system management (1993)
12. OASIS. eXtensible Access Control Markup Language (XACML) v2.0 (2005)
13. Reeder, R.W., Karat, C.-M., Karat, J., Brodie, C.: Usability Challenges in Security and Privacy Policy-Authoring Interfaces. In: Baranauskas, C., Abascal, J., Barbosa, S.D.J. (eds.) INTERACT 2007. LNCS, vol. 4663, pp. 141–155. Springer, Heidelberg (2007)
14. Trojer, T., Katt, B., Wozak, F., Schabetsberger, T.: An Authoring Framework for Security Policies: A Use-Case within the Healthcare Domain. In: Szomszor, M., Kostkova, P. (eds.) e-Health. LNICT, vol. 69, pp. 1–9. Springer, Heidelberg (2011)

Modelling a User Authorisation and Data Access Framework for Multi-specialty Research Systems in Secondary Health Care

Ire Ogunsina¹, Sarah N. Lim Choi Keung¹, Lei Zhao¹, Gavin Langford³,
Edward Tyler¹, and Theodoros N. Arvanitis²

¹ Department of Primary Care Clinical Sciences

² School of Electronic, Electrical and Computer Engineering

University of Birmingham, Edgbaston, Birmingham B15 2TT, United Kingdom

³ Birmingham and Black Country Comprehensive Local Research Network
{i.ogunsina, s.n.limchoikeung, l.zhao, e.tyler,
t.arvanitis}@bham.ac.uk, Gavin.Langford@uhb.nhs.uk

Abstract. Patient identification and consequent recruitment in clinical trials is normally preceded with searches on electronic health record (EHR) systems. Query results may be collated across multiple health organisations and specialties. In such scenarios, a prime concern is the possibility of systems and their users inadvertently or otherwise impinging on the privacy of patients. Access to patient data is crucial for research purposes, but the degree of access must be controlled in such a way that it conforms to agreed legal, organisational and ethical policies. In this paper, we present a proposed model for managing a dynamic matrix of roles and data access privileges within the context of research systems in secondary care.

Keywords: Role Based Access Control (RBAC), secondary care research, electronic healthcare records.

1 Introduction

Patient data is arguably the most essential resource in health care systems. The recommendations of the Caldicott guardian stipulate that access to patient data must strictly be on a “need-to-know” basis [1]. Caldicott-compliant systems need to be effective in granting and restricting access to patient data and resources according to system and user specifications. This work is part of a larger project involved with the design and development of a research system to be used for clinical studies across multiple specialties in secondary care. System users (subsequently referred to as the subjects) will typically be health practitioners with proficiency in at least one specialty and affiliated to one or more health organisations, typically NHS Trusts.

Apart from the patient data resource, license protected resources also exist for the computation of clinical data. Access to all resources need to be verified along several lines to ascertain rights and permissions. The paper continues with an overview of the

classic RBAC, some known limitations and possible solutions. Sections 3 and 4 focus on the context of our work, the reasoning and rationale behind the proposed model.

2 Overview of the Traditional RBAC

Originally proposed by Sandhu [2], RBAC is a standard of access control whose popularity is proof of its effectiveness. RBAC is based on a grouping mechanism, known as roles. A subject's role, therefore, determines what resources they have access to. This approach breaks the tight coupling between the subject and the permissions to resources. From a data modelling perspective, the inclusion of the role entity resolves the many-to-many relationships between subject and rights or access to resources [3]. Within the health enterprise, the effectiveness and appropriateness of the RBAC approach as a privileges management infrastructure has been successfully demonstrated by Slevin and Macfie [4].

2.1 Problems

Although RBAC has proven its effectiveness and good utilisation record across a wide industrial spectrum, the approach is not to be considered as a panacea for all access control issues [2]. The traditional RBAC has been shown to have limitations, making it unsuitable for scenarios requiring complex access requirements [5]. Some of the main problems, often associated with RBAC, are related to potential policy conflicts and inconsistencies with authorisation of subjects with multiple roles. The objective of our work is to model an effective framework for granting or denying access requests to patient data and resources. This framework must take into consideration the combination of the roles, specialties, Trusts, resources and actions requested. The scope of the proposed model does not include more challenging scenarios, such as emergency access requests.

3 Making Access Decisions

3.1 Policies

Multiple role assignments to a subject can potentially contribute to the problem of separation of duty and role precedence [5]. Our approach lends from the PERMIS authorization infrastructure project [6] where the subject, for example, is an object which could have an attribute-key, role with the attribute-value, researcher-only, or an attribute-key first_name, where the attribute-value may be Fred. The proposed model is also a policy driven model where access criteria are contained in separate policies. *Trust policies* for instance, would exist to specify the different relevant directives that may exist for the subject within the context of their affiliation with the health organisation. Similarly, *specialty policies* exist to provide directives relating to subject specialties that could influence the access decision. By default, each attribute has its policy. Our model contains an access decision framework, which provides a decision based on the aggregation of the different policies.

3.2 Resources

Apart from patient data, another requirement of the model is the management of system resources and tools. An example is licensed quality of life questionnaires which may exist in paper or electronic format. A typical case is the Hospital Anxiety and Depression Scale (HADS) [7] score in the COPD specialty. These licenses are usually bought by the Trust and used in calculating patient data.

The RBAC model of the now defunct NHS National Programme for IT (NPfIT) project [8] includes additional concepts for suitability within the health enterprise. Within each policy entity, directives would specify the concept status such as *Legitimate Relationships (LR)* - ensuring that patient identifiable data is only accessible if the subject is involved in the patient’s care. Non patient identifiable data may be accessed as allowed within the *Sealed Envelopes (SE)* segments of policy documents. In most access request scenarios, it would be impossible to be granted access without *Patient Consent (PC)* status being true within policy documents.

As shown in Figure 1, the scalable model generates a decision outcome as a product of aggregated policies. This approach gleans from Blobel’s more detailed model [9] as well as HL7’s security policy information model [10].




Fig. 1. Policy-based model for resource access decisions

4 Conclusion and Related Work

4.1 Related Work

Farzad and Yu [11] extended Crook’s RBAC model [12] which modelled the concepts of responsibility, operation and context in addition to the role concept as criteria for object access and permission, bears good resemblance with our approach work described with particular focus on including the model during the knowledge

engineering phase. Although the work by Slevin and Macfie [4] involves a single specialty and Trust in a clinical environment, it highlights challenges common to access control mechanisms within healthcare systems and possible solutions.

4.2 Conclusions and Future Work

Traditional RBAC may be extended to include policies and constraints that will augment its suitability in complex scenarios. Our proposed model is a policy-oriented framework which handles access decisions in a flexible way within the context of multi-specialty research systems in secondary health care. It would be beneficial to see how integrating Cassandra, [1] a role based, trust management system into the permissions management infrastructure would facilitate our model. Furthermore, it would be interesting to research on aligning our model fully with the HL7's Privacy, Access and Security Services (PASS) Access Control Services model [10].

Acknowledgements. The work was supported by the England's National Institute for Health Research (NIHR) and the Birmingham and Black Country Comprehensive Local Research Network (BBC CLRN).

References

- [1] Becker, M.Y., Sewell, P.: Cassandra: Flexible Trust Management, Applied to Electronic Health Records. In: Computer Security Foundations Workshop, pp. 139–154 (2004)
- [2] Sandu, R.S., Coyne, E.J., Feinstein, H.L., Youman, C.E.: Role-Based Access Control Models. *IEEE Computer* 29, 38–47 (1996)
- [3] Ferraiolo, D.F., Kuhn, D.R.: Role Based Access Controls. In: 15th National Computer Security Conference, pp. 554–563 (1992)
- [4] Slevin, L.A., Macfie A.: Role Based Access Control for a Medical Database. In: IASTED-Software Engineering and Applications Conference, pp. 19–21 (2007)
- [5] Covington, M.J., Moyer, M.J., Ahamad, M.: Generalized Role-Based Access Control for Securing Future Applications. Technical Report GIT-CC-00-02. Georgia Institute of Technology (2000)
- [6] PERMIS. FAQ (2011), <http://sec.cs.kent.ac.uk/permis/documents/FAQ.shtml>
- [7] Snaith, R.P.: The Hospital Anxiety and Depression Scale. *Health Qual. Life Outcomes* 1, 29 (2003)
- [8] National programme for IT (Npfit), <http://www.gpchoice.org/npfit.aspx>
- [9] Blobel, B.: Authorisation and Access Control for Electronic Health Record Systems. *International Journal of Medical Informatics* 73, 251–257 (2004)
- [10] HL7. Privacy, Access and Security Services (PASS) Access Control Services Conceptual Model. Release 1 (2010), <http://hssp-security.wikispaces.com/PASS+HL7+Balloted+Documents>
- [11] Farzad, F., Yu, E., Hung, P.C.K.: Role Based Access Control Requirements Model with Purpose Extension. In: Workshop on Requirements Engineering, pp. 207–216 (2007)
- [12] Crook, R., Ince, D., Nuseibeh, B.: Modelling Access Policies Using Roles in Requirements Engineering. *Information and Software Technology* 45(14), 979–991 (2003)

Software Engineering-Inspired Approach to Prevention Healthcare

Victor Rentea, Andrei Vasilateanu, Radu Ioanitescu, and Luca Dan Serbanati

Faculty of Engineering Taught in Foreign Languages, Politehnica University of Bucharest,
Bucharest, Romania

victorrentea@gmail.com, andrei.vasilateanu@upb.ro,
ioanitescuradu@yahoo.com, luca@serbanati.com

Abstract. The paper extracts proven principles from software engineering to be applied in the field of prevention and wellness preserving. For instance, the concern-oriented partitioning principle could support the analysis and design of healthcare information systems. Results of such an approach are multifaceted systems that integrate all stakeholders' concerns. For such a system, the paper proposes a multi-agent system, where a patient avatar negotiates with expert systems assisting the user to monitor his health state and supply proven prevention alternatives. The aggregated patient-centered virtual organization also aims at stimulating the patient empowerment through personalized information based on their conditions.

Keywords: patient empowerment, prevention, virtual health assistant, patient-centric virtual organization, multi-agent systems.

1 Introduction

In the past two decades, much of the growth in health expenditures has been attributable to chronic conditions in the context of global population ageing. Future health reform should rely on stimulating patient empowerment [1], i.e., people to take responsibility of their health status and act proactively upon maintaining it [2]. One of the most promising applications of patient empowerment is in the field of prevention, aiming to avoid even the debut of chronic conditions, with an alarming increasing prevalence in young people [3]. However, the existing healthcare information systems are not suitable for this purpose because they focus on acute disease care and favor diagnosis and treatment based mainly on the current symptoms [4]. The paper approaches the healthcare area from the original and seminal perspective of software analysis and design, applying proven methods for tackling partial knowledge problems in the prevention field. Using a concern-oriented approach to health monitoring and prevention, a multi-agent system is designed around the Personal Health Record (PHR) to provide continuous assistance throughout the whole lifespan of the user.

Section 2 presents a brief state of the art in the fields this paper relies upon. To address some of the identified problems, the conceptual design of a multi-agent system is presented in Section 4, building upon a theoretical basis described in Section 3. Conclusions and future work are presented in section 5.

2 State of the Art

The use of agents in healthcare has a wide range of applications, including decision support, medical data management, resource planning and pervasive care [5]. Numerous prototype agent systems were proposed which demonstrate clear strengths compared to classic programming, such as support for expert systems or personalization.

Significant improvements in clinical practice have been demonstrated through the use of decision support systems [6]. Such systems also present new opportunities for patient empowerment through feedback and recommendations.

3 Theoretical Background

The goal of our research is to compare the software development process with the management of the patient's health state in the attempt to apply proven methods from the former to innovate in the prevention field. Although apparently very different in nature, the two domains have many similarities: both produce unique, unrepeatable work, both make critical decisions relying on some abstract model built through iterative analysis and both emphasize the continuous nature of their work, i.e. treatment and maintenance, respectively. Furthermore, both medical assistance and software development requests are triggered by a change in an unpredictable universe, health status and business environment, respectively.

The IEEE defines the concerns for a system as “those interests which pertain to the system's development, its operation or any other aspects that are critical or otherwise important to one or more stakeholders” [7]. Concern-Oriented Software Development [8][9] is an approach to partitioning a software system in a concern-driven way, with the aim of supporting the specification of stakeholders' concerns and their composition into a working system. The concerns about a system are represented as *facets*, usually interrelated and overlapping. A key supposition is that the sum of all the facets identified is the system to be developed, because the composite representation subsumes all the knowledge about the system available at that moment [8]. Thus, the system analyst has to consider all the stakeholders' concerns and design the software system to harmonize all the different facets. Similarly, the medical specialties have different, sometimes overlapping or even conflicting views and concerns about the health state of a patient that should be harmonized in the final medical advice.

In software development several requirements often recur in similar forms, e.g., the requirement for security, but to implement it correctly, worldwide-agreed standards and proven solutions are reused. In medical practice, the proven knowledge is built on clinical trials [10] that quantify the benefic or detrimental effect of habits and lifestyle on the health state in the form of a Relative Risk (RR%), and embedded in clinical guidelines and protocols that indicate the best therapeutic attitude.

After delivery, a software system enters the maintenance phase throughout which its various problems and limitations emerge in the ever-changing business environment. Similarly, in healthcare, while treating a patient multiple adverse effects together with unexpected medical conditions may begin to manifest. However, regular monitoring could predict acute and possibly dangerous manifestations.

4 The Patient-Centric Virtual Organization

The system represents a Patient-Centric Virtual Organization (PCVO) and consists of three parts: the PHR, the *Virtual Health Assistants* (VHA) and the patient *avatar*.

Conceptually, the **PHR** provides views on the patient's current health state and clinical history, including both self-gathered data and authoritative information retrieved from the Electronic Health Record (EHR). **VHAs** are collaborating expert systems whose role in the organization is to predict and prevent health deterioration based on the PHR data. The system is designed using a concern-oriented approach, with each VHA dedicated to a medical specialty. This organization enables distributed development and responsibility of the prevention Knowledge Bases (KB) containing proven advices. In the PCVO, **the avatar** is the leading agent that coordinates the activity and the only entity that directly communicates with the user. One of its main tasks is to capture the user profile and to adapt to it in order to increase its influence on her/his life style. On the other hand, in terms of virtual/real distinction, it is a boundary agent that maps the real into the virtual [4] and provides the required information to the VHAs. It is a design objective for our system to be low-intrusive, essentially remaining a reactive system and using proactive communication channels only for critical health conditions. A typical usage scenario is presented below (Fig. 1):

- The user contacts the avatar which in turn asks the VHAs for information requests.
- Each VHA responds taking into consideration the known state of the respective anatomic system, predispositions and interactions with other VHAs.
- The avatar aggregates and provides the user a personalized view on these requests.
- The user responds based on current symptomatology and/or by executing several quick self-tests (e.g. glucose level) and her/his responses are stored in the PHR.
- To build a clearer image, a VHA can request further patient information.

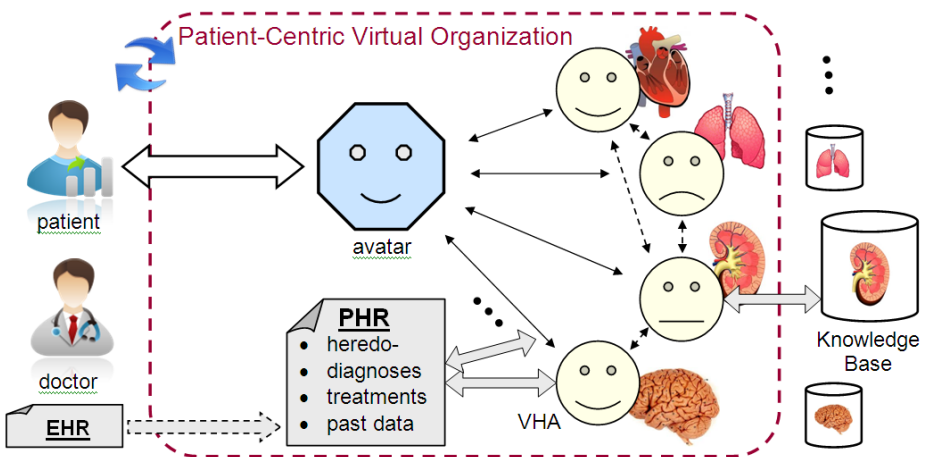


Fig. 1. Patient-Centric Virtual Organization

- The different recommendations are negotiated between VHAs, e.g., potassium salt is good for patients with high blood pressure, but questionable in case of arrhythmia.
- The final recommendations are presented to the user and integrated in the PHR.

5 Conclusions

Software development was compared to healthcare provisioning and proven methods were applied in the field of medical prevention. The paper presents the conceptual design of a concern-oriented multi-agent system for monitoring the health state and stimulating the patient empowerment. Unlike most existing agent systems [5], the control loop does not include any medical staff, thus allowing for a quicker adoption. What differentiates our approach from other medical expert systems [6] is the emphasis on prevention targeted at persons with no previous symptomatology. Incipient stages of an otherwise unsuspected illness are detected and help is provided to prevent its debut. Future work will address the behavior of the avatar and the negotiation among the VHAs that determines the questions and recommendations for the patient.

The user of the system benefits from a continuous and unobtrusive monitoring achieved through simple and personalized questions. From the user's perspective, the system is in continuous evolution following his current health state and providing proper advices. On a social dimension, an important expected impact is on stimulating the patient empowerment and disseminating basic medical information, through pertinent advices and continuous care provided by the system. However, the greatest potential of the system relies in the opportunity to avoid the very development of a chronic condition, one of the today's major health concerns.

References

1. Thorpe, K., Ogden, L., Galactionova, K.: Chronic conditions account for rise in medicare spending from 1987 to 2006. *Health Aff.* (February 2010) doi:10.1377/hlthaff.2009.0474
2. Anderson, R.M., Funnell, M.M.: Patient empowerment: reflections on the challenge of fostering the adoption of a new paradigm. *Pat. Ed. and Counsel.* 57, 153–157 (2005)
3. Schleiffenbauma, B.E., et al.: Unexpected high prevalence of metabolic disorders and chronic disease among young male draftees – the Swiss Army XXI experience. *Swiss Med. Wkly* 136, 175–184 (2006)
4. Serbanati, L.D., Ricci, L.R., Mercurio, G., Vasilateanu, A.: Steps towards a digital health ecosystem. *Journal of Biomedical Informatics* (in press) doi: 10.1016/j.jbi.2011.02.011
5. Isern, D., Sánchez, D., Moreno, A.: Agents applied in health care: A review. *International Journal of Medical Informatics* 75, 145–166 (2010)
6. Kawamoto, K., Houlihan, C.A., Balas, E.A., Lobach, D.F.: Improving clinical practice using clinical decision support systems: a systematic review of trials to identify features critical to success. *BMJ.* 330, 756 (2005)
7. IEEE Recommended Practice for Architectural Description of Software-Intensive Systems. *IEEE Std.* 1471 (2000)

8. Serbanati, L.D.: Integrating Tools for Software Development. Prentice Hall (1992)
9. Bogdan, C., Serbanati L.D.: Concern-oriented and Ontology Based Analysis of Information Systems. In: Proc. ITAIS Conference 2008, in Information Systems: People, Organizations, Institutions, and Technologies, pp. 245–253. Springer, Heidelberg (2010)
10. Timmermans, S., Mauck, A.: The promises and pitfalls of evidence-based medicine. *Health Aff.* 24, 18–28 (2005)
11. Arrow, K.J.: Uncertainty and the welfare economics of medical care. *Am. Econ. Rev.* 53, 941–973 (1963)

Evaluation of a Web-Based Patient Portal for Chronic Disease Management

Stacey Guy, Alexandria Ratzki-Leewing, and Femida Gwadry-Sridhar

Lawson Health Research Institute,
Commissioners Rd E. 801, N6C 5J1 London, Canada
stacey.guy@sjhc.london.on.ca
alexandria.ratzkileewing@lawsonresearch.com
femida.gwadry-sridhar@lhsc.on.ca

Abstract. Chronic disease directly affects more than 9 million Canadians. Efficient strategies are needed to cope with the demand on health care services and to increase patient adherence to treatment. Emerging web 2.0 technologies present viable options for patient engagement in health care. We undertook a pilot project to assess the feasibility of two chronic disease management patient portals. A total of 35 patients participated in the assessment. Portals were evaluated for participant expectations, motivations, usability, and recommendations for future iterations. Findings suggest the features of this portal were useful. Important issues to participants include access to their medical record, communication with health care professionals and other participants regarding topics of interest, keeping track of biometrics, and keeping up with the latest clinical studies.

Keywords: Diabetes, Prostate Cancer, Web 2.0, Chronic Disease, Health Portal.

1 Introduction

Currently, more than 9 million people are suffering from chronic disease in Canada [1]. Accounting for nearly 87% of all disability in the country and consuming over 67% of all healthcare costs, chronic disease poses an incredible burden on the Canadian healthcare system. By 2015, the World Health Organization predicts that chronic disease will account for 89% of all Canadian deaths [2].

In Ontario, Canada, the impact is just as severe. Approximately 1 in 3 people suffer from at least one chronic disease in the province, costing the healthcare system a total of 80 billion dollars annually [3]. The impact of adverse effects of chronic disease are especially salient in Southwestern Ontario where rates of chronic disease, particularly prostate cancer and type II diabetes, are disproportionately higher than in other regions of the province [4]. A 2011 report by the Canadian Cancer Society found prostate cancer to be the most frequently diagnosed cancer in Ontario followed by breast cancer and colorectal cancer [5].

Canada, like most developed countries have a cohort of aging baby-boomers. In Southwestern Ontario the prevalence of diabetes and cancers continues to grow at an

alarming rate. However notwithstanding current strategies to mitigate the prevalence and severity of the diseases, Ontario's healthcare system is underperforming with regards to chronic disease management and treatment [3]. Increasing chronic disease patient participatory care needs requires innovative use of communication technology, including interactive, tailored programs with feedback and social support through networks, not to mention access to care.

1.1 Study Objective

In conjunction with a web portal provider - we developed, implemented and evaluated a web-based chronic disease management system for patients suffering from diabetes (My Diabetes Wellness Portal™) and prostate cancer (ProPortal™), respectively (MedManager Interactive Corp., Waterloo, ON). The goal of this pilot project is to test the effectiveness and usability of an interactive web-based patient portal in providing prostate cancer, and diabetes patients with the knowledge guidance and education they need to help them understand their disease.



Fig. 1. My Diabetes Wellness Portal™ interface. Features of this portal can be seen in the side bar located on the left-hand side.

2 Methods

We designed and implemented a pilot study where patients from each disease cohort were allocated to one group per portal and followed for 6 months. All patients were given access to usual care (print material, advice from their physician and other health providers) and new sources of education, via the web-based portal. Patients had access to self-monitoring tools, and the ability to track disease-related metrics. Our ethics committee (University of Western Ontario) found this approach acceptable (REB #16100E).

Two disease cohorts were targeted through convenience sampling. Patients diagnosed with prostate cancer, including first and second line, of any age, were approached to participate. Patients with type II diabetes were invited to participate.

This cohort was chosen as the disease requires many variables to be managed to achieve optimal care. In both populations patients needed to have access to a computer and the internet. Patients with comorbidities were not excluded. This allowed us to have a fully representative population. We were not concerned with the age range between cohorts (prostate cancer patients are generally an older cohort) as we were not comparing cohorts. Rather this range allowed us to examine possible barriers to use of web 2.0.

Patients were recruited by the physicians and educators at the London Health Sciences Centre. One clinic was located in the Cancer Program, the other in Endocrinology Care. These clinics are ideal as a large number of patients from practices in South-western Ontario are referred for treatment. In addition to face-to-face recruitment, mail outs were sent to patients. We aimed to recruit a total of 50 patients for each condition.

The patient letter of information and consent clearly explained the nature of what data was going to be stored and how this was going to be stored. The data stored is encrypted, and the portal company did not have access to any data. Data is stored behind a secure hospital firewall and backed up nightly. The principal investigator, physician, research assistant and healthcare provider on the team (a social worker, and a pharmacist) had access to the portal. Patients were given randomly generated logins which they changed in order to ensure confidentiality. Disease specific data for the site is provided through manual entry (information entered by the patient) and HL7 data transfer.

An interactive, guided help video is available on the portal. Upon signing in, the help video would pop up. A toll-free helpline was established for participants experiencing difficulties. This was manned by the portal provider. Participants were also able to call the research assistant if an issue arose.

We initiated the study in September 2009. Each portal was monitored by a healthcare provider. The portal enables participants to track their disease-related metrics (e.g., diabetes patients could download readings from their blood glucose monitor) and visualize the data via graphs. External notifications based on goals set by patients are automatically sent. Evidence-based educational material, chosen by healthcare providers on the team, was accessible through the portal. Patient-to-patient and patient-to-provider interaction was available through a community forum and short messaging service.

The feasibility of this portal is assessed by a telephonic survey and a focus group. A 10 minute telephonic survey was administered to patients from both portals. This survey consists of 28 likert scale items and 4 open-ended questions. It was created to assess patient experience with the portals over the course of 3 months. Questions were derived from a questionnaire developed by Evangelista et al., (2006) [6] as well as an expert panel of software developers, and healthcare providers (specifically, team members). Items evaluated: motivation to use the portal, expectations, usability, aesthetics, specific features, support team service, and benefit to health. Responses to the survey were analyzed according to frequency and were subsequently grouped into the themes.

Focus groups were held, at study closure, with participants from each portal, to provide in-depth understanding of portal experience. A focus group is a qualitative research data collection method. Focus groups are particularly useful for exploring opinions, preferences and experiences of a study [7, 8]. Focus groups "have an advantage for researchers in the field of health and medicine...they can encourage participation from people reluctant to be interviewed on their own or who feel they have nothing to say" p.299[9]. Focus group validity is recognized by considering the participants' responses as "an accurate representation of the perceptions of reality for the group members and therefore valid" p.489 [10]. According to Calder [7] enough focus group sessions have been held when it is possible for the moderator to anticipate what will be said next. A semi-structured interview schedule was developed, aimed at developing iterations of the portal that would provide optimal support for managing a disease. The themes explored within the focus groups included: motivations and expectations of the project, usability of the portal, reasons for usage or non-usage of the community forum, and suggestions for the next portal prototype. Two facilitators were present at each focus group (SG, and FG-S). These sessions were audio recorded.

3 Analysis

Data from the telephonic questionnaire was entered into SPSS (Chicago, IL) where the frequency of items was computed. Open-ended questions were grouped into themes initially laid out by the areas delineated in the survey. In addition, usage statistics were collected through the backend of the portal. This provided data on number of logins, time spent on each task etc.

The two focus group sessions (one for each cohort) were audio recorded and the raw data was transcribed by a moderator (S.G). This data was analyzed according to thematic analysis. Qualitative thematic analysis provides a rigorous method of analysis across which a gathered data set will be searched and organized in to pre-empted and emergent themes (repeated units of meaning or patterns) [11; 12] This analysis requires that initial codes be generated after transcription, searching for patterns across the data set, reviewing the themes, defining and naming themes, as well as reporting issues considered relevant to the research question.

4 Results

Thirty-four out of 64 participants completed the survey (at the time of the survey, 64 participants were enrolled in the study). Ten male and five female My Diabetes Wellness Portal™ (MyDWP) participants completed the survey – the majority of which (5 participants) were between the ages of 40 to 50 years. The majority of ProPortal™ (ProP) participants (n=19) who completed the survey were between the ages of 73-83 years (8 participants). At study closure, when the focus groups were conducted, a total of 99 participants (46 MyDWP participants, and 53 ProP participants) had consented to the study. Five ProP, and 2 MyDWP participants took

part in the focus groups. Findings from both the survey and focus groups are reported within themes in the table below.

Table 1. Survey and focus groups findings across both cohorts

Themes	Findings
Motivations	Expand knowledge base and receive Canadian content Sense of community and social network. Help others. Find out how to improve overall health.
Expectations	Involvement of primary physician. Complete medical record history pertaining to disease to be available via portal. System to be available to hospital personnel.
Usability	Easy to navigate, well-organized, clear, caught-on quickly. Feelings of frustration at missing medical information led to discontinued use.
Portal Usage	Time spent on MyDWP: average 15.60 minutes with average participant login of 1.64 per day. Time spent on ProP: 17.58 minutes with average participant login of 0.66 times per day. Frequency of login related to checking for new information and postings.
Privacy & Security	Felt personal information secure and protected.
Features	
Community Forum	Lack of participation by participants and healthcare team. Wanted to see more activity. Recommendations: Discussion led by healthcare provider on topic chosen by participants.
Tracking Tools	Well-liked and used. Need to combine items of likeness.
External Notification	‘No new messages’ led to turning off feature.
Personal Health Record	MyDWP: liked ability to upload glucometer readings. However, some glucometers were not compatible with the system. Delay between results and appearance on portal. Results uploaded were close to unreadable. ProP: Wanted portal pre-populated with biometric data.

Table 1. (continued)

Future Portal Recommendations	
Design	Ability to design guest accounts. Visual identifiers for healthcare team. One main portal with functionality to choose co-morbidity. Receive notification of portal updates and reminders. Relational graphing of biometrics. Glossary of acronyms. Intuitive system.
Pedagogy	Video clips of talks/programs/recent news. “Frequently Asked Questions” Information on new study findings and products.
Medical Records	Clinicians to make recommendations based on info entered by participants.
Economic Costs	Portal of worth and would pay a fee for that. Build into insurance of chronic disease.

5 Discussion

With 1 in 20 diabetes patients in Ontario experiencing major complications within a year, [13] & 1 in 7 Canadian men at risk for prostate cancer [14], it is essential we explore the feasibility of tailored, interactive web-based portals to encourage patient self-monitoring. Regardless of age (majority between 61-83 years old) 58 participants logged in and used the portal. The majority of participants derived benefit from the tools provided, however they wanted a more interactive social medium.

Future forays into web-based self-management programs need to engage participants in evaluation to ensure sustainability. Participant recommendations include emphasis on the social aspects of the portal, as well as integration of pre-populated medical records for tracking. Self-management opportunities will support future actions based on engaging patients in evaluation and improvements to these portals. Their suggestions include, specifically increasing forum activity- for chats and provider-patient interactions, posted clinician recommendations and pre-populated medical records for tracking.

5.1 Study Limitations

Despite the initial sample size of 64 participants (size of the sample at the telephonic interview time period), only 33 patients completed the survey and only 7 participated in the focus group (an effort was made to contact all participants). This concern summons question bias. Did non-compliant users not respond to the survey? Would their answers have provided different information than what was collected? Additionally, questions regarding the motives of focus group participants and what

implications these motivations may have had on study results must also be considered. Repeating the trial with a larger sample size will improve the validity and generalizability of the study; in painting a more accurate picture of the MyDWP, researchers can accurately assess the portal's effectiveness. In the future, it may be beneficial to apply the technology adoption model (TAM) to determine user acceptance of the portal and to make study outcomes more generalizable to the research population.

6 Conclusion

This pilot study marks an important journey into e-based chronic disease management in Canada. As the role patients with chronic disease play in their 'healthfulness' (as opposed to illness) becomes larger, cost-effective avenues to explore self-management become crucial to the survival of our healthcare system. By learning from the findings discussed in this paper, researchers will be able to deploy future iterations of portals that encompass more of what patients want to see. Issues of importance to participants include access to their medical record, communication with health care professionals and other participants regarding topics of interest, keeping track of biometrics, and keeping up with the latest clinical studies.

Acknowledgments. This work was funded by the Ministry of Research and Innovation, The Health Technology Exchange. We would like to thank the following investigators on the project: Dr. Glenn Bauman, Dr. Merrill Edmonds, and MedManager Interactive Corp.

References

1. Canadian Academy of Health Sciences.: Health System Transformation to Meet the Burden of Chronic Disease: The Challenge (2008), http://www.cahs-acss.ca/e/pdfs/CDM_Challenge.pdf
2. World Health Organization.: Preventing Chronic Diseases: A Vital Investment. World Health Organization Press, Geneva (2005)
3. Ontario Health Quality Council.: 2008 Report on Ontario's Health System. Ontario Health Quality Council, Toronto (2008)
4. The University of Western Ontario.: Facing Facts (2011), <http://www.robarts.ca/facing-facts>
5. Canadian Cancer Society.: Ontario cancer Statistics (2011), <http://www.cancer.ca/ontario/about%20cancer/cancer%20statistics/ontario%20cancer%20statistics.aspx>
6. Evangelista, A., Stromberg, A., Westlake, C., Ter Galstanyan, A., Anderson, N., Dracup, K.: Developing A Web-based Education and Counseling Program for Heart Failure Patients. *Prog. Cardiovasc. Nurs.* 21, 196–201 (2006)
7. Asbury, J.E.: Overview of Focus Group Research. *Qual. Health Res.* 5, 414–420 (1995)
8. Barbour, R.S.: Making Sense of Focus Groups. *Med. Educ.* 39, 742–750 (2005)

9. Kitzinger, J.: Qualitative Research – Introducing Focus Groups. *Br. Med. J.* 311, 299–302 (1995)
10. Carey, M.A.: Concerns in the Analysis of Focus Group Data. *Qual. Health Res.* 5, 487–495 (1995)
11. Braun, V., Clarke, V.: Using Thematic Analysis in Psychology. *Qual. Res. Psych.* 743, 77–101
12. Huberman, A.M., Miles, M.B.: *The Qualitative Researcher’s Companion*. Sage, Thousand Oaks (2002)
13. Ontario Health Quality Council.: 2010 Report on Ontario’s Health System. Ontario Health Quality Council, Toronto (2010)
14. Canadian Cancer Society.: Prostate Cancer Statistics at a Glance (2011), http://www.cancer.ca/Canada-wide/About%20cancer/Cancer%20statistics/Stats%20at%20a%20glance/Prostate%20cancer.aspx?sc_lang=en

Cardio Online Reader/COR: A Web 2.0-Based Tool Aimed at Clinical Decision-Making Support in Cardiology

Vendula Papíková and Miroslav Zvolský

Centre of Biomedical Informatics, Institute of Computer Science,
Academy of Sciences of the Czech Republic,
Pod Vodárenskou věží 2, 182 07 Prague, Czech Republic
{papikova, zvolsky}@euromise.cz

Abstract. There is a wide acceptance of the fact, that processes of clinical decision-making has to be based on sound scientific evidence. But in the context of a rapidly growing amount of new information, it is increasingly difficult task. That is why specialized tools and resources enabling to quickly and efficiently search and disseminate relevant evidence are needed. Web 2.0 is a platform that provides an armamentarium with great potential to contribute to solving this task. This paper presents a newly developed information resource aimed at support of clinical decision making in cardiology called Cardio Online Reader/COR. The COR provides the best available evidence from database MEDLINE/PubMed by means of an online application equipped with tools and services typical for Web 2.0. A beta-version of the application is freely accessible at <http://neo.euromise.cz/cor>.

Keywords: EBM, Web 2.0, medical information sources, clinical decision-making support.

1 Introduction

There is a wide acceptance of the fact, that the processes of clinical decision-making has to be based on reliable scientific evidence, nowadays derived from clinical research trials [1], [2], [3], [4]. A general agreement exists also around the statement, that medical doctors have been challenging a huge increase of information, which is still unstopably growing. They do not know about important advances, and feel overwhelmed by new scientific information [5], [6].

1.1 Information Overload

According to the study of Richard Smith medical knowledge will increase four times during a professional lifetime [6]. Nowadays, a total of 20,000 randomized controlled trials (RCTs) are published each year and approximately 50 new RCTs are published every day. Therefore, to keep up to date with RCTs alone, a general practitioner (GP)

would have to read one study report every half hour, day and night. In addition to RCTs, about 1000 papers are also indexed daily on MEDLINE from a total of about 5000 journal articles published each day [7]. In addition to that, the quality of most of the information is very poor: most published information is irrelevant and/or the methods are not good [7]. Medical doctors have to select and appraise the sought information, which requires specific knowledge, skills and experience. But previous studies unveiled, they are not good at finding new information, and do not know how to evaluate it when it is found [5], [6]. Information retrieval therefore is not a trivial task for most of them.

Even when clinicians have time to read some of a new literature, it is difficult to identify which information will be most useful in clinical practice and to recall the most up-to-date findings when they need them [7]. Moreover, push technology to disseminate information has magnified the problem to unwanted information [8]. Unsolicited information received through the mail alone can amount to kilograms per month [7]. And so busy clinicians are now caught in an information paradox - overwhelmed with information but unable to find the knowledge they need when they need it [8].

Considering this reality, it is not surprising, that medical practitioners, particularly GPs, are overloaded with information[7] and - at the same time - they face a serious problem in keeping up to date [5], [6], [7].

1.2 Evidence-Based Clinical Practice

Evidence-based medicine (EBM) movement has been promoting the use of results of clinical research in medical practice. But in the above described context, it is an increasingly difficult task. It has been published a whole range of studies referring about difficulties and obstacles of adequate utilization of the latest and best available evidence in the clinical practice (e.g. [2], [9]).

One of the biggest obstacles is clinicians' lack of time [10]. A questionnaire study of general practitioners was carried out in the former Wessex region, England. Randomly selected sample of 25% of all general practitioners (452), of whom 302 replied. The study revealed that the major perceived barrier to evidence based practice was lack of personal time [11]. Time limitations dictate that evidence based practitioners also rely heavily on conclusions from pre-appraised resources, which apply a methodological filter to original investigations and therefore ensure a minimal standard of validity (e.g. Cochrane Library or ACP Journal Club) [12].

Another clinicians' barrier to evidence-based practice is lack of specific skills[12]. They have difficulty finding, assessing, interpreting, and applying current best evidence [13].

1.3 Clinical Information Sources

Doctors need to be linked to the medical research literature in a way that allows them to routinely obtain up-to-date, outcomes-based information [7]. Most of the questions generated by doctors can be answered, usually from electronic sources, but it is time consuming and expensive to do so-and demands information skills that many doctors do

not have [6]. Though new resources focused on clinical doctors' information needs (e.g. Clinical Evidence or Evidence Updates) have been created, new information tools still are needed [4], [6]. They should respect the "3Rs" of evidence communication, which are reliability, relevance, and readability [4]. They are likely to be electronic, portable, fast, easy to use, connected to both a large valid database of medical knowledge and the patient record, and a servant of patients as well as doctors [6].

Some of the features mentioned above can be accomplished by means of the Internet. Doctors are increasingly proficient with the Internet. Searching with Google came out as useful means to formulate a differential diagnosis in difficult diagnostic cases [14]. Yet doctors are seeking new methods of information discovery because of the limitations of search engines [15]. Also the use of MEDLINE/PubMed to answer daily medical care questions is limited because it is challenging to retrieve a small set of relevant articles and time is restricted [9].

The Web 2.0 proved a potent platform able to provide right tools for the above mentioned tasks. Its applications and services are characterized by features enabling collaboration, information sharing and aggregation, composition of independent services and provision of rich user interaction [16]. Using an RSS doctors can fight information overload. RSS feeds enable to them to organize new web content sent to them by various medical websites in a single interface of an RSS reader [15]. All of the facts described above motivated us to develop an information system aimed at targeted dissemination of the best available evidence from the cardiology by means of tools and services of the Web 2.0.

2 Cardio Online Reader/COR

The COR (Cardio Online Reader) application is based on domain focused records of scientific publications, which are presented using Web 2.0 technologies. The application functions as an online RSS reader and database of selected types of scientific articles. In the process of selecting information we put the accent on their high reliability and clinical relevance according to principles of evidence-based medicine. The articles have been gained from the biomedical database MEDLINE/PubMed. The automated script periodically loads selected records from free accessible interface of the PubMed and stores them in the own fully searchable database of the COR. Afterwards the most recent articles are displayed at top positions of the COR title web page. The user interface of the COR was developed with an accent on ease of use and simplicity of control.

The goal of the COR is to simplify tracking and searching for methodologically valid and clinically relevant publications to disseminate the latest piece of knowledge from the clinical research to the clinical practice, and to create a space for discussion about these findings and articles. These goals are reached by following attributes and functions:

- The online reader presents articles selected from the MEDLINE/PubMed database.
- The articles have been chosen according to criteria of evidence-based medicine, specifically methodological reliability and clinical relevance. For that reason content of the application is created from five types of articles. They are Guidelines

(G), Practice Guidelines (PG), Systematic Reviews, (SR), Systematic Reviews with Meta-Analysis, (SR&MA), and Randomized Controlled Trials (RCT). Separate types of articles are marked by variously colored graphical elements.

- A default page of the application contains a chronological list (in the reverse order, latest articles are on top of the list) of abbreviated records in the following format: article title – article authors – the date of enlisting to the PubMed”. Besides that the number of comments and user rating for each article are also displayed.
- A directly accessible detailed display of the record consists of bibliographical data, generally the abstract of the article, link to the full text (placed at its original location), and link to the original MEDLINE/PubMed record (which can be accessed in a new browser window). Furthermore each article can be rated using a five-point scale within the detailed display of the article record.
- A search form can be filled out by keywords, MeSH terms, author name, and type of an article. Search results can be limited by a time filter including a graphical calendar. The full text search works with Boolean terms. The interface of the COR application provides also a form enabling a direct retrieval of PubMed using the Entrez system.
- The articles are tagged by MeSH terms, which are displayed as a tag cloud showing the relative frequencies of MeSH terms (“MeSH cloud”). A simple MeSH cloud consists from about sixty the most frequent terms within the selected type of articles (or within all articles, if no type of article is selected). The cloud is displayed in the right column of every page. A complete MeSH cloud (or a MeSH list in alphabetical order) can be accessed by provided link. These lists are built from all MeSH terms contained in all articles enlisted in the COR database. An user can simply add each MeSH term to the search form by clicking on it.
- There are implemented web services enabling viral propagation of the whole application, as well as the individual articles. Users can use the Facebook social network or the Twitter microblog as well as social news websites, social bookmarking websites, and many blogging platforms. The application is connected with web services providing social bookmarks including tools for archiving and sharing expert publications as Connotea or CiteULike.
- Users can trace activities related to the COR through Facebook or Twitter, YouTube, blog, and/or. They can also subscribe to website actualizations via integrated RSS feeds.

3 Discussion

The presented web application uses freely accessible source of biomedical bibliographic information and brings the added value of domain specific focus (cardiology). Its development is concentrated on quality, simplicity, and usability.

We plane to develop a fully individualized interface, predefined filters, list of favorite articles, authors or MeSH tags, and other advanced functions and tools for registered users in next versions of the application. There is also a need for at least basic registration process to ensure the chance to archive the authorship of comments

and ratings. A big deal will be to adjust the application to future trends in information sharing, as they develop spontaneously.

4 Summary

We introduce the Cardio Online Reader/COR, which is a medical information resource focused on cardiology. It is intended to support clinical decision-making and medical care based on evidence ("Evidence-Based Health Care"). The application is therefore designed to facilitate both monitoring and searching for scientific evidence to doctors, and to streamline the transfer of the latest findings of clinical research into the clinical practice. The content has been acquired from one of the leading biomedical databases, MEDLINE/PubMed, and has been equipped with tools and services specific for Web 2.0. It also offers space for communication of users about the articles by means of comments or via the website forum. So emerging content generated by a community of users-professionals ("user-generated content") will be an added value to complement the core of a purely scientific content of the database. A beta-version of the COR is available at <http://neo.euromise.cz/cor>.

Acknowledgements. This work was partly supported by the project 1M06014 of the Ministry of Education CR and the research plan AV0Z10300504.

References

1. Eddy, D.M.: Evidence-based medicine: a unified approach. *Health Affairs (Project Hope)* 24(1), 9–17 (2005)
2. Moyer, V.: Evidence based medicine: is it practical? *Arch. Dis. Child.* 89(5), 399–400 (2004)
3. King, J.F.: A short history of evidence-based obstetric care. *Best Practice & Research. Clinical Obstetrics & Gynaecology* 19(1), 3–14 (2005)
4. Straus, S., Haynes, R.B.: Managing evidence-based knowledge: the need for reliable, relevant and readable resources. *CMAJ* 180(9), 942–945 (2009)
5. Williamson, J.W., German, P.S., Weiss, R., Skinner, E.A., Bowes, F.: Health science information management and continuing education of physicians. A survey of U.S. primary care practitioners and their opinion leaders. *Annals of Internal Medicine* 110(2), 151–160 (1989)
6. Smith, R.: What clinical information do doctors need? *BMJ: British Medical Journal* 313(7064), 1062–1068 (1996)
7. Glasziou, P., Del Mar, C., Salisbury, J.: *Evidence-Based Medicine Workbook*, 1st edn. BMJ Publishing Group, London (2003)
8. Gray, J.A.M.: Where's the chief knowledge officer? To manage the most precious resource of all. *BMJ: British Medical Journal* 317(7162), 832–840 (1998)
9. Hoogendam, A., Stalenhoef, A., Robbe, P., Overbeke, A.J.: Analysis of queries sent to PubMed at the point of care: Observation of search behaviour in a medical teaching hospital. *BMC Medical Informatics and Decision Making* 8(1), 42 (2008)
10. Oxman, A.D., Guyatt, G.H.: The science of reviewing research. *Ann. N Y Acad. Sci.* 703, 125–134 (1993)

11. McColl, A., Smith, H., White, P., Field, J.: General practitioners perceptions of the route to evidence based medicine: a questionnaire survey. *BMJ* 316(7128), 361–365 (1998)
12. Guyatt, G.H., Meade, M.O., Jaeschke, R.Z., Cook, D.J., Haynes, R.B.: Practitioners of evidence based care: Not all clinicians need to appraise evidence from scratch but all need some skills. *BMJ: British Medical Journal* 320(7240), 954–955 (2000)
13. Haynes, R.B., Haines, A.: Barriers and bridges to evidence based clinical practice. *BMJ: British Medical Journal* 317, 273–276 (1998)
14. Tang, H., Ng, J.H.K.: Googling for a diagnosis—use of Google as a diagnostic aid: internet based study. *BMJ* 333(7579), 1143–1145 (2006)
15. Giustini, D.: How Web 2.0 is changing medicine. *BMJ: British Medical Journal* 333(7582), 1283 (2006)
16. Karkalis, G.I., Koutsouris, D.D.: E-health and the Web 2.0. In: itab 2006, Congress Center Du Lac, Ioannina, Greece (2006)

An Agent Based Pervasive Healthcare System: A First Scalability Study^{*}

Johannes Krampf¹, Stefano Bromuri¹, Michael Schumacher¹, and Juan Ruiz²

¹ Institute of Business Information Systems,
University of Applied Sciences Western Switzerland
{johannes.krampf, stefano.bromuri, michael.schumacher}@hevs.ch,
<http://aislab.hevs.ch/>

² Department of Endocrinology, Lausanne University Hospital, 1011 Lausanne
juan.ruiz@chuv.ch

Abstract. Gestational Diabetes Mellitus (GDM) occurs during pregnancy due to an increased resistance to insulin caused by the growth of the baby. It appears after the 24th week of pregnancy and it is treated with diet counselling and insulin treatment. In this paper we present the complete implementation of a Pervasive Healthcare System (PHS) based on intelligent agents to support continuous monitoring of pregnant women affected by GDM. Our infrastructure is composed of a mobile interface connecting to a distributed multi-agent system which in turns is connected to a patient management system. This stores the data produced during the monitoring phase and present them to the doctors in charge of the patient. Our system's scalability is then evaluated to show the strong and weak points of our approach.

Keywords: Personal Health Systems, Agents, Pervasive Healthcare.

1 Introduction

Gestational diabetes mellitus (GDM)[10] affects 3%–10% of all pregnant women with no history of diabetes before pregnancy and manifests itself in high blood sugar levels during pregnancy. Current treatment guidelines [13] consist in diet adjustment and in anti-diabetic medicines such as insulin and metformin. In particular, the patient starts the treatment by simply monitoring the levels of glucose 4 times per day, with one *preprandial* observation and one *postprandial* observation in the morning, and two *postprandial* observations after the lunch and after the dinner. Such values are then written in a notebook that is handed to the doctors twice weekly. According to the behaviour of the physiological values the doctors may introduce further checks at lunch and dinner, and, if the glucose values are outside the boundaries, start the treatment with metformin or insulin. If not treated, GDM may have severe risks for the mother, who may develop high blood pressure and protenuria (preeclampsia) [14], and for the baby,

^{*} This work has been partially funded by the Hasler Stiftung and by the Nano-Tera Consortium.

who may become large for the gestational age (macrosomia), with complications at delivery and later in life.

Rather than checking the patient once or twice weekly, a better monitoring may allow doctors to assess the situation of the patient and propose the correct treatment. One approach to continuous and intelligent patient monitoring is represented by pervasive healthcare [12]. The goal of a pervasive healthcare system (PHS) is to break the boundaries of hospital care, allowing patients to be monitored while living their day-to-day life and to keep in touch with healthcare professionals. Due to its distributed nature, a PHS is faced with three main challenges: scalability, accuracy and security. Scalability is important for PHSs as these systems must be able to serve many patients at the same time, without experiencing disruption due to high loads. Secondly, an accurate PHS should be able to filter information efficiently in order to save time to the healthcare professionals and produce alerts only when needed, with a good trade-off between false positives and false negatives. Finally, security is also an important dimension for a PHS as it deals with medical data, which is sensitive data.

In this paper we present a PHS to monitor patients affected by GDM. A previous version of this system was presented in [2], where we modelled a distributed agent-based PHS. We have chosen agents as a modelling abstraction for our PHS as they are understood to be autonomous software entities, that act proactively and pursue a set of goals [15] in an intelligent way, by applying AI reasoning techniques. Using multi-agent systems (MAS) abstractions to model PHSs is beneficial as this programming paradigm is well suited for distributed systems, due to the autonomy property of the agents, and thanks to distributed event based approach that these systems take into consideration to model the interactions between the agents and the other available resources [3].

In [2] we have already provided a first validation of the accuracy of the notifications provided to the health professionals by our intelligent agents. On one hand, in this paper we present the full implementation of our PHS for GDM, evaluating the scalability of our system and illustrating how healthcare professionals can utilise the functionalities of our tool. On the other hand, the security of our PHS will be evaluated in future publications as the system is currently being audited for security at the Lausanne University Hospital, although in this paper we also present how we secured the interfaces of our PHS. The remainder of this paper is structured as follows: Section 2 discusses the functionalities of the components of our system; Section 3 discusses an evaluation of our PHS in terms of its scalability; Section 4 puts our work in comparison with relevant related works; finally Section 5 concludes this paper and draws the lines for future work.

2 A Personal Health System for GDM

Fig. 1 shows that our system is composed of three main components, which are the *Mobile Infrastructure* (MI), the *Agent Environment* (AE) and the *Patient Management System* (PMS). Furthermore, these components are interfaced between each others by means of a mediator component, realised as a Web service




Fig. 1. The Pervasive Health System Logic Architecture and the Security Interfaces

Data Gateway connector that accepts HTTPS requests. The MI component collects the physiological data of the patient and delivers such data to the AE component and to the PMS component.

The AE component utilises logic programming to model intelligent agents that filter the data submitted to the PMS and provide alerts in case of significant events, such as a possibility of preeclampsia in the patient or a high level of blood sugar that requires a treatment adjustment. The AE system is subdivided in cells associated to an area of a real city where the patients connect with their mobile phones to produce their physiological data, that are then evaluated by intelligent agents. The patients are represented in the AE as *avatars* that can communicate to a personal intelligent agent, embodied in the AE. This representation of the patient is convenient as we can reuse the AE communication and notification facilities to interact with the intelligent agents situated in it. To every patient avatar we associate an intelligent agent whose cognitive architecture will be explained later in this Section. Finally the PMS allows the doctors to visualise the patient’s data, to modify its treatment and to visualise the alerts produced by the AE.

The three tier logic architecture shown on the top of Fig. 1 translate then to a four tier architecture as shown on the bottom of Fig. 1. In particular, the mobile phone and the Web browser represent the *presentation layer*, the reverse proxy and the Web server represent the *Web application layer*, the agent environment represents the *business logic layer* while the database represents the *data layer*.

The Web application layer accepts outside secure connections only on the HTTPS port. It connects business logic and data layers. Caretakers and patients use client authenticated HTTPS to connect to the system. A second authentication factor is provided by the combination of user name and password. Internal components use local in-memory or mutually authenticated TLS connections to communicate with each other. The data base partition is encrypted to protect against physical access to the hard disk. User access to resources is restricted by membership in one of the three groups users, caretakers and administrators. Access to patient data is further restricted by an access control list which only allows caretakers who treat a patient to access this patient's data. All actions are logged including IP address, user name, resource and success of the action to provide an audit trail.

2.1 The Mobile Infrastructure

The MI component is based on Android, and it is provided to the patient through mobile phones and tablets to introduce their physiologic data and symptoms associated to GDM: blood pressure, blood glucose, weight, pulse, blurred vision, epigastric pain, oedema in the legs, dyspnoea, chest pain, head ache. Such an interface allows the patients to see if their physiological values meet the targets for the week, with a set of pie charts as shown in Fig. 2. At the same time the MI is also built with a synchronisation approach to avoid data loss: whenever the connection with the AE is impossible, the MI saves the data in a local database. When the connection with the AE is possible again, the data stored in the mobile phone is submitted for storage in the PMS.




Fig. 2. The Mobile Interface

2.2 The Agent Environment and the Patient Management System

Our PHS makes use of intelligent agents to analyse and filter the physiological data produced by the patients. In particular, we decided to include the GOLEM agent platform in our system as it currently implements the patterns of Distributed Event-Based Systems (DEBSs) as described in [2]. By means of these patterns, the agents in the GOLEM agent platform can subscribe to events produced by the patients and the GOLEM platform will take care of notifying such events when they take place.

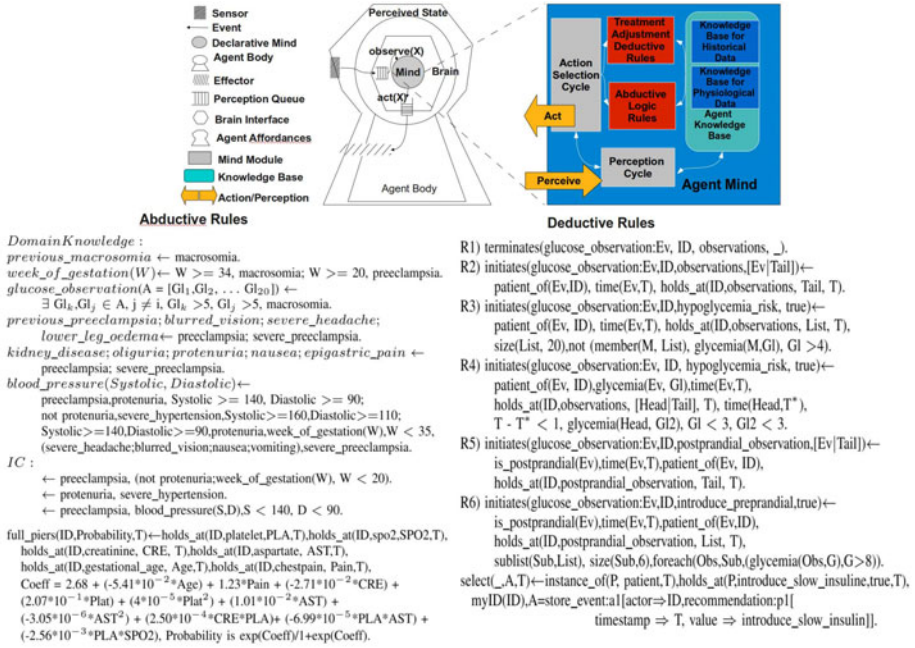


Fig. 3. Caretaker Agent Mind Architecture

In the particular case of our PHS, the patients are represented in the AE as avatars that can communicate with a personal caretaker agent, whose architecture is reported in Fig. 3. Every agent is deployed in GOLEM in a container. A GOLEM container represents a portion of the distributed agent environment which in this case is associated with a portion of the real environment, in order to distribute the load of the requests of the patients. This topology was chosen because we imagine that this system could work in synergy with the actual cellular network. As described in [2], every caretaker agent has a cognitive model with a deductive and an abductive part, whose specification is shown in Fig. 3.

The deductive rules are specified in Event Calculus [8], to describe the evolution in time of the patient physiological values. Such rules specify how the treatment of the patient should evolve. For example rule R6 specifies that, when the patient had high glucose in the postprandial observations, then the agent suggests to introduce further preprandial observations. Similarly, if the patient is already in a 6 checks per day regime, then the agent suggests the doctors to introduce a slow insulin in the morning to tackle the values that are out range. The abductive rules take into consideration the symptoms of the patient to provide alerts of macrosomia or preeclampsia to the doctors. In particular, for preeclampsia, we also provide the probability of adverse outcome using the fullPiers model [14], also reported in Fig. 3. To be able to provide this probability, the agent connects to the PMS using the GOLEM middleware to download the blood samples needed by the fullPiers model and introduced in the system

through the PMS. Further details about the agent cognitive model and its accuracy are reported in [2] and we refer the interested reader to this publication.

The Patient Management System allows healthcare professionals to visualise and analyse data as well as to introduce new data gathered during a patient's visit. The PMS is a hybrid application incorporating both elements of a classic server side Web application and a modern AJAX-powered client side Web application.




Fig. 4. The Patient Overview Page

After logging in to the PMS, healthcare professionals can visualise a dynamic patient page, shown in Fig. 4. The patient page is divided into a static side bar on the left and a changing content area on the right side of the window. When showing different information about the patient, an HTML fragment will be loaded and replace the content area. Graphs are handled completely on the client side. There is a maximum of about 1700 data points, assuming 6 daily blood sugar measurements during 40 gestational weeks, for a graph. Processing and displaying this number of data points is almost instant on modern browsers and allows for better interactivity than server side processing would. Healthcare professionals can interact with the graphs by showing a specific time period, changing this period and changing the view on the data. There exist two further views on the data beside the point view in Fig. 4, which are a percentage view, showing the data points in a week which are below, inside and above the normal range of a physiological value, and a distribution view that shows the 9th, 25th, 50th, 75th and 91th percentile for each week as well as outliers. These views were created to fulfil the different information needs of doctors at the Lausanne University Hospital.

3 Evaluation

To evaluate our solution, we measured the performance for HTTPS requests with different requirements on the application and database. Our goal in evaluating our PHS was to understand if the system could support the traffic load of

the patients of an hospital of a medium sized city, such as the city of Lausanne in Switzerland, where we plan to perform field tests. Also, another goal of this evaluation is to understand what is the maximum amount of patients that we can serve before having to introduce load balancing techniques in the PHS. We therefore perform our evaluation on those components representing a bottleneck of the current architecture. We do not perform an evaluation on the agent environment as this is based on the GOLEM platform, whose performances have been previously evaluated in [3], showing that the system can scale up with the number of GOLEM containers spawned for the application. For the tests, we ran our PHS on a 3 GHz Intel Core 2 Duo processor, 4 GB RAM and Ubuntu 10.04.




Fig. 5. Requests/second results and 90th percentile response time results

The tests were performed using the ApacheBench [1] utility, which works by performing a defined number of requests to a specified URL and measuring various values such as response time or transferred bytes for each request. After finishing the benchmark, ApacheBench shows a statistical analysis of all requests, showing transfer rate and mean time per request. To prevent the benchmark utility from influencing the results, it was executed on a notebook with a 1.6 GHz Intel Core 2 Duo processor, 1 GB RAM running Linux with a direct cable connection to the server, to minimise influences caused by network latency. We tested different usage scenarios stressing different parts of the system. All requests were executed at 1 to 100 concurrent requests to simulate different usage load. Values were recorded during a 60 seconds stress testing period.

The curves on the left of Fig. 5 show the response time to different concurrent calls performed in the system. In particular, we show the time to: retrieve the data for glucose in the PMS; retrieve a patient history; insert glucose values from the mobile phone; retrieve a patient summary page; access the login page and patient page. We also assume that concurrent calls coming from different patients are distributed in different cells of the agent environment. The requests per second values increase with higher concurrency levels until a plateau is reached. At this point the server becomes overloaded and the response time increases. The graph of requests to the login page shows the plateauing behaviour when reaching about 1400

requests per second. By looking at the detailed values we discovered that the response time increase between 400% and 500% when comparing 20 and 100 concurrent connections while the processed requests per second are virtually unchanged. The curves on top of Fig. 5 also show the plateauing behaviour of the requests with database activity. The maximum value for requests per second is reached between 10 and 20 concurrent requests. This suggests a database related limit in concurrency. When increasing the concurrency from 10 to 15, the glucose insertion from the mobile phones experiences a sharp fall from 130 to 20 requests per second. The message queue in the Web service data gateway interface is the limiting factor here: messages are acknowledged in order to ensure message delivery and with many concurrent requests messages cannot be acknowledged fast enough.

The 90th percentile response time charts on the right of Fig. 5 show the expected maximum response time for 90% respectively of all requests. The response time begins to be above 1 second and noticeable by a user at about 50 concurrent requests for requests with medium database activity and at about 15 concurrent requests for database heavy requests.

Each patient will transmit a number of values each day: Twice daily blood pressure and four to six times blood sugar. Furthermore, the patient will transmit one weight value per week and she will report symptoms when she experiences them. We are interested in the maximum number of requests in a short time period and will make the pessimistic assumptions that during one second of usage of her mobile phone the patient transmits symptoms, blood pressure, blood sugar and weight at the same time in the morning. The maximum number of requests/second for a user is therefore 4 requests/second. The worst case scenario is all 10 patients of a planned pilot study making their 4 requests concurrently in the same second, leading to 40 requests/second with 10 concurrent connections. When producing Fig. 5 we found that the system is capable of 132 requests/second for a concurrency level of 10. As at the Lausanne University Hospital, that serves the Canton Vaud in Switzerland, there are a maximum of 5–6 patients with GDM at the same time, the system we defined is viable to deal with the load experienced by a big sized university hospital. To estimate the maximum number of users the system can serve, we modify our assumptions to assume an uniform distribution of the 4 requests over the course of 30 minutes. We will furthermore use 20 requests/second as the system's performance due to performance drop off at higher concurrency levels. This results in about 0.002 requests/second per user (Eq. 1) and 9000 patients (Eq. 2) with 20 concurrent connections (Eq. 3). This allows us to consider usage for the whole canton Vaud. The canton has a population of 700,000 [11] and 9.4 births per 1000 inhabitants per year. This results in 6580 births per year (Eq. 4) which means that the system can theoretically monitor all pregnant women in the Canton of Vaud.

$$\frac{4}{30 * 60} \frac{request \times user}{second} = \frac{1}{450} \frac{request \times user}{second} \quad (1)$$

$$20 \frac{request}{second} / \frac{1}{450} \frac{request \times user}{second} = 9000 user \quad (2)$$

$$(4 \frac{request}{user} \times 9000 user) / 30 \times 60 second = 20 \frac{request}{second} \quad (3)$$

$$700000 inhabitant \times \frac{9.4}{1000} \frac{birth}{inhabitant \times year} = 6580 \frac{birth}{year} \quad (4)$$

4 Related Work

From the related work stand point, several attempts have been done in the past to combine agent technology with the healthcare domain. The systems described by Huang et al. in [7] and by Hammond and Sergot in [6] use symbolic reasoning over clinical workflows to manage oncological patients within a healthcare institution and to simplify the management of clinical trials. Larson et al present Guardian in [9], an early attempt to provide an agent-based system for medical monitoring and diagnosis. Guardian uses a tuple space based approach where cognitive agents with a properly programmed knowledge base, provide a diagnosis for situations such as liver failure and hypothermia. In [4] Ciampolini et al present a distributed MAS to deal with distributed diagnosis performed by heterogeneous distributed abductive agents. In Ciampolini's approach the diagnosis is provided in term of probabilities, although they do not consider a realistic model for their experiments. The ASPIC project [5] has developed an architecture based on argumentation theory for an autonomous agent that single- and multi-agent healthcare applications can use. Evaluation scenarios focus on the management and treatment of people with heart disease. With respect to the systems reported above, our contribution is twofold: first of all we developed a practical system that takes into consideration scalability and security issues following the needs of medical doctors at Lausanne University Hospital; secondly, for our intelligent reasoning agents, we also utilise clinical models like the fullPiers [14], whereas the systems mentioned above lack this approach.

5 Conclusion and Future Works

In this paper we presented a fully implemented Personal Health System for the monitoring of GDM and alerting of treatment adjustment suggestions and continuous diagnosis of conditions related to GDM. This system is composed of a mobile infrastructure used by patients, a distributed agent environment and a patient management system for medical professionals in charge of the patients. We evaluated the performances in terms of scalability of our system demonstrating the feasibility of the approach in the case of GDM. Future work implies the evaluation of security interfaces amongst the different layers of our infrastructure and the produced data. Another approach to future work is to have a different PHS to doctors notification system, where the notifications are not just produced in the Patient Management System, but also imply the submission of emails or SMS to the medical professionals.

References

1. ApacheBench, <http://httpd.apache.org/docs/2.2/programs/ab.html>
2. Bromuri, S., Schumacher, M., Stathis, K., Ruiz, J.: Monitoring gestational diabetes mellitus with cognitive agents and agent environments. In: Proceedings of the 2011th IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT 2011) (August 2011)

3. Bromuri, S., Stathis, K.: Distributed Agent Environments in the Ambient Event Calculus. In: DEBS 2009: Proceedings of the Third International Conference on Distributed Event-based Systems. ACM, New York (2009)
4. Ciampolini, A., Mello, P., Storari, S.: Distributed medical diagnosis with abductive logic agents. In: ECAI 2002 Workshop on Agents in Healthcare, Lione (2002)
5. Fox, J., Glasspool, D., Modgil, S.: A canonical agent model for healthcare applications. *IEEE Intelligent Systems* 21, 21–28 (2006)
6. Hammond, P., Sergot, M.: Computer support for protocol-based treatment of cancer. *The Journal of Logic Programming* 26(2), 93–111 (1996)
7. Huang, J., Jennings, N.R., Fox, J.: Agent-based approach to health care management. *Applied Artificial Intelligence* 9(4), 401–420 (1995)
8. Kowalski, R., Sergot, M.: A logic-based calculus of events. *New Gen. Comput.* 4(1), 67–95 (1986)
9. Larsson, J.E., Hayes-Roth, B.: Guardian: An intelligent autonomous agent for medical monitoring and diagnosis. *IEEE Intelligent Systems* 13, 58–64 (1998)
10. Serlin, D.C., Lash, R.W.: Diagnosis and management of gestational diabetes mellitus. *Am. Fam. Physician* 80, 57–62 (2009)
11. Swiss Statistics. The population of Switzerland 2009 (2010), <http://www.bfs.admin.ch/bfs/portal/en/index/themen/01/22/publ.Document.136821.pdf>
12. Varshney, U.: *Pervasive Healthcare Computing: EMR/EHR, Wireless and Health Monitoring*. Springer Publishing Company, Incorporated (2009)
13. Vogel, N., Burnand, B., Vial, Y., Ruiz, J., Paccaud, F., Hohlfeld, P.: Screening for gestational diabetes: variation in guidelines. *Eur. J. Obstet. Gynecol. Reprod. Biol.* 91, 29–36 (2000)
14. von Dadelszen, P., et al.: Prediction of adverse maternal outcomes in pre-eclampsia: development and validation of the fullPIERS model. *Lancet* 377, 219–227 (2011)
15. Wooldridge, M.: *MultiAgent Systems*. John Wiley and Sons (2002)

An Agent-Based Approach to Real-Time Patient Identification for Clinical Trials

Gareth Tyson¹, Adel Taweel¹, Simon Miles¹,
Michael Luck¹, Tjeerd Van Staa², and Brendan Delaney¹

¹ King's College London, UK

² General Practice Research Database, UK

Abstract. Patient recruitment for clinical trials is expensive and has been a significant challenge, with many trials not achieving their recruitment goals. One method that shows promise for improving recruitment is the use of interactive prompts that inform practitioners of patient eligibility for clinical trials during consultation. This paper presents the ePCRNI-IDEA recruitment system, which utilises an agent-based infrastructure to enable real-time recruitment of patients. In essence, whenever patients enter a clinic, the system compares their details against eligibility criteria, which define the requirements of active clinical trials. If a patient is found to be eligible, a prompt is raised to notify the user. In this way, it becomes possible for recruitment to take place quickly in a cost effective manner, whilst maintaining patient trust through the involvement of their own health care practitioner.

1 Introduction

Clinical trials are the gold standard by which medical research is evaluated. They are used to study various aspects of medical science, as well as being a vital stage in the deployment of new drug treatments. Currently, however, such trials are frequently unsuccessful at recruiting sufficient patients. A review of the UK Medical Research Council found that only 31% of trials actually recruited to their planned target, with 30–40% of costs arising during the recruitment phase alone [1]. This is because discovering and contacting eligible potential recruits is both logistically and legally challenging. Consequently, many research projects take far longer to complete than is desirable, resulting in an unnecessary burden for those who could potentially benefit from the results.

The main challenge for patient recruitment lies in locating and contacting patients in a sufficiently timely manner to allow them to participate. However, the ease with which this can be done varies dramatically with the type of trial; for instance, recruitment can be challenging for trials that have high recruitment targets or complex eligibility criteria. Currently, recruitment is performed in a highly laborious manner, which is ill-suited to the above situations. It often involves a human recruitment agent visiting clinics in an attempt to locate suitable patients (e.g. asking practitioners or searching local medical records). This creates significant overhead as it is both slow and costly, as well as non-scalable

for most trials. For example, a trial investigating rare ailments might need an extensive number of visits to locate sufficient patients.

In consequence, it is of paramount importance to address the recruitment challenges within trials to ensure the future efficacy of medical research. It is therefore necessary to find a scalable way by which eligible patients can be discovered. To address this, we propose replacing human agents with software agents that permanently reside within clinics, with the aim of searching for patients who might satisfy trial eligibility criteria. Through this, the agent could inspect patient information in real-time to ascertain eligibility before presenting notifications to local practitioners. Importantly, by using a software agent this could be done rapidly within a consultation before a patient has left, thereby shortening the recruitment lifecycle (as well as enabling trials based on incidental cases). The paper's contributions are therefore as follows:

- A critique of existing clinical trial recruitment approaches, highlighting that current techniques are both slow and expensive.
- An agent-based distributed architecture called ePCRNI-IDEA that enables real-time recruitment of patients, whilst avoiding the key limitations of existing approaches.
- A procedure by which software agents can guide the recruitment of patients to their most appropriate trials.

The rest of the paper is structured as follows; first the background to the research is discussed in Section 2 before detailing the ePCRNI-IDEA recruitment system in Section 3. Following this, a discussion is presented in Section 4, alongside future work and conclusions.

2 Background and Related Work

2.1 Recruitment in Clinical Trials

Clinical trials are a challenging stage in the research of clinicians due to the complexity of recruiting patients for participation. Many types of trials can suffer from such difficulties; for instance, trials that have potential recruits who are widely distributed over many clinics (e.g. primary care) are extremely difficult to recruit for due to the intense resource requirements. Similarly, trials with certain types of patient eligibility criteria can be difficult to recruit for; for example, a trial dealing with incidental/acute conditions would need a practitioner to immediately identify patients in real-time. This can be highly challenging, especially when dealing with complex eligibility criteria or trials that require immediate actions (e.g. a change of drug treatments). This has a significant impact on medical research and stunts potential life-saving advances. Recruitment is currently performed in a number of ways; examples include:

- *Recruitment Visits*: Using recruiters to visit or contact practices so that they can search local medical records and/or ask local practitioners.

- *Advertisements*: Using posters, web sites, mailing lists or newspaper adverts describing what is required for the trial.
- *Practitioners*: Using practitioners to suggest patients who might be eligible. Alternatively, practitioners may be asked to look for patients (during consultations) who match the criteria in real-time.

Clearly, the above approaches are slow and often quite expensive. The predominant approach of *employing recruiters* to contact and visit clinics is laborious and often infeasible, especially when dealing with the above types of trials. For instance, different clinics often use different database systems, thereby requiring a high degree of training for recruiters; something that is exacerbated by the limited in-house IT training in many clinics. Further, attempts to standardise this database search procedure (e.g. MIQUEST) are often poorly understood by researchers, whilst such standards can often be undermined by variations in coding and medical dictionaries (e.g. if multiple databases use different semantics). Consequently, there is little automated support for finding eligible patients, often leaving recruitment very much as an ad-hoc process that falls outside of recognised standards; Moreover, there is little infrastructure available that can be exploited by clinicians to assist in the stage.

As an alternative to the recognised approaches, Embi et. al. [2] propose the use of a Clinical Trial Alert (CTA) system, through which practitioners could be notified while they are in consultation with a patient who is eligible for a trial. Their work shows that significant increases in recruitment could be achieved. However, their pilot study was only deployed in a single clinic with a single trial, in an attempt to study the improvements in recruitment. Other similar techniques have also seen only limited large-scale testing, e.g. [3]. Consequently, a number of issues (e.g. scalability) have not been investigated, leaving the sole focus on the recruitment outcomes. To address these concerns, this paper therefore focusses on more infrastructural aspects for enabling a large-scale trial alert system, specifically through the use of intelligent agents.

2.2 Agents in Healthcare

Agents have emerged as a prominent technology for handling a range of real-world problems [4]. An agent can be defined as “a computer system that is situated in some environment, and that is capable of autonomous action in their environment in order to meet its delegated objectives” [5]. Agents in healthcare have seen widespread investigation; in general, their use attempts to address endemic issues such as *(i)* distributed information and expertise, *(ii)* unpredictable dynamics, and *(iii)* uncertainty in reasoning and data.

For example, MAID [6] is an agent-based system for integrating heterogeneous data sources within a hospital environment. The hospital studied had 24 departments, each using their own information systems. To address this, agents were constructed to interoperate with each system to monitor changes and retrieve data for insertion into a central repository. In a subsequent work,

HealthAgents [7] went beyond MAID to also enable decision support, specifically for diagnosing brain tumours. A range of agent-based systems have also been proposed for handling distributed expertise. These includes using agents to enable better communication between healthcare workers based on ambient information, e.g. their role, location etc. [8], as well as using agents to remotely monitor patients [9][10]. These systems also often involved data analysis; S(MA)²D [10], for instance, uses statistical analysis to cluster patients into similar groups. This ability to scalably perform data analysis in real-time, clearly, also shows potential for enabling the type of eligible patient identification discussed previously. Despite this, so far little work has been performed into using agents to improve clinical trial recruitment. Consequently, the rest of this paper explores exploiting the properties of agents to enable real-time recruitment of patients to trials.

3 ePCR-IDEA Recruitment System

This section presents the ePCR-IDEA recruitment system, which is designed to enable real-time identification of eligible patients. We first present the overall architecture before describing the individual components, focussing on the behaviour of the clinic-based agents.

3.1 Overview

The core goal of the ePCR-IDEA recruitment system is to improve patient recruitment. In order to do so, researchers must formally define the eligibility criteria of participants, then distribute it to software agents that reside locally on practitioners' computers. These software agents listen to interactions between the practitioner's local Electronic Healthcare Record (EHR) database and the user in an attempt to locate patients who are eligible for trials. Importantly, this occurs in real-time during the consultation, thereby allowing a pop-up to be generated, notifying the practitioner of the patient's eligibility. In this way, the patient can be instantly consulted regarding the trial and, if interested, recruited via a web interface. The key architectural entities in the system are as follows:

- *LEPIS*: An agent that resides at primary care practices and investigates the eligibility of any present patients, termed the Local Eligible Patient Identification Service.
- *CCS*: A point of storage and distribution that allows clinical researchers to inject new trials into the system, termed the Central Control Service.
- *CTMS*: A website that handles the actual recruitment process once an eligible patients has been discovered, termed the Clinical Trial Management System.

Prototype implementations of all these components have already been developed. Fig. 1 provides an overview of these, as detailed in the rest of this section.




Fig. 1. Overview of ePCRN System

3.2 Central Control Service (CCS)

A *trial store* maintains a repository of active trials within the system, as entered by any clinical researchers wishing to recruit patients. Trials are stored using a standard model (the PCROM standard [11]), which defines the data format used to represent the various aspects of the trials. This includes a description of the trial as well as the eligibility criteria, which can be based on a variety of aspects ranging from a pre-computed list of eligible patient identifiers to complex diagnosis information. The trial store is realised within the ePCRN-IDEA system through the Central Control Service (CCS), which is a service used to manage all the trials. When a new trial is created, it is injected into the system via the CCS and stored in a MySQL database back-end before being made accessible to the appropriate (and authorised) parties — namely, the recruitment agents. These are accessed securely using an encrypted SQL connection initiated by each agent to the CCS.

3.3 Local Eligible Patient Identification Service (LEPIS)

The Local Eligible Patient Identification Service (LEPIS) is a Java-built software agent that resides on every practitioner's PC (who is authorised to recruit patients). The purpose of this agent is to actively discover any eligible patients who might be seen by the practitioner. It therefore operates as a replacement for the existing human agents who visit clinics in an attempt to locate suitable patients. Every LEPIS agent is required to obtain two distinct sets of information. First, it must acquire information about patients in real-time as they enter the clinic. Second, LEPIS must also acquire the necessary trial information (from the CCS) so that it can compare patients against each trial's eligibility criteria; ideally, eligibility should also be computed locally to avoid unnecessary delays or bottlenecks.

Accessing Patient Information. To discover the eligibility of a patient for a trial, it is first necessary for LEPIS to gain access to any pertinent information

related to the patient. Whenever a patient enters a clinic, the practitioner opens his or her medical record using the Electronic Healthcare Record (EHR) database on their desktop computer. This medical record contains a range of demographic and medical information about the patient, allowing practitioners to retrieve information about the patient during the consultation, as well as enter new information. This offers an existing platform through which a recruitment agent can access information about patients. We have modified a popular EHR system, *Vision* [12], to interact with LEPIS. Whenever a patient record is opened or modified, the information is passed to LEPIS (through a standard file using a shared XML schema) so that it can attempt to locate trials for which the patient may be eligible. Information is coded using standard Read Codes and Multilex Drug Codes to allow LEPIS and the EHR to understand each other. LEPIS is therefore given real-time access to information about any patients who are currently in consultation. Although, evidently, eligibility criteria is limited to those attributes provided by the EHR, which can vary based on both policy and EHR implementation (many EHR vendors exist).

Accessing Trial Information. To allow a LEPIS agent to compute a patient's eligibility, it must first gain access to trial information. Ideally, this should be stored locally to enable real-time eligibility checks within a consultation. However, evidently, this is largely infeasible with the huge number of active trials running; e.g. clinicaltrials.gov currently lists well over 100,000 trials. Consequently, it is necessary for each agent to independently select the most appropriate trials for its clinic and practitioner. Each agent thus maintains a set of trials T of size n , as limited by the host's local resources. An agent therefore selects n based on the capabilities of its host, by performing eligibility checks on a random set of trials repeatedly for one second; n is then set as the number of iterations. It then keeps a persistent record of all patient identifiers, Read Codes and Multilex Codes provided by the EHR to build up a profile of the clinic. Using this information, n trials are retrieved from the CCS through the following process:

1. LEPIS attempts to retrieve a set of n trials from the CCS containing:
 - (a) p trials that includes a known patient registered within the clinic ($p = < n$);
 - (b) if $p < n$, c trials that includes coded information previously encountered within the clinic ($c = < n - p$); and
 - (c) if $p + c < n$, r randomly selected trials ($r = < n - p - c$).
2. Remove any trials that are fully recruited.

These two steps are repeated throughout an agent's lifetime with a configurable interval, which is set to 24 hours by default.

Computing Eligibility and Generating a Popup. When LEPIS acquires patient information from the EHR, it must compare it against the eligibility criteria of any known trials. This is a simple process that currently involves iteratively computing eligibility for each known trial and then selecting a random one if multiple are found. A popup is then generated to notify the user. Fig. 2 shows a screenshot of the user interface.

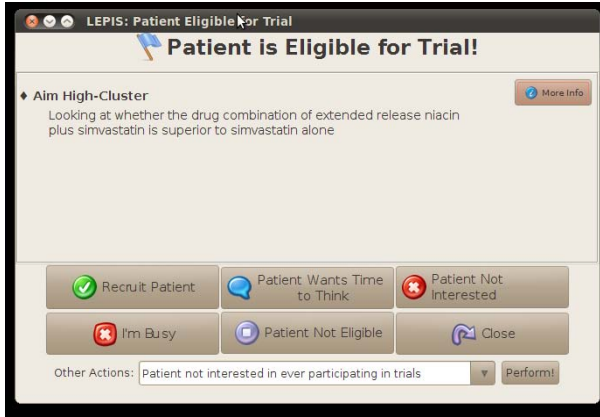


Fig. 2. Screenshot of LEPIS Popup

3.4 Clinical Trial Management System (CTMS)

If a patient is interested in being recruited for a particular trial, it is then necessary to actually perform the recruitment procedure. This is not handled by the local agent; instead, an external website is used, called the Clinical Trial Management System (CTMS). The CTMS is securely accessed by the practitioner and then used to register the patient's interest in being recruited. Any necessary steps can then be taken, e.g. contacting the patient, recording information etc.

4 Discussion, Future Work and Conclusion

The ePCRN-IDEA recruitment systems differs substantially from previous work in that it is agent-based. Our early evaluation shows that the agent-based approach is promising and has several potential advantages over traditional client-server approaches (e.g [2]). Through the use of agents, intelligence is decentralised within the system so that both computation and decision making is independently performed by each agent.

Primarily, we have used the approach to enable superior scalability. This is critical, especially for patient recruitment in primary care. Even for common diseases, eligible patients are thinly spread across many practices, which potentially number several thousand. In the UK, alone, there are more than 10,000 practices. A client-server approach could be realised in two ways: either all patient information could be transmitted to a server, where it centrally computes eligibility; or all trial information could be transmitted to all clinics for local computation. The former is non-scalable as well as dangerous in terms of privacy and security. The latter, however, is also highly non-scalable as the number of trials (e.g. greater than 100,000) alongside the size of each trial description (e.g. 0.5 MB) makes it impossible for all clinics to know of all trials. Consequently, to address

this, we embed intelligence within the agents to learn how to best select trials for their host clinic, exploiting the local knowledge (and computational abilities) of each agent, rather than burdening a central point.

From our initial phase-1 prototype we have identified a number of future lines of work. First, we aim to complete a full system deployment within the UK primary healthcare system, thereby enabling a detailed quantitative evaluation. Beyond this, we also intend to extend the agent capabilities. Key research lines include, (i) *inter-agent collaboration*: allowing agents to build societies to better enable information and resource sharing (e.g. based on disease areas, localities); (ii) *interface adaptation*: allowing agents to learn (and share) the behaviour of users to adapt interaction; and (iii) *trial negotiation*: allowing agents to negotiate with each other to best distribute trials based on runtime conditions.

References

1. McDonald, A., Knight, R., Campbell, M., Entwistle, V., Grant, A., Cook, J., Elbourne, D., Francis, D., Garcia, J., Roberts, I., Snowdon, C.: What influences recruitment to randomised controlled trials? a review of trials funded by two uk funding agencies. *Trials* 7(1), 9 (2006)
2. Embi, P.J., Jain, A., Clark, J., Bizjack, S., Hornung, R., Harris, C.M.: Effect of a clinical trial alert system on physician participation in trial recruitment. *Arch. Intern. Med.* 195(19) (2005)
3. Rollman, B.L., Fischer, G.S., Zhu, F., Belnap, B.H.: Comparison of electronic physician prompts versus waitroom case-finding on clinical trial enrollment. *Journal of General Internal Medicine* 23(4) (2008)
4. Luck, M., McBurney, P., Preist, C.: Agent Technology: Enabling Next Generation Computing (A Roadmap for Agent Based Computing). AgentLink (2003)
5. Wooldridge, M.: Introduction to Multiagent Systems. John Wiley & Sons, Inc. (2009)
6. Cruz-Correia, R., Vieira-Marques, P., Costa, P., Ferreira, A., Oliveira-Palhares, E., Araújo, F., Costa-Pereira, A.: Integration of hospital data using agent technologies - a case study. *AI Commun.* 18 (August 2005)
7. González-Vélez, H., Mier, M., Julià-Sapé, M., Arvanitis, T., García-Gómez, J., Robles, M., Lewis, P., Dasmahapatra, S., Dupplaw, D., Peet, A., Arús, C., Celda, B., Van Huffel, S., Lluch-Ariet, M.: Healthagents: distributed multi-agent brain tumor diagnosis and prognosis. *Applied Intelligence* 30 (2009)
8. Rodríguez, M.D., Favela, J., Preciado, A., Vizcaíno, A.: Agent-based ambient intelligence for healthcare. *AI Commun.* 18, 201–216 (2005)
9. Koutkias, V.G., Chouvarda, I., Maglaveras, N.: A multiagent system enhancing home-care health services for chronic disease management. *IEEE Transactions on Information Technology in Biomedicine* 9, 528–537 (2005)
10. Rammal, A., Trouilhet, S., Singer, N., Pécatte, J.-M.: An adaptive system for home monitoring using a multiagent classification of patterns. *Int. J. Telemedicine Appl.* 2008, 3:1–3:8 (2008)
11. Speedie, S.M., Taweel, A., Sim, I., Arvanitis, T.N., Delaney, B., Peterson, K.A.: The primary care research object model (pcrom): A computable information model for practice-based primary care research. *Journal of the American Medical Informatics Association* 15(5) (2008)
12. Vision, <http://www.inps4.co.uk/vision/>

The Use of Social Bookmarking by Health Care Students to Create Communities of Practice

Ed de Quincey, Avril Hocking, Josephine O’Gorman, Simon Walker, and Liz Bacon

e-Centre, University of Greenwich
Old Royal Naval College, Park Row London, UK
e.de.quincey@gre.ac.uk

Abstract. Teaching and learning health and social care in a digital age produces many challenges for students and their teachers. A common hurdle for healthcare students and practitioners is the sheer amount of information that they have to make sense of. Another challenge is where this information is captured and stored, with people utilising personal, as well as institutionally owned devices. A potential solution to these problems is the use of social bookmarking applications such as “delicious”, where users can create a centralised repository of online resources, share them with other users, and view what others are bookmarking. This paper describes research conducted at the University of Greenwich involving 160 participants across three Schools and 5 modules, including Health and Social Care who were encouraged to integrate social bookmarking into their learning and teaching. Participants were instructed to tag their resources with an appropriate module code tag e.g. NURS1297 so that a repository of module specific bookmarks was created. Over a 4 month period, 160 users created 1430 bookmarks with 5032 tags. Further analysis of the bookmarking behaviour is discussed along with reflections on the suitability of social bookmarking to create digitally literate health care communities of practice.

Keywords: Social bookmarking, tagging, eLearning, Web 2.0, communities of practice, connectivism, heutagogy, digital literacy.

1 Introduction

A common problem for students studying health related subjects in Higher Education (HE) is that they can access the web on any number of devices which do not necessarily have automatic syncing of bookmarked web pages enabled. There is also a related issue of sharing, as well as discovering relevant online resources [1]. A potential solution to these problems is a social bookmarking application, where users can create a centralised repository of bookmarked resources, share them with other users and view what others are bookmarking. The popularity of these systems has become of increased interest to information architects and has prompted a number of studies into the use social bookmarking and the related field of collaborative tagging [2],[3],[4]. Results from these studies suggest that tagging and bookmarking share

similar features to more traditional indexing systems [2] but also contain extra dimensions such as tags related to time e.g. “toread” and task or users’ emotional responses to a document e.g. “cool”, which conventional indexing systems do not support [2].

One of the most popular online resources that supports social bookmarking is “delicious”, a social bookmarking website which “allows users to tag, save, manage and share web pages from a centralized source” [5]. The main advantage of these services over traditional bookmarking systems such as those included in the majority of web browsers e.g. “favorites”, is that the need for synchronizing bookmarks between multiple computers and browsers is no longer required. This is a common problem for students in HE, in particular those from health and social care courses who are on placements and who access the web on any number of machines and devices, in any number of locations. Although a number of browsers have bookmark synching capabilities (either in-built or via plugins), this is often not available on HE or NHS maintained machines due to security or privacy concerns. A related benefit is the ability to categorise resources in multiple categories, by the use of tags, which is only partially supported by web browsers in the form of “bookmark folders”. A consequence of the use of tagging that is potentially useful for students, lecturers and practitioners is the ability to discover and share resources via these tags. For example tagging resources using the course codes of a degree program gives students a central access point (e.g. <http://delicious.com/tag/NURS1297>) for all relevant online information for that course. If all students then use the appropriate course code tag when bookmarking during their own personal research, both students and lecturers will collaboratively produce a list of online resources for that course.

This paper describes the results of a project that aimed to utilise the delicious website to investigate the use of social bookmarking and tagging in an educational setting by students and lecturers.

2 Research Questions

The main objectives for this project were to investigate the best methods for integrating social bookmarking into everyday practice for both students and teachers, and to then determine behavioral usage and motivations. The following were the proposed research questions for the project:

1. What patterns of user tagging activity emerge through analyses of tagging frequency and co-word analysis? (based on [2])
2. What patterns of user bookmarking activity emerge through analyses of the resources bookmarked and the tags used to bookmark them?
3. What is the temporal distribution of bookmarking during an academic semester?
4. What types of tags are being used i.e. do students/lecturers utilise task and time related tags?

5. What are the levels of use of social bookmarking in relation to resource discovery i.e. do students browse/follow fellow students bookmarks and tags to discover resources?
6. What are students and lecturers perceptions of the advantages and disadvantages of social bookmarking and tagging?
7. What are the motivations for using social bookmarking services?
8. What features are currently missing from social bookmarking websites?

The following section describes the two methods that were used to answer the research questions described above.

3 Methodology

3.1 Participants

In total 160 people, comprising 5 lecturers and 155 students on 5 modules across the Schools of Computing and Mathematical Sciences, Engineering and Health and Social Care, participated in this study. Courses ranged from Masters (Level 7) to first year undergraduate (Level 4) and comprised of around 10 students to over 100. Students ranged from novice computer users to relative experts i.e. students who already had an undergraduate computing related degree. Courses ran from various points in September to the end of the semester in December 2010.

3.2 Materials

All students were provided with introductory materials, which were developed with a pilot group, and delivered through lectures, tutorials and YouTube videos. These resources introduced learners to the social bookmarking project and to the delicious social bookmarking tool.

An initial interesting finding was the use of the term “bookmarking”. During the first introductory talk to a group from the School of Health and Social Care, one of the researchers compared “social bookmarking” to “browser based bookmarking”, outlining the advantages and disadvantages of both approaches. It became apparent however that students did not understand what “browser based bookmarking” was, primarily due to the fact that the majority of students used Internet Explorer as their default web browser, which utilises the term “favorites” instead of “bookmarks”¹. Following this, the term favorites was used in the introductory material, alongside bookmarking, to avoid potential misunderstanding.

A further finding at this stage was the need for two different sets of instructions. A step by step guide was produced with detailed instructions and screenshots of each stage of the sign up, adding bookmarks and tagging resources process (totalling around 14 pages). Although a number of students appreciated these materials, others

¹ The term “bookmark” is used in the majority of other web browsers such as Firefox and Chrome.

suggested that they would prefer a one-page set of instructions, which highlighted the key stages, which was then produced.

3.3 Procedure

In September 2010, all lecturers introduced social bookmarking to their students using the materials described above. The only prescribed usage was that students and lecturers were instructed to tag any resources that were related to a particular module with the appropriate module code e.g. NURS1297. This then enabled a module specific set of resources to be created and made available at a single URL on the delicious website e.g. <http://delicious.com/tag/NURS1297>. Students were encouraged to use delicious to store and find useful resources for each module, with lecturers employing a number of strategies to motivate continued use. These included the production of a number of different visualisations via the tool described in section 3.3.1, as shown in figure 1 below which were shown to the students during lectures and tutorials to demonstrate the course's current bookmarking activity.



Fig. 1. Visualisation of bookmarking activity in the form of a tag cloud from NURS1297, a course entitled “Principles of Learning Disability Nursing across the lifespan”

3.4 Automated Collection and Analysis Tool

One particular advantage for researchers in this field is that in addition to a user facing service, delicious also enables programmatic access to the information stored on the site. This is enabled by an Application Programming Interface (API), which supplies a number of XML/JSON based web services. This means that the collection and analysis of users' bookmarking and tagging behaviour can then be automated.

As described by Kipp and Campbell [2], the basic component of delicious is the bookmark entry made by each user upon encountering a website of interest. In addition to the URL of the website, the user can enter a title, some notes and a number of tags. All of these details along with the username and the date the bookmark was added can be accessed via the delicious API. An analysis package was therefore created with PHP and MySQL that stored any bookmarks that were tagged with the relevant module codes, along with the username, the additional tags, any notes and the timestamp.

This analysis package produces descriptive and visualised statistics, co-word matrices and frequency data for all bookmarks created as part of this study, as well as for individual courses.

Descriptive statistics included: the total number of bookmarks created by users; the most popular bookmarked resources and websites; the total number of tags and the total number of unique tags.

Visualisations included: annotated time line charts showing the temporal distribution of bookmarks; tag clouds highlighting the most frequently used tags; geographical location of bookmarked resources (estimated by the reported location of the server of the site that had been bookmarked).

4 Results

4.1 Descriptive Statistics

160 users across the 3 Schools created 1,430 bookmarks with 5,032 tags from August 2010 to the end of January 2011. A certain amount of agreement between respondents was demonstrated with only 1,069 unique tags being used (21%) and the 1,430 bookmarks being comprised of 882 distinct url's. 58% of bookmarks (829) contained notes about the resource that had been added by the user. The figure below shows the normalised temporal distribution of bookmarking activity for the duration of the project. There were initial peaks of activity in August, with the pilot group and then in September when social bookmarking was introduced to students in all of the modules. In the first full week of the semester (w/c 27th September) 468 bookmarks were created.




Fig. 2. Normalised temporal distribution of bookmarking activity

Activity gradually declined from the initial launch in September with 488 bookmarks being created in October, 180 in November and 71 in December (although this includes the Christmas holiday). Although all of the modules stopped in December, the delicious website was still being used by students with 42 bookmarks being created in January 2011.

The types of tags used were dependent on the domain of the module but generally give a good descriptive overview of the subject area of that discipline. For example,

when considering the co-word matrices for each tag, the tag “NURS1297”, relating to the module “Principles of Learning Disability Nursing across the lifespan”, co-occurred most frequently with the tag “learningdisabilities” (40 times). Within that module, other frequently co-occurring tags were “communicationdifficulties”, “intellectualdisability” and “learningdisabilities”.

The number of tags used per bookmark ranged from 1 to 17, the distribution being shown in Figure 3 below. The majority of bookmarks (266) were tagged with a single tag, with the average number of tags per bookmark being 3.5.

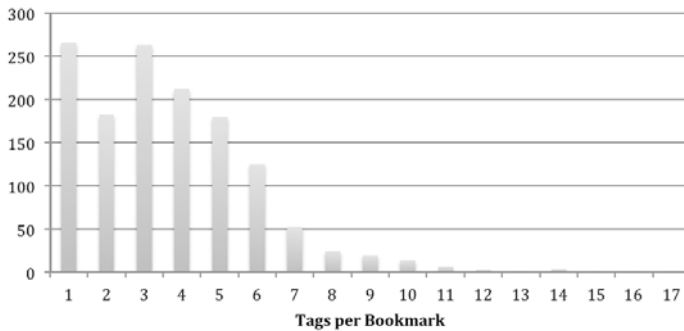


Fig. 3. Total number of Tags per Bookmark

The following table shows numbers of users with how many bookmarks in total they created e.g. 2 users created between 50 and 59 bookmarks.

Table 1. Number of bookmarks created by users

Number of Bookmarks	Number of Users
140-149	1
50-59	2
30-39	1
20-29	12
10-19	36
1-9	108

Although the majority of users created between 1 and 9 bookmarks, over 59% of users (94) created 5 or more tagged resources during the duration of the project.

Analysing the health related courses separately revealed similar activity levels and behaviours to the results as a whole, across the 3 Schools.

4.2 Survey Results

At the end of each module in December 2010, students were asked to complete an online survey to measure the impact on learning and teaching. 81% of respondents said that they used delicious to bookmark web pages as good resources with 46%

saying they used delicious to share resources. Around 30% indicated that they used delicious to find relevant resources. Interestingly only 49% found resources via the module code tag with 77% finding resources via tags related to the module i.e. subject related tags. Around 70% of students viewed other students' bookmarked resources.

84% of respondents stated that they would use the delicious website again, with 89% of those students saying they would use it for University related activities. Pleasingly, 70% of respondents said that they would recommend delicious to a friend and 68% said they'd recommend it to other learners. 52% would recommend it to professional colleagues. Although not a main objective of this project, 49% of students felt that using delicious had improved their ICT skills.

The qualitative feedback generated from the questionnaire was generally positive, relating to both storage and discovery of information.

5 Discussion and Conclusion

The results from this study indicate that social bookmarking has a number of positive outcomes with regards to teaching and learning, across a number of disciplines related to eHealth. Students and lecturers found the tool to be useful for storing, sharing and discovering resources. Through the process of using specific course code tags, both students and lecturers have created their own learning communities or communities of practice. The formation of these learning communities enables them to share information relating to their own specific course subjects, with their peers. The year on year building of a repository of information in effect provides students with an online searchable database for its members to access [6]. It also provides students the opportunity to discover and share the views and perspectives of their fellow members.

In this study, students were not limited to accessing their own learning communities; there were opportunities for students to create and join other learning communities. This was achieved by users utilising their own specific descriptive words when tagging (as "module codes" were chosen in the project to establish learning communities), or by joining existing ones. The latter can be achieved by the students finding others who use identical tags as themselves, relating to the students' own interest. This means individuals can share resources with other 'like minded' people who have similar, or the same interests [7]. Through the tagging process students can discover additional resources that they may not have necessarily found themselves (77% of students reported that they found resources in this way), thus leading to a group of people forming their own learning community [6].

The advantages of learning communities include individual users being able to access these groups at a time and place which is convenient to them, on a 24/7 basis. In this study, students were not restricted as to when and where they could access and use these learning communities, which became increasingly important when students were on placements. Additionally, through the process of using "notes", students can share views with one another about the various resources users have bookmarked, helping students to develop their critical evaluation skills.

It is noted however, the level and extent of participation as a member within each of these learning communities is dependent on the individual. Some students chose only to

participate a few times, whilst others used their communities more often. Users do not receive any feedback about whether their tagged resources have been used by others [6], perhaps explaining why there are some who chose only to consume the resources rather than contribute to the community. Although bookmarking activity decreased during the duration of the project, a key indicator of success is the building of the repository itself, as opposed to the number of contributors. The work of Ortega et al. [8] suggests that in collaborative resource creation applications such as Wikipedia, there is a great level of inequality “with less than 10% of the total number of authors being responsible for more than the 90% of the total number of contributions”. From this study, supposing that creating over 10 tagged bookmarks is a reasonable level of contribution, then 33% of users achieved this level. If we reduce that level to 5 tagged bookmarks then 59% of users have made an active contribution.

The utilisation of social bookmarking is an example of how learning communities can be created where its members can store and share a collective range of resources for others to share. This could be seen as encouraging not only the development of learner independence and autonomy but a range of related graduate attributes which are valued by employers. Further to this, the emergence of these communities demonstrates how social interaction within health care subjects in HE is progressing.

Future work will include identifying specific improvements that could be made to the delicious website and functionality that is currently missing e.g. improved support for critical evaluation of resources. The project will be continued in the following academic year with other modules as well as an investigation into how social bookmarking can be used within VLE's such as Moodle.

References

1. Jenson, J.: It's the information age, so where's the information? Why our students can't find it and what we can do to help. *College Teaching* 52(3), 107–112 (2004)
2. Kipp, M.E.I., Campbell, G.D.: Patterns and Inconsistencies in Collaborative Tagging Systems: An Examination of Tagging Practices. In: Proceedings of the 2006 Annual Meeting of the American Society for Information Science and Technology, Austin, November 3-8 (2006)
3. Kipp, Margaret E.I.: @toread and Cool: Tagging for Time, Task and Emotion. In: Proc. Information Architecture Summit 2007 (2007)
4. Golder, S., Huberman, B.A.: The structure of collaborative tagging systems. HP Labs Technical Report (2006), <http://www.hpl.hp.com/research/idl/papers/tags/>
5. delicious.com (2011), <http://delicious.com/about>
6. Benbunan-Fich, R., Koufaris, M.: An empirical examination of the sustainability of social bookmarking website. *Information Systems E-Business Management* 8, 1310148 (2010)
7. Mason, R., Rennie, F.: *E-Learning & Social Networking Handbook: The resources for Higher Education*. Taylor & Francis (2008)
8. Ortega, F., Gonzalez-Barahona, J.M., Robles, G.: On the Inequality of Contributions to Wikipedia. In: Proceedings of the 41st Annual International Conference on Hawaii International Conference on System Sciences, p. 304 (2008)

Engagement in Online Medical Communities of Practice in Healthcare: Analysis of Messages and Social Networks

David Fowler¹, Martin Szomszor¹, Simon Hammond¹,
John Lawrenson², and Patty Kostkova¹

¹ City eHealth Research Centre, City University, London, UK
² Optometry and Visual Science, City University, London, UK
David.Fowler.1@city.ac.uk

Abstract. We report on preliminary analysis on user engagement in two online medical communities of practice. Despite the communities being independent of each other, and dealing with different domains (field epidemiology and therapeutic prescribing for optometrists), there are some clear similarities in the networks of users, and in patterns of replies to user postings. We also draw some initial conclusions to help maintain user engagement in these and similar sites, and we suggest some future lines of research.

1 Introduction

Professional communities of practice (CoPs) have been the cornerstone for sharing scientific knowledge and professional discourse. The internet has dramatically changed the way communication and peer networking is managed: little overhead and flat structures, easy online recording of scientific discussions, higher frequency of postings, and virtually unlimited geographical coverage of the CoPs. However, online communities may be vulnerable to stagnation and failure if the support tools are not suitable, or if key members of the community are not able to take an active role.

1.1 Communities of Practice

The term “community of practice” has many definitions, although it originates in the work of Lave and Wenger [1]. We will use the definition of CoPs as “groups of people who share a concern, a set of problems, or a passion about a topic, and who deepen their knowledge and expertise in this area by interacting on an ongoing basis.” [2] (also cited in [3]). CoPs may be deliberately created or spontaneously emerge, and be highly structured or informal [3].

1.2 The Two Communities: FEM Wiki and MSU

We examine two independent CoPs developed around medical scientific internet portals: FEM Wiki (<http://www.femwiki.com>), dealing with field epidemiology, and Medicines Support Unit for Optometrists (MSU, <http://www.med-support.org.uk>),

supporting therapeutic prescribing by optometrists. The user bases are geographically dispersed (mainly throughout the UK for MSU and throughout Europe for FEM Wiki). Both sites provide centrally authored information to specialists, and have means for user discussion. Each was created to order, but FEM Wiki is more highly structured than MSU. In FEM Wiki, users can directly edit the content, but to guarantee quality, changes must be approved before the changes are made official. In MSU, changes can be suggested informally via the forum.

2 Social Network Analysis

We collected the messages that were posted on the discussion forums of the communities, and extracted networks of users. Each node corresponds to a user, with arcs linking the nodes of users who were involved in the same discussion (Fig. 1).




Fig. 1. The user networks extracted from FEM Wiki (*left*), and MSU (*right*). The nodes with most connections are highlighted.

Fig. 2 shows the number of connections for each node in the networks, and Table 1 summarises some key statistics. Each network has a number of users who are involved in many discussions; these seem to be mainly senior project leaders or administrators. There is an almost linear decline to users who were only involved in one or two discussions (possibly they only had a specific question that was answered to their satisfaction). Although the networks that are extracted are not a complete picture of the knowledge sharing activities in the communities (for example, members may share knowledge in person or via other media and the network does not measure the quality of contributions), it may give a reasonable approximation. Users with many connections are involved in many discussions, and therefore may have more knowledge and experience to share.




Fig. 2. The numbers of neighbours for each node in Fig. 1

Table 1. Summary statistics for the two communities

	FEMWiki	MSU
Nodes	23	20
Edges	73	62
Average Degree	6.348	6.200
Diameter	3	4
Average Path Length	1.798	1.816
Graph Density	0.289	0.326

3 Message Analysis

In addition to examining the connections between users, we also looked at the characteristics of the messages. There were striking similarities between FEM Wiki and MSU in the distribution of replies to messages. Table 2 shows that the majority of posts have a small number of replies (the median is 2 for both communities). This seems to be typical behaviour for online forums, e.g. [4].

Table 2. Summary statistics for the numbers of replies to FEM Wiki and MSU posts

	FEM Wiki (33 posts)	MSU (45 posts)
Min	0	0
Q1	1	1
Median	2	2
Q3	3	3
Max	18	16

4 Discussion and Future Work

We have shown that there are underlying similarities in the user network structure and distribution of numbers of replies to posts of two independent online CoPs. The two sites also vary in their organisational structures and editing processes, so these results might suggest some properties that are shared more widely between online CoPs. This should provide some useful lines of enquiry, although it will require access to data from a larger number of online CoPs. We will also need to investigate how the properties of CoPs vary with size, as our examples were both in the small to medium range.

The type of analysis in this paper may be helpful in identifying users whose contributions are critical to keeping an online community active. If such users become less active (for example, through pressures of other work), there is a risk that the community will stagnate, and lose other users. There is some evidence that this has happened recently with the MSU site (although with MSU there was another possible cause for loss of activity: a spam attack on the discussion forum may have driven away some users).

We are interested in tracking the activity of online CoPs over time to see how the user networks vary, investigating what factors may affect the activity, and whether there is an identifiable “critical point” at which community activity breaks down. We are currently redesigning the MSU site, and plan to promote the site again to existing and prospective users in order to increase activity.

Finally, we plan to investigate the factors that affect the user response to forum messages. Section 3 showed high level similarities, and it will be interesting to see which types of posts attract most discussion, and to draw comparisons between sites.

References

1. Lave, J., Wenger, E.: *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, Cambridge (1991)
2. Wenger, E., McDermott, R., Snyder, W.M.: *Cultivating Communities of Practice: A Guide to Managing Knowledge*. Harvard Business School Press, Boston (2002)
3. Hara, N., Shachaf, P., Stoerger, S.: Online communities of practice typology revisited. *Journal of Information Science* 35(6), 740–757 (2009)
4. Mishne, G., Glance, N.: Leave a reply: An analysis of weblog comments, In: *Third Annual Workshop on the Weblogging Ecosystem (part of WWW 2006)* (2006)

Towards Delivering Disease Support Processes for Patient Empowerment Using Mobile Virtual Communities

Bert-Jan van Beijnum, Pravin Pawar, Lamia Elloumi, and Hermie Hermens

Faculty of Electrical Engineering, Mathematics and Computer Science
University of Twente, Enschede, The Netherlands
{b.j.f.vanbeijnum,p.pawar,l.elloumi,h.j.hermens}@utwente.nl

Abstract. In existing healthcare systems, the focus is on clinical processes to assess the health condition of patients, making clinical decisions and applying therapeutic procedures all under control of health professionals. However, patients with chronic diseases usually face many disease related problems that are not immediately under control or supervision of a health professional. Taking patients with chronic cardiovascular diseases as an example, these patients are recommended to make a number of lifestyle changes: increase physical activity, change diet habits, quit smoking and adhere to a medication intake regime. In this paper we propose the use of Mobile Virtual Communities to enhance the empowerment of patients by providing the ICT mediated social support functionalities that assist patients to realize the lifestyle changes needed.

Keywords: mobile virtual communities, telemedicine, cardiovascular diseases, disease support processes, patient empowerment.

1 Introduction

In the past decade, the European Union has, and still is, investing a lot in changing the way in which health services are delivered. Telemedicine is seen as a solution for the problems in current and future healthcare delivery. Originally started with remotely monitoring of the health condition of patients with chronic diseases, the emphasis today is becoming more on the self-management of patients. A project that embraces this vision is the BraveHealth project [1]. This project targets patients with chronic cardiovascular diseases, these diseases have the highest contribution to the European mortality rate (about 2 million per year), and account for about 192 billion Euro in health expenditure and about 270 million lost working days.

This paper leverages on our previous work [2] in which the research challenges and opportunities for the mobile virtual communities (MVC) in telemedicine are described. In this paper we outline the vision for CVD patient empowerment using MVCs. E-Support Groups or virtual communities for patients exists already for some time. These virtual communities mostly focus on emotional support and informational support. In this paper, we focus on the MVCs for CVD patient empowerment including instrumental and feedback support as well.

This paper is organized as follows. Section 2 describes related work. Section 3 presents requirements for the MVC. Section 4 presents proposed concept of patient empowerment using MVC. Section 5 elaborates MVC platform architecture in order to fulfill targeted functionalities. Section 6 includes a brief discussion.

2 Related Work

Today, many dedicated virtual health communities exist, we have found over 40 different sites (English language only). *Patients Like Me* [3] is a social network for patients, the focus of this network is on giving and receiving support and information from the peers with similar health profiles. *Face To Face Health* [4] is a social network to find and connect patients with similar health experiences on a one-to-one basis. Main objective is to support storytelling and share experiences. *IBM's Patient Empowerment System* [5] is based on, and extends IBM's long tradition in the medical sector. The IBM vision is to merge the medical network with a social network in order to empower patients. *DailyStrength* [6] is a social network website where patients provide one another with emotional support by discussing their struggles and successes with each other. The site contains multiple online communities that deal with different medical conditions or life challenges. Medical professionals are also available to contact and treatments for a variety of illnesses and problems on some of the sites listed above. The virtual community sites we investigated focus on providing emotional and informational (peer) support.

3 MVC Platform Requirements

A robust architecture and technical platform is necessary to support patient mobility and empowerment needs of CVD patients. We have identified a collection of functional and non-functional requirements of such a platform based on a preliminary analysis of CVD patient needs, application scenario studies [7] and related work [8, 9, 10]. Here we only highlight some of the main functional requirements

- *Community member roles*: The platform must support user roles relevant for providing support, e.g. patient, cardiologist, nurse, practitioner, relative and friend. Also, creation of new roles and associated access rights must be supported by the platform.
- *Access*: The MVC platform must be accessible using devices that have become standard today, hence including smart phones and tablets. Human-machine interfaces must adapt according to device capabilities.
- *Language*: the MVC platform must support multiple languages.
- *Ease-of-use / usability*: The user-interface must be intuitive and easy to use. User role, skills, age, gender and any disability may affect the user-interface requirements.
- *Information access control*: Users must be able to audit who may access their health related information and they must be able to exercise control over who is allowed to see their health related information.

4 Patient Empowerment Using the Mobile Virtual Communities

The envisioned MVCs aim at accelerating disease support processes to empower CVD patients in achieving their health goals, for instance to adhere to the recommendations regarding physical activity, diet and medication [11]. Relevant types of support are known from social theory [8,9] and include informational, emotional, instrumental and appraisal / feedback support [8,9]. We see the added value of our approach especially in the domain of instrumental and feedback support from e.g. peers and aids in keeping patients motivated and ways to track their own and group performance with respect to their health goals.

The expected patient empowerment outcomes by using the prospective MVC functionalities are: being better informed, feeling more confident in the relationship with their physician, improved acceptance of the disease, feeling more confident about the treatment, enhanced self-esteem, enhanced social well-being and increased optimism and control. These outcomes are adopted from the study of patient empowerment in the online patient groups [9]. In the patient-centric view of MVCs, each community specifically addresses a particular type of CVD lifestyle facet. In each community, members take on a particular role and each role performs corresponding functions the combined result of which is focused on patient empowerment.

5 MVC Platform Architecture and Modules

Regarding the logical architecture of the MVC, we adopt the well-known and proven three-tier architectural pattern. To provide all the intended MVC disease support processes, we identified the need for the following modules:

- *MVC Platform Management*: This module is responsible to provide functionalities for managing the platform. The platform management tasks include deployment of community services, platform performance monitoring and platform auditing functionalities.
- *MVC Template Management*: The platform needs to be able to accommodate multiple communities, based on disease facets and required support type. Through this module new templates can be defined which are to be instantiated in order to create a community. The community template specifies roles and services that are to be part of communities and the rules that govern the operation and service use.
- *MVC Generics*: The MVC generics refers to the services that can be reused in different contexts. E.g. chat capabilities, the management of publication (for the purpose of information support or for educational purposes) are generic services that may appear in various communities.
- *Specific MVCs*: Based on the above modules, dedicated communities are created to address specific disease facets. These communities may be enhanced with new functionalities specific for the disease facet addressed.

6 Discussion

In this paper we have described ongoing work on delivering patient empowerment processes for CVD patients with the aim to maintain or improve patients' health status by focusing on CVD lifestyle disease facets. The disease facets include: diet, physical exercise/activity, medication compliance, and smoking cessation. The concept of patient empowerment, requirements, architecture and modules of proposed MVCs are elaborated. Presently, we are working on the implementation of MVC modules and performance metrics by which patients can objectively track their own performance and that of peer patients in their community. In addition, metrics for group performance are under development. In future phases of the project, the effectiveness in terms of patient empowerment using instrumental and feedback support will be evaluated in a clinical trial.

Acknowledgments. This study is supported by EUROPEAN UNION FP VII project number 248694 (BRAVEHEALTH).

References

1. Bravehealth: Patient Centric Approach for an Integrated Adaptive, Context Aware Remote Diagnosis and Management of Cardiovascular Diseases, <http://bravehealth.eu>
2. Beijnum, B.J.F., van Pawar, P., Dulawan, C., Hermens, H.: Mobile Virtual Communities for Telemedicine: Research Challenges and Opportunities. *International Journal of Computer Science & Applications* 6(2), 38–49 (2009)
3. Patientslikeme, <http://www.patientslikeme.com/>
4. FacetoFaceHealthCommunity, <http://www.facetofacehealth.com/>
5. Made in IBM Labs: IBM Reinvents the Patient Portal (March 2011), <http://www-03.ibm.com/press/us/en/pressrelease/33944.wss>
6. DailyStrength, <http://www.dailystrength.org/>
7. Fedele, F., Sterfanis, P.D., Ribeiro, V.: Application Scenario Studies, BraveHealth project Deliverable 1.1 (September 2010), http://bravehealth.eu/pdf/Deliverable%201.1%20Application%20Scenario%20Studies_FINAL.pdf
8. Tardy C.H.: Social Support Measurement. *American Journal of Community Psychology* 13(2) (1985)
9. van Uden-Kraan, C.F., et al.: Self-Reported Differences in Empowerment Between Lurkers and Posters in Online Patient Support Groups. *Journal of Medical Internet Research* 10(2) (2008)
10. Maloney-Krichmar, D., Preece, J.: A Multilevel Analysis of Sociability, Usability, and Community Dynamics in an Online Health Community. *ACM Transactions on Computer-Human Interaction* 12(2), 201–232 (2005)
11. Dickstein, K.: ESC Guidelines for the diagnosis and treatment of acute and chronic heart failure 2008. *European Heart Journal* 29, 2388–2442 (2008)

E-Health Readiness Assessment for E-Health Framework for Africa: A Case Study of Hospitals in South Africa

Alfred Coleman¹, Marlien E. Herselman², and Dalenca Potass³

¹Tshwane University of Technology, Soshanguve Campus, South Africa

²Living Labs, Meraka Institute, CSIR, South Africa

³School of Information and Communication Technology,

Nelson Mandela Metropolitan University, South Africa

colemana@tut.ac.za, mherselman@csir.co.za,

Dalenca.Pottas@nmmu.ac.za

Abstract. This study assessed e-healthcare readiness of rural and urban hospitals in North West Province of South Africa. Outcome of assessment led to creation of e-health architectural framework for e-health solutions. Assessment was conducted in usage of ICT in patient healthcare record system, processes and procedures in consultation among healthcare professionals, prescription of medication, referral of patients and training of healthcare professionals in ICT usage. The study was in two phases and six hospitals were selected. E-healthcare readiness assessment focusing on need, technological, engagement and social acceptance readiness were assessed. Data collected used group interviews and qualitative questionnaires. Findings showed that computers were not used for clinical duties and no e-health solutions were found. E-health Maturity Level was at level zero. Recommendations and compilation of Provincial E-Health Framework (PEHF) were made. The findings were unexpected and therefore, of great benefit to healthcare institutions which intend to implement e-health initiatives in hospitals.

Keywords: Technology readiness assessment, e-health, patient electronic health records, electronic prescription.

1 Introduction

HIV pandemic in most African countries has not only caused considerable strain on various national healthcare systems, but has increased the number of orphans, reduced productive human capital and productivity, eroded knowledge and skills, put pressure on national budgets, increased the poverty-stricken populace and reduced the quality of life, health and wellness[1]. E-health as one of the supportive systems within the healthcare system has great potential to address the challenges facing healthcare systems in developing countries. The recognition of Information and Communication Technology (ICT) in healthcare is not an end unto itself but a means to an end. The successful introduction of ICT in healthcare requires the examination of complex political, organisational and infrastructural factors, including a readiness factor [2].

The White Paper on transformation of the public healthcare system in South Africa (SA) reveals that majority of the population has inadequate access to basic healthcare services and that the greater percentage of this population lives in rural communities [3]. The White Paper also indicates that majority of South Africans receive their medical care from government-run clinics and hospitals [4]. The North West Province Health Department (NWPHD) has invested substantial sums of money in the district and provincial health services, health facilities management and healthcare support services, including ICT, in an effort to improve the work processes of healthcare professionals in order to promote efficient delivery of healthcare services. Despite the high investment in the healthcare system, many healthcare professionals in NWP do not receive the benefits for which ICT can provide to improve on their work processes.

This article reports on how ICT can be used to promote efficiency in the work processes within the hospitals to deliver quality healthcare services to the people in North West Province (NWP). To operationalise this objective, an e-health readiness assessment of rural and urban hospitals was conducted. Outcome of the assessment led to compilation of e-health framework to improve healthcare professionals' work processes. The assessment investigated effective and efficient use of ICT in patient record systems, processes in consultation, prescription, referrals and training.

This research has contributed to the body of knowledge in e-health in two ways. Firstly, a set of recommendations has been drawn from the research data to help bridge the gap between current levels of ICT usage in healthcare to attain a higher level where there is integration of e-health adoption in the hospitals. Secondly, a PEHF for the hospitals in NWP has been compiled, based on the data drawn from the research.

1.1 E-Health and E-Healthcare Readiness Assessment Theories

The World Health Organization [5] defines e-health as 'being the leveraging of ICT to connect providers, patients and governments; to educate and inform healthcare professionals, managers and consumers; to stimulate innovation in healthcare delivery and health system management and to improve our healthcare system'. Eysenbach [6] refers to e-health as 'a concerted effort undertaken by leaders in healthcare and hi-tech industries to fully harness the benefits available through convergence of the Internet and healthcare'. The advent of e-health seems fitting to address both opportunities and challenges in the healthcare sector but the question is, are hospitals in NWP ready to embrace the full potential of e-health?

Currently there are six assessment models which are commonly used to assess e-health readiness in the health environment. These models [3], [8], [7], provide different dimensions which can be utilized in assessing e-healthcare readiness. These are core readiness, engagement readiness, structural readiness and non-readiness. Another component, Technology Acceptance Model [9], is included for use in developing countries, especially in rural areas. The principal components adopted for this e-health readiness assessment included technology acceptance construct which is vital to the introduction of any technology in rural hospitals.

2 Methods

Qualitative research design employing a multiple case study approach was used. Data from selected urban and rural hospitals were collected, using group interview and qualitative questionnaire. The questionnaire instrument was used in addition to the group interview because it promoted reliability and validity of the data. The data were analyzed by utilizing [10] case study analysis template and [11] guidelines for case study analysis. The group interview and questionnaire items were formulated in accordance with the categories of background/history of hospitals, hospital infrastructure, ICT access level, including ICT availability, accessibility and usability; and e-health solutions including availability, accessibility and usability.

There are five regions in the province and all the regions were taken into account by selecting a hospital from each region. The selected hospitals were Rustenburg, Taung, Ganyesa, Klerksdorp, Christian and Reivilo Hospital. Rustenburg and Klerksdorp Hospitals are urban hospitals while Taung, Christiana and Reivilo Hospitals are rural hospitals.

2.1 Participants

A purposive sampling was applied to select respondents from the hospitals in order to achieve the goals of the study. A total of 48 respondents were selected. The selected respondents from each hospital had the following categories: 2 administrators, 2 general doctors, 2 professional nurses and 2 assistant nurses. The respondents were informed prior to the group interviews and the completion of the questionnaires that participation in the research was voluntary and that any information provided would be treated as confidential.

2.2 Research Instruments

Group Interview

A group interview was conducted in each hospital. The interviewees were of diverse age, ethnicity, gender and educational levels. The purpose of the interview was to determine how the healthcare professionals perceive the usefulness and potential benefits of e-patient record, e-prescription, e-consultation, e-referrals and e-training systems. The interview questions are attached to the full research report.

Questionnaire

The questionnaire instrument consisted of two sets of questions. The first set of questions was administered to hospital administrators whilst the second was for general doctors, nurses and assistant nurses. The purpose of the first set of questions was to establish the background history, the settings of the hospitals and the existing ICT infrastructure in these hospitals. The second set of questions was to establish baseline data for processes and procedures in keeping patient health records, consultation among healthcare professionals, prescription and referral processes. The questions were drawn from e-health readiness assessment framework [9].

3 Results

This section presents the findings and proposes some recommendations which served as guiding principles for the development of e-health solution.

3.1 Questionnaire Results

Availability of Healthcare Professionals in Rural and Urban Hospitals

The findings indicate that there is a widespread shortage of doctors in rural hospitals in NWP. An average ratio of 1 doctor to 18000 patients for rural hospitals as compared to 1 doctor to 9000 patients in urban hospitals was found. These hospitals are all state-owned healthcare institutions.

ICT Availability and Accessibility in Rural and Urban Hospitals

The findings reveal that urban hospitals have more ICT equipment than rural hospitals. Internet connection is more reliable in terms of connectivity and speed in urban hospitals. In rural hospitals, connectivity and speed of Internet services are often affected by poor telephone lines and interruption of electricity power supply. The average ratio of number of computers to doctors in rural hospitals is 1:3 while it is 1:2 in urban hospitals. However, data obtained from the e-health readiness assessment revealed that there were no computers in the doctors' consultation rooms. It was ascertained that doctors do not use computers for their clinical duties. Again nurses and assistant nurses do not use computers for their clinical work. These computers are used by hospital administrative staff for capturing patients' demographic information and revenue collection. This means that even though there are computers in the hospitals, they are used for activities which are not directly related to clinical work like e-patient health record keeping.

Healthcare professionals' use of computer was also assessed. Healthcare professionals in urban hospitals use computers/Internet daily for searching for information but not for carrying out their clinical duties. Healthcare professionals in rural hospitals use computers weekly or monthly.

3.2 Group Interview Results

E-health Solution, Availability and Accessibility in Rural and Urban Hospitals, and Perceived Usefulness in Rural and Urban Hospitals

In comparing urban and rural hospitals in terms of e-health solution availability and accessibility, the following findings emerged: There is no e-patient health record system, e-consultation, e-prescription, e-referral and e-training systems in both urban and rural hospitals in NWP. Participants indicated that installation of e-patient health record system will save time when recording patient information, diagnosing and in the treatment of patients. One participant stated "Patients who do not get their records after a long time, the e-patient record system will help to safeguard this". It became evident that e-patient record system will improve communication and eliminate patients who

collect medication from multiple hospitals with duplicate paper prescription. One respondent stated, “It will eliminate ghost patients and before a patient arrives at the pharmacy, the pharmacist will have the information about the patient”.

It emerged that referrals of patients from a lower level hospital to a higher level hospital were done through the use of paper notes. Patients are given referral letters (except in emergency cases) which they take to the referred hospital. Patients often misplace these letters and end up not going to the referred hospital. One participant indicated, “E-referral will help because you refer patients but on the way they get lost or don’t go. If it is done electronically, the referred hospital will know that such a time we expect this patient”.

There is Internet facility in both urban and rural hospitals but the Internet is limited to searching information and sending e-mails. Both urban and rural hospitals have PAAB system which is used to collect and send patient demographic information to NWPHD’s head office monthly. Urban hospitals have tele-radiography facility which is used for sending x-ray images between Klerksdorp and Rustenburg Hospitals.

4 Discussion

4.1 Availability of Healthcare Professionals in Rural and Urban Hospitals

An average ratio of 1 doctor to 18000 patients for rural hospitals as compared to 1 doctor to 9000 patients in urban hospitals was noted. Since these are state-owned healthcare institutions, the ratio of doctors to number of patients, both in rural and urban hospitals, is higher than that quoted by the [12] report. The report [12] focused on both private and state-owned hospitals in SA, while this research investigated only state-owned hospitals in NWP in SA. The shortage of doctors in rural hospitals has to drive investment into e-health technologies like e-consultation, to enable non-doctors healthcare professionals such as nurses and midwives to perform more advanced functions through synchronous consultation with professional doctors.

4.2 ICT Availability and Accessibility in Rural and Urban Hospitals

It was found that more ICT equipment is needed in both rural and urban hospitals if e-health solutions are to be implemented. Internet connection also needs upgrading.

ICT infrastructure in both urban and rural hospitals is not integrated to be able to work together within and across urban and rural hospitals. This confirms research conducted in the Eastern Cape Province of SA [13].

4.3 E-Health Solution, Availability and Accessibility in Rural and Urban Hospitals, and Perceived Usefulness in Rural and Urban Hospitals

The findings showed that the Internet facility in both rural and urban hospitals is limited to searching information and sending e-mails. Urban hospitals have tele-radiography facility which enables them to send x-ray images. These findings place both the urban and rural hospitals at Level 0 (The Baseline) on the E-health Maturity Curve as illustrated in Figure 1 below.

Source: [14]




Fig. 1. E-health Maturity Curve

Figure 1 illustrates that the level of e-health application in rural hospitals is at the Presence stage which is classified as Level 0 (The baseline). At this stage there is a non-interactive website where the main intent is to disseminate information. Thus, rural hospitals are able to receive and send information from and to NWPHD through e-mails and the PAAB system.

The level of e-health application in urban hospitals is at the Interaction stage which is also at Level 0 (The baseline). However, the Interaction stage offers services that are more advanced than the Presence stage. Thus, urban hospitals are able to send and receive x-ray images through the tele-radiography facility. Despite the availability of these ICT infrastructure facilities, the ICT systems are not integrated to be able to work together across departmental and organizational boundaries to use information that has been provided by the patient.

5 The Need for Provincial E-Health Framework (PEHF)

A major setback which became evident is that despite the availability of the ICT facilities and other ICT infrastructure in both urban and rural hospitals, the ICT systems are not integrated to work together within and across the hospitals to allow healthcare professionals to gain the benefits of e-health solutions and applications. These findings place both rural and urban hospitals at Level 0 on the E-Health Maturity Curve (cf Fig1). Therefore, it is imperative that a special e-health framework be compiled based on these findings to move the e-health application usage in these hospitals from Level 0 to Level 2 (Healthcare 2.0) on the E-Health Maturity Curve by ensuring that Level 1(Integration) is effectively and efficiently achieved. Therefore, the PEHF was compiled (cf Fig 2).




Fig. 2. PEHF Architecture

The PEHF architecture emphasizes the integration of services within local hospitals and avails these services to the web through a provincial e-health services hub. This component incorporates the business services of the hospitals and transforms them from physical and paper phase to the electronic phase through the service hub. This represents the migration from the physical normal business and paper services of the hospitals to web-based services.

6 Conclusion

Having reviewed the literature on e-health and e-healthcare assessment, conducted the e-healthcare assessment in the selected hospitals, and considering the findings derived from the hospitals’ assessment, this article provides the following conclusion:

- ICT infrastructure in both urban and rural hospitals is not integrated to work together within and across hospitals to allow healthcare professionals to gain the benefits of e-health solutions and applications
- E-health applications in the work processes in the hospitals are at the Presence and Interaction stages on E-Health Maturity Curve which is Level 0 (The Baseline).

Based on the above conclusion, the following recommendations are made:

- ICT systems within each hospital and across hospitals in NWP need to be integrated in order to facilitate e-consultation by using an integrated network which will assist healthcare professionals to consult with peers and specialists for professional advice and information; and

- Inter-operability of systems to bring together the diverse systems and data sources into a coherent, controlled environment, create a network of healthcare delivery and incorporate existing e-health applications like PAAB and tele-radiography into one e-health framework.

In conclusion, this article has provided the PEHF which will move the work processes of healthcare professionals from Level 0 to Level 2 (Healthcare 2.0) on the E-health Maturity Curve, leading to efficient and effective delivery of healthcare services to the people in NWP.

References

1. Botswana Government: Status of the National Response to the UNGASS Declaration of Commitment on HIV/AIDS. National Aids Coordinating Agency, Gaborone (2003)
2. Jennett, P., Jackson, A., Healy, T., et al.: A Study of a Rural Community's Readiness for Telehealth. *J. Telecare* 9, 259–263 (2003)
3. Department of Health (DoH). Documents and Reports.: North West Department of Health, http://www.nwhealth.gov.za/nwhealth_web/doc_reports.php
4. Rasmeni, M.: Budget Speech for the Department of Health for the Financial Year 2007/2008. Presented by the MEC for Health, at the North West Provincial Legislature, http://www.nwhealth.gov.za/nwhealth_web/dir_prog.php
5. World Health Organization (WHO):. Developing Health Management Information Systems: A Practical Guide for Developing Countries, Geneva (2004)
6. Eysenbach, G.: What is e-Health? *J. Med. Internet Research* 3(2), 20 (2001)
7. Campbell, J.D., Harris, K.D., Hodge, R.: Introducing Telemedicine Technology to Rural Physicians and Settings. *J. Fam. Pract.* 50, 419–424 (2001)
8. Demiris, G., Oliver, D.R.P., Porock, D., Courtney, K.: Home Telehealth: The Missouri Telehospice Project: Background and Next Steps. *Home Health Care Technology Report* 1(49), 55–57 (2004)
9. Ojo, S.O., Olugbara, O.O., Emuoyibofarhe, O.J.: Formal Model for e-Healthcare Readiness Assessment in Developing Country Context. In: 4th International Conference on Innovations in Information Technology, IIT 2007 (2007)
10. Creswell, J.W.: *Qualitative Inquiry and Research Design*, 2nd edn. Choosing among Five Approaches. Sage Publications, London (2007)
11. Klein, H.K., Myers, M.D.: A Set of Principles for Conducting and Evaluating Interpretive Field Studies in Information Systems. *Management Information Systems Quarterly* 23(1), 67–94 (1999)
12. United Nations Framework Convention on Climate Change (UNFCCC):. Development and Transfer of Technologies. In: Workshop on Enabling Environment for Technology Transfer, <http://www.unfccc.int/resource/docs>
13. Nkqubela, L.R., Herselman, M.E.: ICT Application as E-health Solutions in rural Healthcare in the Eastern Cape of South Africa. *Health Information Management J.* 39(1) (2010)
14. Jarvis, B., Fortunov, I., Bachour, T.: Connect Health Framework Architecture and Design Blue Print, <http://www.microsoft.com/.../articles>

The Physical Activity Loyalty Card Scheme: Development and Application of a Novel System for Incentivizing Behaviour Change

Ruth F. Hunter¹, Michael Davis², Mark A. Tully^{1,3}, and Frank Kee^{1,3}

¹ Centre for Public Health, Queen's University Belfast, Northern Ireland

² Centre for Secure Information Technologies (CSIT), Queen's University Belfast

³ UKCRC Centre of Excellence for Public Health, Queen's University Belfast

ruth.hunter@qub.ac.uk

Abstract. In the Public Health White Paper “Healthy Lives, Healthy People” (2010), the UK Government emphasised using incentives and “nudging” to encourage positive, healthy behaviour changes. However, there is little evidence that nudging is effective, in particular for increasing physical activity. We have created a platform to research the effectiveness of health-related behaviour change interventions and incentive schemes. The system consists of an outward-facing website, incorporating tools for incentivizing behaviour change, and a novel physical activity monitoring system. The monitoring system consists of the “Physical Activity Loyalty Card”, which contains a passive RFID tag, and a contactless sensor network to detect the cards. This paper describes the application of this novel web-based system to investigate the effectiveness of non-cash incentives to “nudge” adults to undertake more physical activity.

Keywords: incentives, nudging, physical activity, behaviour change, Web 2.0, RFID tags, sensor technologies

1 Introduction

The prevalence of physical inactivity, rising obesity levels and associated health conditions in children and adults is rising [1,2,6]. A continuing trend of a more sedentary lifestyle has led to stark projections of future prevalence of obesity, morbidity and mortality [1,2,12]. Recent European figures show that only 31% of adults currently meet the physical activity recommendations [16]. This level of inactivity has directly contributed to the rising rate of obesity worldwide. In the UK, most adults are already overweight and by 2050, 60% of men and 50% of women could be clinically obese, costing an extra £45.5 billion per year in treating obesity-related disease [5].

More innovative interventions are required to halt the global increase in physical inactivity and obesity by sustaining healthy lifestyle behaviours for all ages. “Nudging” and using incentives to promote positive, long-term healthy

behaviour changes is now the UK government's preferred strategy for promoting public health [15]. Nudging is defined as:

... any aspect of the choice architecture that alters people's behaviour in a predictable way without forbidding any options or significantly changing their economic incentives. To count as a mere nudge, the intervention must be easy and cheap to avoid. Nudges are not mandates. Putting the fruit at eye level counts as a nudge. Banning junk food does not. [19]

Nudging is not new and draws on behavioural economics and social psychology to explain why people may behave in ways that deviate from rationality, invoking various mechanisms of behaviour change including the changing of default options, message framing and the provision of social norm feedback. Given the vicious cycles that confound our attempts to change our behaviour [18], it is not surprising that a recent framework has underlined the value of combining approaches that can simultaneously affect our capability, opportunity and motivation to change behaviour [14].

1.1 Previous Research

While there is growing evidence of effectiveness of these approaches for some behaviours, the evidence base for nudging to promote physical activity is sparse. Certainly, more behavioural economists are beginning to research the role of incentives for physical activity, and both financial [3] and non-financial [9] incentives have been shown to increase levels of physical activity, at least in the short term.

The Public Health White Paper [15] highlighted a number of nudge-type initiatives, such as the Step2Get scheme, developed by Intelligent Health Ltd.¹ Based on individual incentives and innovative technology to monitor physical activity levels, students were provided with a swipe card which they touched on receivers placed on lampposts along a safe walking route to school. Each completed walk to school was converted to points and these were redeemed as rewards (*e.g.*, cinema vouchers) at the end of the week. Post-intervention results showed an 18% modal shift towards walking to school. Unfortunately, there are no follow-up results available to determine if the large extrinsic incentive (cinema vouchers) had led to a long-term intrinsic behaviour change (increased physical activity).

This paper describes the development and application of a novel web-based system for incentivizing behaviour change using innovative technology. The paper will also outline the study design and preliminary findings of a randomized controlled trial investigating the effectiveness of incentives (redeemable vouchers for retail outlets) as a nudge to encourage adults to be more physically active.

¹ <http://www.intelligenthealth.co.uk/>



(a) Swiping the PAL Card at a sensor in the park

(b) The PAL Scheme website

Fig. 1. Elements of the PAL Scheme

2 The Web-Based System

We have developed a web-based system in conjunction with The Physical Activity Loyalty Card Scheme, based on guidance from previous research [7,8,11,17,20]. The primary purpose of this system is to increase physical activity and promote a healthy lifestyle. The system incorporates a number of behaviour change tools and nudge elements, which are described in more detail in this section.

2.1 Incentives and Nudge Elements

The PAL Scheme system integrates an innovative physical activity tracking system (Intelligent Health Ltd.) with web-based monitoring and nudge elements (Fig. 1). The tracking system uses Near-Field Communication Technology (NFC) and a “Physical Activity Loyalty Card” (PAL Card) which contains a passive Radio Frequency Identification (RFID) tag. This is a contactless proximity technology: when users present their card within 10 cm of the NFC sensor, a transaction is logged. As the sensors are placed in the outdoor environment, they are housed within a wooden post. The post also contains a mobile phone, which sends the transaction (Card ID, Sensor ID, and Timestamp) via SMS to a data centre. In the back-end system, the transactions are assembled into paths (see Fig. 3a for a visualisation of the network paths). As we know the elapsed time between sensor swipes and the distance between sensors, we can calculate distance and average walking speed for each session. We also know the height and weight of each user: these can be used to estimate the number of calories burned.

The system has been set up to give users real-time feedback on various aspects of their recent bout of physical activity, including minutes completed, distance

covered and calories expended (Fig. 1b). Each user has a personal account where they can view their own physical activity data. Feedback is presented graphically to show users' daily and aggregated physical activity parameters. Personal goals can be set using this feature and shown on the graphical display. This system can be used to incorporate some nudge elements—for example, minutes of activity are converted to points and redeemed for various rewards.

We incorporated a feature that enabled automated messages to be tailored to the individual users' weekly physical activity level. For example, if a user had fallen below the recommended weekly physical activity level, a prompt was sent to encourage them to undertake more activity with links to some of the behaviour change tools featured on the website to help them.

2.2 Behaviour Change Tools

The website contains a number of interactive tools offering personal advice, support and aids for behaviour change. Two tools have been adapted to offer individually tailored advice to the participant in order to increase their physical activity levels. These include the Stages of Change Questionnaire [13] and an adapted questionnaire for identifying and overcoming barriers to becoming more physically active. In addition, the website contains tools for planning physical activity opportunities as part of their working week.

2.3 Other Features

The website contains health promotion material for leading a healthy lifestyle, including advice on physical activity, diet and smoking. We used Facebook and Twitter to disseminate health promotion messages and regular encouragements to keep participants engaged with the scheme. Additionally, the website incorporated a number of Web 2.0 features (e.g. forums, user comments) to receive feedback on the system. Participants were also able to use the social networking features to provide social support—for example, planning walks with other users.

Finally, the web-based system also acts as a comprehensive research tool, incorporating features to support each stage from recruitment to data analysis. This includes electronic data collection, processing and aggregation.

3 Application of the System

3.1 Trial Design

We have built upon the non-randomized scheme highlighted in Sect. 1.1 and designed an intervention with several nudge components (including modest incentives for physical activity participation). We also draw upon evidence-based approaches from the behavioural science literature, including self-monitoring, the provision of personal feedback and goal-setting resources. The aim of the intervention is to provide an extrinsic incentive to nudge individuals to develop a

long-term intrinsic behaviour change. The system has been piloted in a 12-week intervention with civil servants who work at Stormont Estate, Belfast, Northern Ireland. Based on a sample size calculation, we recruited 406 employees aged 18–65 years old, who work on the Estate at least four days a week (minimum of 6 hours per day) to participate in a randomized controlled trial investigating the effectiveness of incentives for encouraging physical activity in adults. Participants were recruited via email, posters and flyers distributed around the workplace. Participants were randomly allocated to one of two groups:

1. Participants received incentives for being physically active. Participants collected points for each minute of physical activity that they completed by swiping their “loyalty card” across sensors placed around the Estate. Points could be earned by walking/running around the Estate, and by attending the gym/exercise classes. Points were reimbursed for rewards at the end of week 6 and week 12. Participants could earn a maximum of 30 points per day, five days per week (equivalent to the physical activity recommendations). The rewards were various redeemable vouchers from a number of local businesses. Participants received feedback regarding minutes of activity, points earned, distance covered and calories burned.
2. Participants did not receive incentives for being physically active. Participants used their “loyalty card” to monitor their physical activity levels and received feedback regarding time, distance, and calories. They were not able to collect points or earn rewards.

3.2 Preliminary Results

There were 406 participants in the trial: 67% were female and the mean age was 43.2 years. Fig. 2a shows the number of unique PAL Card users during the 12-week intervention. There were 250 users in the first week. This gradually declined from week 4, with approximately 100 users during the final week.

Fig. 2b shows the number of times each participant used their PAL Card per week. Most users undertook activity (used their PAL Card) between 1 and 5 times per week. This data will be further analysed to investigate the effectiveness of incentives for increasing physical activity and the predictors of card usage.

Fig. 3 represents the PAL Card transactions as network graphs. In Fig. 3a, the vertices represent RFID sensors and edges represent patterns of mobility. Sensors R1–R5 were placed along paths in the outdoor environment; there were two additional sensors on the door to the Gym and the Fitness Studio. Edge weights are proportional to the frequency that each path segment is traversed. This allows us to analyse the mobility patterns of the participants in the study.

Fig. 3b shows the social connections between participants in the study. We used a method similar to [10] to infer social connections based on users who presented their cards at the same sensor within a few seconds of each other. If two users present their cards in the same place at the same time, this suggests a possible social connection. If this happens frequently for a given pair of users, a connection is almost certainly present. The graph shows that certain users

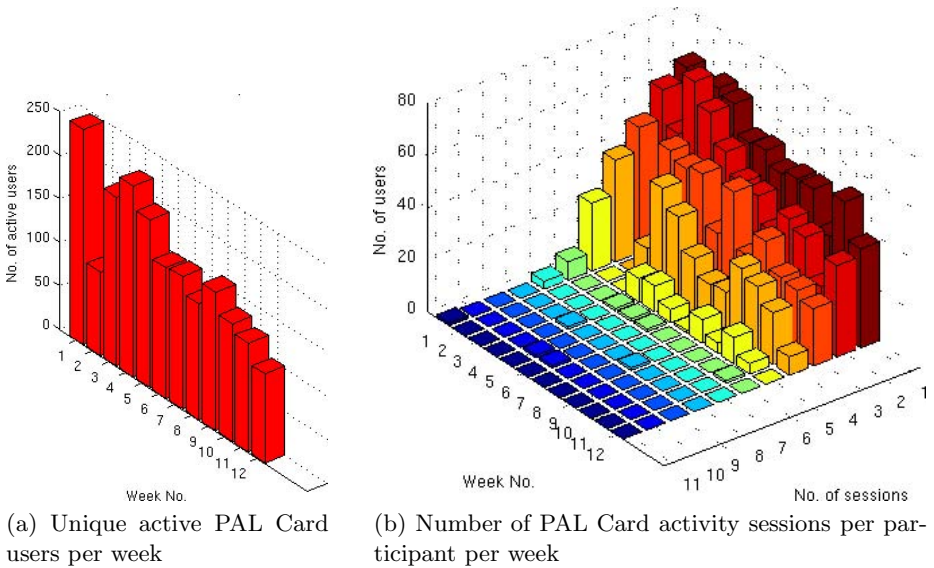


Fig. 2. PAL Card activity during the 12-week intervention

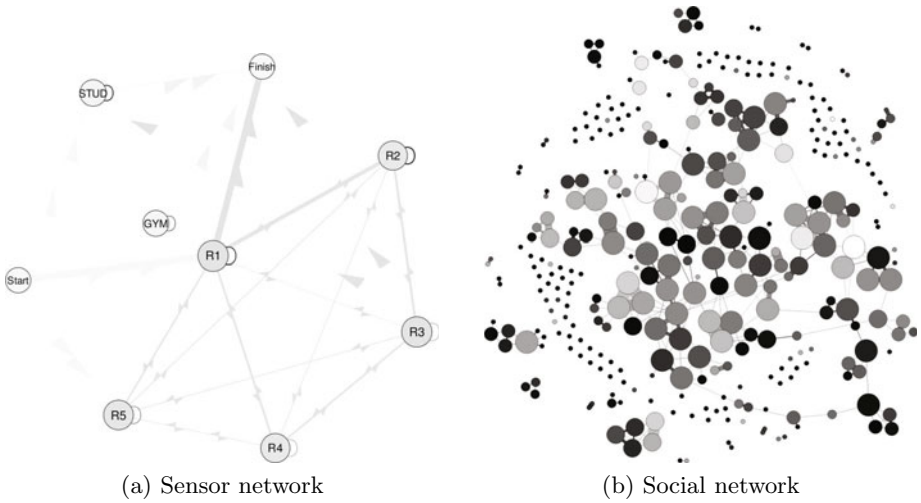


Fig. 3. PAL Card transactions represented as network graphs

formed clear clusters, whereas others preferred to exercise on their own. The size of each node indicates the degree of connectedness; the shade represents the level of physical activity (darker shades indicate less activity and lighter shades indicate more activity). We propose to use this kind of graph to investigate whether physical activity behaviour changes percolate through social networks. Beyond visual analytics, we intend to investigate the formal properties of the graph representations using graph-mining techniques such as [4].

4 Future Applications

There is scope for this web-based system to be used to investigate the effectiveness of nudge interventions including the influence of social norms, competition, and other lessons learnt from the behavioural economics literature. This innovative technology can be applied to future schemes and research trials investigating the use of incentives and nudges to encourage positive lifestyle behaviour changes in terms of, for example, smoking and diet, and in various settings—for example, school and workplace.

5 Conclusions

In this paper, we have described a sophisticated system which has been used to investigate the effectiveness of incentives for encouraging adults to be more physically active. Preliminary results indicate that this is a useful tool for this purpose. The data will be further analysed to investigate the effectiveness of incentives, the predictors of users of the system and the influence of peer relationships in increasing physical activity in adults. We plan to develop a set of software tools for the development and deployment of future nudge and behaviour change interventions that can be used by other researchers and public health practitioners.

Acknowledgements. We would like to acknowledge funding from the National Prevention Research Initiative (NPRI) and their funding partners (Alzheimer’s Research Trust; Alzheimer’s Society; Biotechnology and Biological Sciences Research Council; British Heart Foundation; Cancer Research UK; Chief Scientist Office, Scottish Government Health Directorate; Department of Health; Diabetes UK; Economic and Social Research Council; Engineering and Physical Sciences Research Council; Health and Social Care Research and Development Division of the Public Health Agency (HSC R&D Division); Medical Research Council; The Stroke Association; Welsh Assembly Government; and World Cancer Research Fund (<http://www.npri.org.uk>). We also wish to acknowledge funding from the Department for Employment and Learning, Northern Ireland and the support of Intelligent Health Ltd., London.

References

1. White house task force report on childhood obesity: Report to the president. solving the problem of childhood obesity within a generation (May 2010), <http://www.letsmove.gov/white-house-task-force-childhood-obesity-report-president>
2. Branca, F., Nikogosian, H., Lobstein, T.: The challenge of obesity in the WHO European region and the strategies for response. Technical report, WHO Regional Office for Europe (2007)
3. Charness, G., Gneezy, U.: Incentives to exercise. *Econometrica* 77(3), 909–931 (2009)

4. Davis, M., Liu, W., Miller, P., Redpath, G.: Detecting anomalies in graphs with numeric labels. In: 20th ACM Conference on Information and Knowledge Management (CIKM 2011), Glasgow (October 2011)
5. DHSSPS. A fitter future for all: An obesity prevention framework for Northern Ireland 2011–2021 (draft), <http://www.dhsspsni.gov.uk/obesity-prevention-framework-2011-2021.pdf> (accessed September 13, 2011)
6. Farhat, T., Iannotti, R.J., Simons-Morton, B.G.: Overweight, obesity, youth, and health-risk behaviors. *American Journal of Preventive Medicine* 38(3), 258–267 (2010)
7. Ferney, S.L., Marshall, A.L.: Website physical activity interventions: preferences of potential users. *Health Education Research* 21(4), 560–566 (2006)
8. Ferney, S.L., Marshall, A.L., Eakin, E.G., Owen, N.: Randomized trial of a neighborhood environment-focused physical activity website intervention. *Preventive Medicine* 48(2), 144–150 (2009)
9. Johannesson, M., Östling, R., Ranehill, E.: The effect of competition on physical activity: A randomized trial. *The B.E. Journal of Economic Analysis & Policy* 10(1), 91 (2010)
10. Lauw, H.W., Lim, E.-P., Pang, H., Tan, T.-T.: Social network discovery by mining spatio-temporal events. *Comput. Math. Organ. Theory* 11, 97–118 (2005)
11. Leslie, E., Marshall, A.L., Owen, N., Bauman, A.: Engagement and retention of participants in a physical activity website. *Preventive Medicine* 40(1), 54–59 (2005)
12. Lieberman, D.A., Chamberlain, B., Medina, E., Franklin, B.A., Sanner, B.M., Vafidis, D.K.: The power of play. In: *Innovations in Getting Active Summit 2011*, vol. 123, pp. 2507–2516 (2011)
13. Marcus, B.H., Selby, V.C., Niaura, R.S., Rossi, J.S.: Self-efficacy and the stages of exercise behavior change. *Research Quarterly for Exercise and Sport* 63(1), 60–66 (1992)
14. Michie, S., van Stralen, M., West, R.: The behaviour change wheel: a new method for characterising and designing behaviour change interventions. *Implementation Science* 6(42) (April 2011)
15. Department of Health. *Healthy lives, healthy people: our strategy for public health in England* (2010), http://dera.ioe.ac.uk/1974/1/dh_122347.pdf
16. World Health Organization. 10 key facts on physical activity in the WHO European region, <http://tinyurl.com/10factsPA-EU> (accessed September 13, 2011)
17. Pressler, A., Knebel, U., Esch, S., Kölbl, D., Esefeld, K., Scherr, J., Haller, B., Schmidt-Trucksäss, A., Krčmar, H., Halle, M., Leimeister, J.M.M.: An internet-delivered exercise intervention for workplace health promotion in overweight sedentary employees: a randomized trial. *Preventive Medicine* 51(3-4), 234–239 (2010)
18. Swinburn, B., Egger, G.: The runaway weight gain train: too many accelerators, not enough brakes. *BMJ* 329(7468), 736–739 (2004)
19. Thaler, R.H., Sunstein, C.R.: *Nudge: Improving decisions about health, wealth and happiness*. Penguin, London (2008)
20. Vandelanotte, C., Spathonis, K.M., Eakin, E.G., Owen, N.: Website-delivered physical activity interventions. *American Journal of Preventive Medicine* 33(1), 54–64 (2007)

Tracking Media Reports on the Shiga Toxin-Producing *Escherichia coli* O104: H4 Outbreak in Germany

Jens P. Linge¹, Jas Mantero², Flavio Fuart¹, Jenya Belyaeva¹,
Martin Atkinson¹, and Erik van der Goot¹

¹ Joint Research Centre of the European Commission,

Institute for the Protection and Security of the Citizen, Ispra, Italy

² European Centre for Disease Control and Prevention (ECDC), Stockholm, Sweden

{jens.linge, flavio.fuart, martin.atkinson,
erik.van-der-goot}@jrc.ec.europa.eu
jas.mantero@ecdc.europa.eu
jenya.belyaeva@ext.jrc.ec.europa.eu

Abstract. In May 2011, an outbreak of enterohemorrhagic *Escherichia coli* (EHEC) occurred in northern Germany. The Shiga toxin-producing strain O104:H4 infected several thousand people, frequently leading to haemolytic uremic syndrome (HUS) and gastroenteritis (GI). First reports about the outbreak appeared in the German media on Saturday 21st of May 2011; the media attention rose to high levels in the following two weeks, with up to 2000 articles categorized per day by the automatic threat detection system MedISys (Medical Information System). In this article, we illustrate how MedISys detected the sudden increase in reporting on *E. coli* on 21st of May and how automatic analysis of the reporting provided epidemic intelligence information to follow the event. Categorization, filtering and clustering allowed identifying different aspects within the unfolding news event, analyzing general media and official sites in parallel.

Keywords: epidemic intelligence; event-based surveillance; *E. coli*; EHEC; MedISys.

1 Introduction

Internet surveillance systems have increasingly been used for early event detection and alerting of emerging public health threats [1]. The Medical Information System (MedISys, <http://medisys.newsbrief.eu>) is a fully automatic public health surveillance system to monitor reporting on emerging public health threats such as human and animal infectious diseases, chemical, biological, radiological and nuclear (CBRN) threats, food & feed contaminations and plant diseases [2]. The system retrieves news articles from the internet, categorizes all incoming articles according to pre-defined multilingual categories, identifies entities such as organizations, persons and locations, clusters articles and calculates statistics to detect emerging threats. Users can screen the categorized articles and display world maps highlighting event locations together with

statistics on the reporting of health threats, countries and combinations thereof. Articles can be further filtered by language, news source, and country.

Articles are classified in a category, if they satisfy the category definition which may comprise Boolean operators, proximity operators and wild cards. Cumulative positive or negative weights can be used with an adjustable threshold. For *E. coli*, a category was introduced in MedISys in 2008 with expert input from DG SANCO and ECDC [3]. It consists of a set of patterns covering several languages. Table 1 shows the definition (as of 27 June 2011), with the weights for each pattern indicated in the second column.

MedISys monitors the volume of news per category and per country to determine sudden changes from the 14-day average number of articles in any given country-category combination (e.g. Germany - *E.coli*). If the number of articles in the last 24 hours for a country-category combination (normalized by weekday fluctuations) is significantly higher than the 14-day average, users are notified using ranking graphs on the web site or email notifications. The statistics are language-independent because of the multi-lingual patterns in the category definitions. This allows users to detect any change in a category even before the event is reported in their own language.

MedISys also clusters all news items within a time window of four hours (or more, depending on the number of recent articles) and presents the largest clusters as Top Stories.

2 Media Reports on the Outbreak

On the 19th of May, the Robert-Koch-Institute was informed about a cluster of three cases of HUS in Hamburg [4]. MedISys detected the first media report in the German newspaper Die Welt on Saturday 21st of May at 12:14 CEST (Central European Summer Time). Several other articles followed the developing story in the afternoon (see Table 2), making reference to various press releases by public health authorities in Germany, e.g. in Hamburg and Lower Saxony (in MedISys, users can distinguish between general news media and official sources). Altogether, 23 German news items triggered the MedISys *E. coli* category, of which 22 were about the outbreak (one irrelevant article was about the water quality in German lakes in the region Oberbergischer Kreis). Furthermore, there were two news reports in Farsi and one report in Mandarin; these reports were from the science pages of Iranian and Chinese newspapers and were not related to the outbreak. Due to the sudden increase in media reports on *E.coli* on Saturday (in comparison to the average value of the last 14 days), MedISys issued an automatic alert and also highlighted *E.coli* in combination with Germany to the users (in the section called "Most active topics").

The Early Warning and Response System (EWRS) of the European Union [5] received a first communication by the German authorities on Sunday 22 May at 11:40. ProMED-mail covered the event in a report on Monday 23 May [6]. All other major early alerting systems (ARGUS, Biocaster, GPHIN, HealthMap, PULS) reported the event as well.

Table 1. Category definition (as of 27 June 2011) showing the patterns with their corresponding weights (using wild cards such as % for zero, one or more characters, + for whitespace/linebreak, and _ for one character). An incoming news article is selected, if the sum of the weights of all triggered patterns exceeds the threshold (which was set to 40 for this category); negative weights are used to exclude irrelevant news items.

Pattern	Weight
escherichia+coli+enfek%	20
koli+basili%	20
pałeczka+okreźnicy	20
koli+basili	20
eşerişiya+koli	20
бактерията+Ешерихия+коли	20
кишечн%+палочк%+ЕНЕС	20
палочк%+ЕНЕС	20
ЕНЕС	20
enterohämorrhagisches+Eschericha+Coli	20
Ehec-Infekt%	20
Ehec-Keim	20
Coliba	-99
Colidiu	-99
Eşerihija+koli	20
escherichia+coli	20
ешерихи%+кол%	20
Eşerihioze	20
大腸桿菌	20
大腸埃希氏桿菌	20
大腸杆菌	20
大腸埃希氏杆菌	20
الإنتريكية+القولونية	20
اشريش—باکلي	20
اکولى	20
اسچريچيا+کولى	20
ىکولای	20
enterokrwtoczn%	20
escherichia+coli+O157_H7	20
escherichia+coli+O104_H4	20
escheric%+coli%	20
e.coli+bacteri%	20
Ешерихия+коли	20
кишечн%+палочк%	20
e-coli%	20
Sukelta	20
Enterohemoraginès	20
Lazdelès	20
Žaminès	20
escherich%	20
Nakkus	20
enteroh%	20
Colibacille	20
Κολοβακτηρίδιο	20
e+coli%	20
Eşerihioze	20
Kolibakteeri	20




Fig. 1. Number of articles in the *E. coli* category per day (data from 1 May to 4 July 2011; all languages)

While the outbreak in northern Germany reached its peak between 21 and 23 May [4], the media reporting showed a different behaviour. As illustrated in Fig. 1, the highest number of articles per day was reached on 2 June 2011 (all languages, all countries). The data can be further filtered by language, country of publication, country mentioned in the text (using a multi-lingual category; see Fig. 2). The filter functionality clearly shows how the media attention changed geographical focus over time, following the developing situation.

In Fig. 2, we can clearly identify key aspects:

- the sudden rise on articles mentioning Germany (21-25 May), when the first cases became public,
- a peak with articles on a Swedish tourist group who got infected in Lower Saxony (26 May).
- the reporting on alleged *E. coli* contaminations in Spanish cucumber, tomatoes and salad (peak on 27 May with 107 articles),
- the reporting on the financial impact on Spanish farmers and the announcement that Spanish cucumbers had tested negative for *E. coli* (peak on 31 May with 300 articles),
- the discussion on trade restrictions in Russia for EU vegetable products (peak on 2 June),
- the announcements by German authorities that bean sprouts were the source of infection (rise in volume on 5-10 June),
- the *E. coli* cluster of cases in Bordeaux, France (peak on 16 June with 203 articles), and
- the reporting on fenugreek seeds imported from Egypt in 2009 and 2010 (from end of June on).





Fig. 2. Number of articles in the *E. coli* category from one of the countries Luxembourg (LU), Sweden (SE), France (FR), Russia (RU), Spain (ES) and Germany (DE)

Fig. 3 illustrates how MedISys presented statistics on the outbreak on its web site (screenshots taken on 26 May). It demonstrates how the deviation from the 14-day average alerts the users to the *E.coli*-Germany combination. All data used for the figures stem from MedISys.


In addition to the *E. coli* category, the outbreak was also visible in filters set up by EFSA [7], e.g. EFSABacteria (which contains *E. coli* as potential pathogen), EFSAEconomics and EFSACommerce in relation to Spain (impact on Spanish farmers) and Russia (trade restrictions) and EFSAFoodFeedSafety. This demonstrates that broader filters targeted at economics and commerce are able to detect changes in media reporting.

Using entity extraction, the main organizations and people mentioned in the articles can be identified by the system. As an example, we extracted a subset of articles on *E. coli* that also mentioned the Robert-Koch-Institute, Commissioner John Dalli, European Commission, ECDC, and EFSA. These entities were selected from an automatically generated list of top entities (according to number of citations). Fig. 4 summarizes the data, highlighting the following aspects:

- Most articles from the early phase mention the Robert-Koch-Institute,
- The European Commission is mentioned during the discussions on the alleged contaminations of Spanish cucumbers, tomatoes and salad; Commissioner John Dalli is cited in the peaks on 2 June and 7 June (in parallel to the European Commission),
- ECDC and EFSA are mentioned in many articles regarding the risk assessment in terms of public health at EU level; a strong peak in reporting on 30 June is due to press releases about scientific reports and risk assessments, e.g. the scientific report [8] lead to more media attention (375 articles mentioning EFSA on 5 July).



AlertLevel (24h)



Previous 14 days average
Alert level:
high medium low

Daily number of articles in this category



Fig. 3. Screenshot of MediSys on 26th of May (above: country distribution; centre: statistics on *E.coli*-country combinations in comparison to 14-day average values; below: daily number of articles for the *E. coli* category)




Fig. 4. Number of articles in the *E.coli* category that also mention one of the entities Robert Koch Institute, Commissioner John Dalli, European Commission, ECDC and EFSA

3 Conclusions

Media monitoring is a well-established technique for Epidemic Intelligence [9]. In the case of the *E.coli* outbreak, the early media reports were identified quickly and accurately; an alert was published on Saturday 21st of May via MedISys. Obviously, the authorities in Germany were informed earlier via the established indicator-based systems in Germany, and the EWRS message published on Sunday 22 notified all EU member states about the outbreak. Thus, media monitoring was less used for early alerting, but rather to identify key aspects of the developing story (which were further disseminated, e.g. in reports for the European Commission). Since MedISys covers both general media and official sites, users were able to analyze their reports in parallel.

Acknowledgments. We thank Prof Dr Nikolaos Stilianakis and Dr Agnès Rortais for fruitful discussions and comments on the manuscript.

References

1. Linge, J.P., Steinberger, R., Weber, T.P., Yangarber, R., van der Goot, E., Al-Khudhairi, D., Stilianakis, N.I.: Internet surveillance systems for early alerting and health threats. *Euro. Surveill.* 14(13), 19162 (2009)
2. Linge, J.P., Belyaeva, J., Steinberger, R., Gemo, M., Fuart, F., Al-Khudhairi, D., Bucci, S., Yangarber, R., van der Goot, E.: MedISys: Medical Information System. In: Asimakopoulou, E., Bessis, N. (eds.) *Advanced ICTs for Disaster Management and Threat Detection: Collaborative and Distributed Frameworks*, pp. 131–142. IGI Global (2010)
3. Mantero, J., Belyaeva, J., Linge, J.P.: How to maximise event-based surveillance web-systems: the example of ECDC/JRC collaboration to improve the performance of MedISys. *JRC Scientific and Technical Reports (EUR collection)*. Publications Office of the European Union (2011) ISBN 978-92-79-19755-0. ISSN 1018-5593, doi: 10.2788/69804

4. Frank, C., Werber, D., Cramer, J.P., Askar, M., Faber, M., Heiden, M.A., Bernard, H., Fruth, A., Prager, R., Spode, A., Wadl, M., Zoufaly, A., Jordan, S., Stark, K., Krause, G.: The HUS Investigation Team: Epidemic Profile of Shiga-Toxin-Producing *Escherichia coli* O104:H4 Outbreak in Germany - Preliminary Report. *N Engl. J. Med.*, June 22 (2011)
5. Guglielmetti, P., Coulombier, D., Thinus, G., Van Loock, F., Schreck, S.: The Early Warning and Response System for communicable diseases in the EU: an overview from 1999 to 2005. *Euro. Surveill.* 11(12), 666 (2006), <http://www.eurosurveillance.org/ViewArticle.aspx?ArticleId=666>
6. ProMED-mail: *E. coli* VTEC - Germany: RFI. ProMED-mail 2011, May 23: 20110523.1566, <http://www.promedmail.org> (accessed July 3, 2011)
7. Linge, J.P., Belyaeva, J.: Customization of MediSys for the monitoring of food and feed hazards (2011), <http://www.efsa.europa.eu/en/supporting/pub/140e.htm>
8. EFSA: Tracing seeds, in particular fenugreek (*Trigonella foenum-graecum*) seeds, in relation to the Shiga toxin-producing *E. coli* (STEC) O104:H4 2011 Outbreaks in Germany and France. EFSA Report. Question No EFSA-Q-2011-00817 (2011), <http://www.efsa.europa.eu/en/supporting/pub/176e.htm> (issued on July 05, 2011)
9. Paquet, C., Coulombier, D., Kaiser, R., Ciotti, M.: Epidemic intelligence: a new framework for strengthening disease surveillance in Europe. *Euro. Surveill.* 11, 665 (2006), <http://www.eurosurveillance.org/ViewArticle.aspx?ArticleId=665>

Syndromic Classification of Twitter Messages

Nigel Collier and Son Doan

National Institute of Informatics, Tokyo, Japan

collier@nii.ac.jp

<http://born.nii.ac.jp/dizie>

Abstract. Recent studies have shown strong correlation between social networking data and national influenza rates. We expanded upon this success to develop an automated text mining system that classifies Twitter messages in real time into six syndromic categories based on key terms from a public health ontology. 10-fold cross validation tests were used to compare Naive Bayes (NB) and Support Vector Machine (SVM) models on a corpus of 7431 Twitter messages. SVM performed better than NB on 4 out of 6 syndromes. The best performing classifiers showed moderately strong F1 scores: respiratory = 86.2 (NB); gastrointestinal = 85.4 (SVM polynomial kernel degree 2); neurological = 88.6 (SVM polynomial kernel degree 1); rash = 86.0 (SVM polynomial kernel degree 1); constitutional = 89.3 (SVM polynomial kernel degree 1); hemorrhagic = 89.9 (NB). The resulting classifiers were deployed together with an EARS C2 aberration detection algorithm in an experimental online system.

Keywords: epidemic intelligence, social networking, machine learning, natural language processing.

1 Introduction

Twitter is a social networking service that allows users throughout the world to communicate their personal experiences, opinions and questions to each other using micro messages ('tweets'). The short message style reduces thought investment [1] and encourages a rapid 'on the go' style of messaging from mobile devices. Statistics show that Twitter had over 200 million users¹ in March 2011, representing a small but significant fraction of the international population across both age and gender² with a bias towards the urban population in their 20s and 30s. Our recent studies into novel health applications [2] have shown progress in identifying free-text signals from tweets that allow influenza-like illness (ILI) to be tracked in real time. Similar studies have shown strong correlation with national weekly influenza data from the Centers for Disease Control and Prevention and the United Kingdom's Health Protection Agency. Approaches like these hold out the hope that low cost sensor networks could be deployed as early

¹ <http://www.bbc.co.uk/news/business-12889048>

² <http://sustainablecitiescollective.com/urbantickurbantick/20462/twitter-usage-view-america>

warning systems to supplement more expensive traditional approaches. Web-based sensor networks might prove to be particularly effective for diseases that have a narrow window for effective intervention such as pandemic influenza.

Despite such progress, studies into deriving linguistic signals that correspond to other major syndromes have been lacking. Unlike ILI, publicly available gold standard data for other classes of conditions such as gastrointestinal or neurological illnesses are not so readily available. Nevertheless, the previous studies suggest that a more comprehensive early warning system based on the same principles and approaches should prove effective. Within the context of the DIZIE project, the contribution of this paper is (a) to present our data classification and collection approaches for building syndromic classifiers; (b) to evaluate machine learning approaches for predicting the classes of unseen Twitter messages; and (c) to show how we deployed the classifiers for detecting disease activity. A further goal of our work is to test the effectiveness of outbreak detection through geo-temporal aberration detection on aggregations of the classified messages. This work is now ongoing and will be reported elsewhere in a separate study.

1.1 Automated Web-Sensing

In this section we make a brief survey of recent health surveillance systems that use the Web as a sensor source to detect infectious disease outbreaks. Web reports from news media, blogs, microblogs, discussion forums, digital radio, user search queries etc. are considered useful because of their wide availability, low cost and real time nature. Although we will focus on infectious disease detection it is worth noting that similar approaches can be applied to other public health hazards such as earthquakes and typhoons [3,4].

Current systems fall into two distinct categories: (a) event-based systems that look for direct reports of interest in the news media (see [5] for a review), and (b) systems that exploit the human sensor network in sites like Twitter, Jaiku and Prownce by sampling reports of symptoms/GP visits/drug usage etc. from the population at risk [6,7,8]. Early alerts from such systems are typically used by public health analysts to initiate a risk analysis process involving many other sources such as human networks of expertise.

Work on the analysis of tweets, whilst still a relatively novel information source, is related to a tradition of syndromic surveillance based on analysis of triage chief complaint (TCC) reports, i.e. the initial triage report outlining the reasons for the patient visit to a hospital emergency room. Like tweets they report the patient's symptoms, are usually very brief, often just a few keywords and can be heavily abbreviated. Major technical challenges though do exist: unlike TCC reports tweets contain a very high degree of noise (e.g. spam, opinion, re-tweeting etc.) as well as slang (e.g. *itcy* for *itchy*) and emoticons which makes them particularly challenging. Social media is inherently an informal medium of communication and lacks a standard vocabulary although Twitter users do make use of an evolving semantic tag set. Both TCC and tweets often consist of short telegraphic statements or ungrammatical sentences which are difficult for uncustomised syntactic parsers to handle.

In the area of TCC reports we note work done by the RODS project [9] that developed automatic techniques for classifying reports into a list of syndromic categories based on natural language features. The chief complaint categories used in RODS were respiratory, gastrointestinal, botulinic, constitutional, neurologic, rash, hemorrhagic and none. Further processes took aggregated data and issued alerts using time series aberration detection algorithms. The DIZIE project which we report here takes a broadly similar approach but applies it to user generated content in the form of Twitter messages.

2 Method

DIZIE currently consists of the following components: (1) a list of latitudes and longitudes for target world cities based on Twitter usage; (2) a lexicon of syndromic keywords used as an initial filter, (3) a supervised machine learning model that converts tweets to a word vector representation and then classifies them according to six syndromes, (4) a post-processing list of stop words and phrases that blocks undesired contexts, (5) a MySQL database holding historic counts of positive messages by time and city location, used to calculate alerting baselines, (6) an aberration detection algorithm, and (7) a graphical user interface for displaying alerts and supporting evidence.

After an initial survey of high frequency Twitter sources by city location we selected 40 world cities as candidates for our surveillance system. Sampling in the runtime system is done using the Twitter API by searching for tweets originating within a 30km radius of a city's latitude and longitude, i.e. a typical commuting/shopping distance from the city centre. The sampling rate is once every hour although this can be shortened when the system is in full operation. In this initial study we focussed only on English language tweets and how to classify them into 6 syndromic categories which we describe below.

Key assumptions in our approach are that: (a) each user is considered to be a sensor in the environment and as such no sensor should have the capacity to over report. We controlled over reporting by simply restricting the maximum number of messages per day to be 5 per user; (b) each user reports on personal observations about themselves or those directly known to them. To control (a) and (b) and prevent over-reporting we had to build in filtering controls to mitigate the effects of information diffusion through re-reporting, particularly for public personalities and mass events. Re-tweets, i.e. repeated messages, and tweets involving external links were automatically removed.

2.1 Schema Development

A syndrome is a collection of symptoms (both specific and non-specific) agreed by the medical community that are indicative of a class of diseases. We chose six syndrome classes as the targets of our classifier: constitutional, respiratory, gastrointestinal, hemorrhagic, rash (i.e. dermatological) and neurological. These were based on an openly available public health ontology developed as part

of the BioCaster project [10] by a team of experts in computational linguists, public health, anthropology and genetics. Syndromes within the ontology were based on RODS syndrome definitions and are linked to symptom terms - both technical and laymen's terms - through typed relations. We use these symptoms (syndromic keywords) as the basis for searching Twitter and expanded them using held out Twitter data.

2.2 Twitter Data

After defining our syndromes we examined a sample of tweets and wrote guidelines outlining positive and negative case definitions. These guidelines were then used by three student annotators to classify a sample of 2000 tweets per syndrome into positive or negative for each of the syndrome classes. Data for training was collected by automatically searching Twitter using the syndromic keywords over the period 9th to 24th July 2010. No city filtering was applied when we collected the training data. Typical positive example messages are: "Woke up with a stomach ache!", "Every bone in my body hurts", and "Fever, back pain, headache... ugh!". Examples of negative messages are: "I'm exercising till I feel dizzy", "Cabin fever is severe right now", "Utterly exhausted after days of housework". Such negative examples include a variety of polysemous symptom words such as *fever* in its senses of raised temperature and excitement and *headache* in its senses of a pain in the head or an inconvenience. The negative examples also include cases where the context indicates that the cause of the symptom is unlikely to be an infection, e.g. headache caused by working or exercising. The training corpus is characterised using the top 7 terms calculated by mutual association score in Table 1. This includes several spurious associations such as 'rt' standing for 'repeat tweet', 'botox' which is discussed extensively as a treatment for several symptoms and 'charice' who is a new pop idol.

The final corpus was constructed from messages where there was total agreement between all three annotators. This data set was used to develop and evaluate supervised learning classifiers in cross-fold validation experiments. A summary of the data set is shown in Table 2. Inter-annotator agreement scores between the three annotators are given as Kappa showing agreement between the two highest agreeing annotators. Kappa indicates strong agreement on most syndromic classes with the notable exception of gastrointestinal and neurological.

2.3 Classifier Models

DIZIE employs a two stage filtering process. Since Twitter many topics unrelated to disease outbreaks, DIZIE firstly requests Twitter to send it messages that correspond to a set of core syndromic keywords, i.e. the same sampling strategy used to collect training/testing data. These keywords are defined in the BioCaster public health ontology [10]. In the second stage messages which are putatively on topic are filtered more rigorously using a machine learning approach. This stage of filtering aims to identify messages containing ambiguous

Table 1. Top 7 terms by syndrome calculated by mutual information score. * indicates a spurious association.

Resp	Gastro	Const	Hemor	Rash	Neuro
throat	stomach	botox*	pain	road	headache
sore	ache	body	hemorrhage	heat	coma
cough	gib	charice*	muscle	arm	worst
flu	feel	jaw	tired	tired	gave
nose	rt*	hurts	pray	rash	giving
rt*	bad	stomach	brain	itcy	vertigo
cold	worst	sweating	guiliechelon*	face	pulpo*

Table 2. Structure of the annotated syndrome corpus of Twitter messages

Syndrome	Positives (P)	Negatives (N)	P/N	Kappa
Respiratory	627	738	0.85	0.67%
Gastrointestinal	489	676	0.72	0.49%
Neurological	549	434	1.26	0.42%
Rash	914	592	1.54	0.86%
Hemorrhagic	320	711	0.45	0.92%
Constitutional	1043	338	3.09	0.78%

words whose senses are not relevant to infectious diseases and messages where the cause of the symptoms are not likely to be infectious diseases. About 70% of messages are removed at this second stage.

To aid in model selection our experiments used two widely known machine learning models to classify Twitter messages into a fixed set of syndromic classes: Naive Bayes (NB) and support vector machines (SVM) [11] using a variety of kernel functions. Both models were trained with binary feature vectors representing a dictionary index of words in the training corpus. i.e. a feature for the test message was marked 1 if a word was present in the test message which had been seen previously in the training corpus otherwise 0. No normalisation of the surface words was done, e.g. using stemming, because of the high out of vocabulary rate with tools trained on general language texts.

Despite the implausibility of its strong statistical independence assumption between words, NB tends to perform strongly. The choice to explore keywords as features rather than more sophisticated parsing and conceptual analysis such as MPLUS [12] was taken from a desire to evaluate less expensive approaches before resorting to time consuming knowledge engineering.

The NB classifier exploits an estimation of the Bayes Rule:

$$P(c_k|d) = \frac{P(c_k) \times \prod_{i=1}^m P(f_i|c_k)^{f_i(d)}}{P(d)} \quad (1)$$

where the objective is to assign a given feature vector for a document d consisting of m features to the highest probability class c_k . $f_i(d)$ denotes the frequency count of feature i in document d . Typically the denominator $P(d)$ is not computed explicitly as it remains constant for all c_k . In order to compute the highest

value numerator NB makes an assumption that features are conditionally independent given the set of classes. Right hand side values of the equation are estimates based on counts observed in the training corpus of classified Twitter messages. We used the freely available Rainbow toolkit³ from CMU as the software package.

SVMs have been widely used in text classification achieving state of the art predictive accuracy. The major distinction between the two approaches are that whereas NB is a generative classifier which forms a statistical model of each class, SVM is a large-margin binary classifier. SVM operates as a two stage process. Firstly the feature vectors are projected into a high dimensional space using a kernel function. The second stage finds a maximum margin hyperplane within this space that separates the positive from the negative instances of the syndromic class. In practice it is not necessary to perfectly classify all instances with the level of tolerance for misclassification being controlled by the C parameter in the model. A series of binary classifiers were constructed (one for each syndrome) using the SVM^{Light} software package⁴. We explored polynomial degree 1, 2, 3 and radial basis function kernels.

2.4 Temporal Model

In order to detect unexpected rises in the stream of messages for each syndrome we implemented a widely used change point detection algorithm called the Early Aberration and Reporting System (EARS) C2 [13]. C2 reports an alert when its test value S_t exceeds a number k of standard deviations above a historic mean:

$$S_t = \max(0, (C_t - (\mu_t + k\sigma_t))/\sigma_t) \quad (2)$$

where C_t is the count of classified tweets for the day, μ_t and σ_t are the mean and standard deviation of the counts during the history period, set as the previous two weeks. k controls the number of standard deviations above the mean where an alert is triggered, set to 1 in our system. The output of C2 is a numeric score indicating the degree of abnormality but this by itself is not so meaningful to ordinary users. We constructed 5 banding groups for the score and showed this in the graphical user interface.

3 Results

3.1 Classifying Twitter Messages

Results for 10-fold cross validation experiments on the classification models are shown in Table 3. Overall the SVM with polynomial degree 1 kernel outperformed all other kernels with other kernels generally offering better precision at a higher cost to recall. Precision (Positive predictive) values ranged from 82.0 to

³ <http://www.cs.cmu.edu/~mccallum/bow/rainbow/>

⁴ <http://svmlight.joachims.org/>

93.8 for SVM (polynomial degree 1) and from 83.3 to 99.0 for NB. Recall (sensitivity) values ranged from 58.3 to 96.2 for SVM (polynomial degree 1) and from 74.7 to 90.3 for NB. SVM tended to offer a reduced level of precision but better recall. In the case of one syndrome (Hemorrhagic) we noticed an unusually low level of recall for SVM but not for NB.

SVM's performance seemed moderately correlated to the positive/negative ratio in the training corpus and also showed weakness for the two classes (Hemorrhagic and Gastrointestinal) with the smallest positive counts. Naive Bayes performed robustly across classes with no obvious correlation either to positive/negative ratio or the volume of training data. Low performance was seen in both models for the gastrointestinal syndrome. This was probably due to the low number of training examples resulting from the low inter-annotator agreement on this class and the requirement for complete agreement between all three annotators.

Table 3. Evaluation of automated syndrome classification using naive Bayes and Support Vector Machine models on 10-fold cross validation. P - Precision, R - Recall, F1 - F1 score. ¹ SVM using a linear kernel, ² SVM using a polynomial kernel degree 2, ³ SVM using a polynomial kernel degree 3, ^R SVM using a radial basis function kernel.

	Naive Bayes			SVM ¹			SVM ²			SVM ³			SVM ^R		
Synd.	P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1
Resp.	90.3	82.4	86.2	85.4	82.5	83.8	83.0	71.0	76.5	86.4	61.3	71.7	66.7	3.2	6.2
Gast.	83.3	75.5	79.2	85.9	78.4	81.8	92.7	79.2	85.4	91.4	66.7	77.1	73.1	39.6	51.3
Neur.	98.2	74.7	84.8	83.2	95.0	88.6	77.9	98.2	86.9	62.4	98.2	76.3	90.0	63.0	74.1
Rash	94.5	76.1	84.3	82.0	90.6	86.0	76.9	91.2	83.4	67.7	94.5	78.9	60.7	100.0	75.5
Hem.	89.4	90.3	89.9	93.8	58.3	71.7	100.0	50.0	66.7	100.0	50	66.7	87.5	43.8	58.3
Con.	99.0	79.8	88.4	83.6	96.2	89.3	83.6	93.3	88.2	78.6	99.0	87.7	76.5	100	86.7

3.2 Technology Dissemination

An experimental service for syndromic surveillance called DIZIE has been implemented based on the best of our classifier models and we are now observing its performance. The service is freely available from an online portal at <http://born.nii.ac.jp/dizie>. As shown in Figure 3.2 the graphical user interface (GUI) for DIZIE shows a series of radial charts for each major world city with each band of the chart indicating the current level of alert for one of the six syndromes. Alerting level scores are calculated using the Temporal Model presented above. Each band is colour coded for easy recognition. Alerting levels are calculated on the classified twitter messages using the EARS C2 algorithm described above. Data selection is by city and time with drill down to a selection of user messages that contributed to the current level. Trend bars show the level of alert and whether the trend is upwards, downwards or sideways. Charting is also provided over an hourly, daily, weekly and monthly period. The number of positively classified messages by city is indicated in Figure 3.2 for a selection of cities.




Fig. 1. Radial graphs showing syndromic alert levels for major world cities. Colour coding on the radial segments indicates the alerting degree automatically assigned to a syndrome in a city based on the previous hour's Twitter counts and the previous 2 weeks as a baseline. The page is updated every hour. Clicking on the graph for a city displays the frequency graph and also the matching tweets for the current hour.




Fig. 2. Number of Tweets by a sample of major world cities classified by DIZIE during the period 2nd March 2011 to 31st August 2011

Navigation links are provided to and from BioCaster, a news event alerting system, and we expect in the future to integrate the two systems more closely to promote greater situation awareness across media sources. Access to the GUI is via regular Web browser or mobile device with the page adjusting automatically to fit smaller screens.

4 Conclusion

Twitter offers unique challenges and opportunities for syndromic surveillance. Approaches based on machine learning need to be able (a) to handle biased data, and (b) to adjust to the rapidly changing vocabulary to prevent a flood of false positives when new topics trend. Future work will compare keyword classifiers against more conceptual approaches such as [12] and also compare the performance characteristics of change point detection algorithms.

Based on the experiments reported here we have built an experimental application called DIZIE that samples Twitter messages originating in major world cities and automatically classifies them according to syndromes. Access to the system is openly available. Based on the outcome of our follow up study we intend to integrate DIZIE's output with our event-based surveillance system BioCaster which is currently used by the international public health community.

Acknowledgements. This work was in part supported by grant in aid support from the National Institute of Informatics' Grand Challenge Project (PI:NC). We are grateful to Reiko Matsuda Goodwin for commenting on the user interface in the early stages of this study and helping in data collection for the final system.

References

1. Java, A., Song, X., Finin, T., Tseng, B.: Why we twitter: Understanding microblogging usage and communities. In: Proc. 9th WebKDD and 1st SNA-KDD Workshop on Web Mining and Social Network Analysis, August 12. ACM (2007)
2. Collier, N., Nguyen, S.T., Nguyen, M.T.N.: OMG U got flu? analysis of shared health messages for bio-surveillance. *Biomedical Semantics* 2(suppl. 5), S10 (2011)
3. Earle, P.: Earthquake twitter. *Nature Geoscience* 3(4), 221–222 (2010), doi:10.1038/ngeo832
4. Sakaki, T., Okazaki, M., Matsuo, Y.: Earthquake shakes twitter users: real-time event detection by social sensors. In: Proc. of the 19th International World Wide Web Conference, Raleigh, NC, USA, pp. 851–860 (2010)
5. Hartley, D., Nelson, N., Walters, R., Arthur, R., Yangarber, R., Madoff, L., Linge, J., Mawudeku, A., Collier, N., Brownstein, J., Thinus, G., Lightfoot, N.: The landscape of international biosurveillance. *Emerging Health Threats J.* 3(e3) (January 2010), doi:10.1093/bioinformatics/btn534
6. Szomszor, M., Kostkova, P., De Quincey, E.: swineflu : Twitter predicts swine flu outbreak in 2009 (December 2009)
7. Lampos, V., De Bie, T., Cristianini, N.: Flu Detector - Tracking Epidemics on Twitter. In: Balcázar, J.L., Bonchi, F., Gionis, A., Sebag, M. (eds.) ECML PKDD 2010. LNCS, vol. 6323, pp. 599–602. Springer, Heidelberg (2010)
8. Signorini, A., Segre, A.M., Polgreen, P.M.: The use of twitter to track levels of disease activity and public concern in the U.S. during the influenza a h1n1 pandemic. *PLoS One* 6(5), e19467 (2011)
9. Wagner, M.M., Espino, J., Tsui, F.C., Gesteland, P., Chapman, W., Ivanov, W., Moore, A., Wong, W., Dowling, J., Hutman, J.: Syndrome and outbreak detection using chief-complaint data - experience of the real-time outbreak and disease surveillance project. *Morbidity and Mortality Weekly Report (MMWR)* 53(suppl.), 28–31 (2004)

10. Collier, N., Doan, S., Kawazoe, A., Matsuda Goodwin, R., Conway, M., Tateno, Y., Ngo, Q., Dien, D., Kawtrakul, A., Takeuchi, K., Shigematsu, M., Taniguchi, K.: BioCaster: detecting public health rumors with a web-based text mining system. *Bioinformatics* 24(24), 2940–2941 (2008), doi:10.1093/bioinformatics/btn534
11. Joachims, T.: Text categorization with support vector machines: Learning with many relevant features. In: Nédellec, C., Rouveirol, C. (eds.) *ECML 1998*. LNCS, vol. 1398, pp. 137–142. Springer, Heidelberg (1998)
12. Christensen, L.M., Haug, P.J., Fiszmann, M.: Mplus: A probabilistic medical language understanding model. In: *Proceedings of the Workshop on Natural Language Processing in the Biomedical Domain*, Philadelphia, USA (July 2002)
13. Hutwagner, L., Thompson, W., Seeman, M.G., Treadwell, T.: The bioterrorism preparedness and response early aberration reporting system (EARS). *J. Urban Health* 80(2), i89–i96 (2003)

Towards Spatial Description-Based Integration of Biomedical Atlases

Nurzi Juana Mohd Zaizi¹ and Albert Burger^{1,2}

¹ Department of Computer Science, Heriot-Watt University, Edinburgh, UK

² MRC Human Genetics Unit, Edinburgh, UK

Abstract. Biomedical imaging has become ubiquitous in both, basic research and the clinical sciences. Technology advances, and the resulting multitude of imaging modalities, have led to a sharp rise in the quantity and quality of such images. Whether for epidemiological studies, educational uses, monitoring the clinical progress of a patient or translational science purposes, being able to integrate and compare such image-based data has developed into an increasingly critical component in the Life Sciences and eHealth domain. Image processing-based solutions have difficulties when the underlying morphologies are too different. Ontology-based solutions often lack spatial precision. In this paper, we describe a compromise solution which captures location in biomedical images via spatial descriptions using so-called fiducial points. The work is discussed in the context of biomedical atlases and includes, in addition to the introduction of the basic method, some experimental performance results.

Keywords: spatial description, integration of biomedical images/atlasses.

1 Introduction

Patients are now routinely undergoing a variety of medical imaging investigations, such as Magnetic Resonance Imaging (MRI) and Computed Tomography (CT) scanning, and the images resulting from these investigations become part of the patients' medical records. The same as well as other imaging techniques, e.g. optical imaging, are also used in preclinical studies and the Life Sciences. The work presented in this paper is rooted in the latter and uses examples from biomedical atlases, where we consider an atlas to consist of the image data components, a set of labels describing structures in the images, and mappings between them. There are questions of data integration within the domain of clinical images, within the Life Sciences image datasets, as well as between human and model organism data. The latter being particularly of interest in the translational sciences.

Biologists have access to a variety of biomedical atlases. Many of these atlases are data sources for the same experimental field, for example, mouse gene expression data. Though storing the same type of data, different experimental designs, varying analysis of results and different update routines have caused the data in these atlases to be different. The consequence is that these atlases may provide

different results even for the same query. It is vital that multiple resources in the same field are used so that full and complete results can be generated for the query. The comparative clinical issue is the integration of different medical images for a single patient, or the comparison of images of multiple patients with the same disease.

A biomedical atlas consists of a graphical model, the ontology associated with the graphical model and a mapping between these two. The ontology contains a collection of anatomical domains and relations between these domains. The graphical model is the image for a mammalian with those anatomical domains. This paper proposes the integration of these data sources by mapping the images of biomedical atlases using spatial descriptions. Given two images $I1$ and $I2$, mapping one image onto another means that, for each anatomical space in image $I1$, we try to find a corresponding space, which has the same intended meaning, in image $I2$. For this study, we circumvent the extra complexity of image segmentation by considering anatomical domains that can be easily segmented. More specifically, we explore 2D image space of mouse embryo domains.

Mappings anatomical spaces concern a number of issues. Different images may have a different number of segmented regions causing one structure to correspond to parts of several structures, and vice versa. Furthermore, even if these images may have the same anatomical structures, the morphology may vary with scale, orientation and the position of the structure. In addition, different biomedical atlases may have the same segmented images but may use different anatomical names causing interoperability issue of finding correspondences anatomical regions between these images. An efficient representation structure is necessary to conceptualize anatomical space of an image to guide the mapping process. This paper explores spatial description-based approach for the linking of images for the integration of biomedical atlases.

Section 2 presents an overview of image mapping approaches. The proposed integration approach is described in Section 3. Section 4 provides experimental results of the proposed approach. In Section 5, a discussion is presented. Finally, a conclusion in Section 6.

2 An Overview of Image Mapping Approaches

This section discusses two approaches for mapping. In particular, its focus is on the following approaches: (1) ontology based mapping (2) image processing based mapping. Ontology based mapping depends on spatial relations between anatomical regions, whereas, mapping using image processing depends on fiducial points.

2.1 Spatial Relations: Ontology Based

This section describes the mapping between anatomical spaces across images using ontologies. Mappings based on an ontology start by segmenting the image of a biomedical atlas according to its anatomical regions. Subsequently, the regions

can be linked to the appropriate concepts in the atlas' anatomy ontology. Two regions are then mapped according to the similarity of their spatial relationships. Given two atlas anatomy ontologies O_1 and O_2 , if anatomical structure A_1 in ontology O_1 has the relationships A_1 adjacent to B_1 , and A_1 is adjacent to C_1 then its equivalent anatomical structure, A_2 , in ontology O_2 , must be adjacent to B_2 and C_2 , where the latter two correspond to B_1 and C_1 , respectively. The integration of biomedical atlases can then be achieved by linking between their respective anatomy ontologies.

The concepts of spatial relations have been well employed in ontologies by both FMA [1] and Bittner et al. 2008 [2] to describe anatomical space in the biomedical domain. In general, spatial relations between anatomical entities are described using relationships from the following categories:

Mereological relations describe the concept of parthood between the whole and its parts, e.g., finger is *part of* hand, hand is *part of* the arm etcetera.

Topological relations describe the concept of adjacency, discreteness, and connectedness among entities. Two entities are described as being *adjacent* when they are close to each other, but not connected. *Discrete* entities are not connected. If two entities have a common anatomical space, such that they partially coincide or are externally attached with one another, they are said to be *connected*., e.g., two entities are externally connected if the distance between them is zero and do not overlap, for example, in human major parts of the joint, the synovial cavity is externally connected to the synovial membrane [2].

Location relations describe the concept of relative location between entities that may coincide wholly or partially without being *part of* one another, for example, the brain is located in (but not *part of*) cranial cavity.

Based on spatial relations, for example, anatomical region x in Figure 1(a) is mapped to the result region y in Figure 1(b) if x is described as:

$$'adjacent(x, midgut), adjacent(x, liver)'$$




Fig. 1. Based on spatial adjacency between (a) anatomical region x with other anatomical regions will map x to (b) the result region y

2.2 Fiducial Points: Image Processing Based

This section discusses the mapping between biomedical atlases based on image processing techniques. These methods start by examining the pixels in an image to classify them into regions, e.g. [3,4]. Classification is by the pixel's intensity level. Subsequently, a registration algorithm is required to identify equivalent regions, across images, based on pixels. In addition, based on the pixel classification, fiducial points can be located. A fiducial point is a point in space in either 2D or 3D, typically an anatomical landmark which is easily recognizable in an image, usually identified by human experts and possibly assisted by auto/semi-automated image processing algorithms. These fiducial points are typically used for registration experimentation image of canonical atlas. Izard and Jedynak [5] describes a registration approach which employs a Bayesian model to detect these points in order to map between regions across images. Registration technique as proposed by Khaissidi et al. 2009 [6] uses the Hough Transform algorithm to align medical images, based on fiducial points extracted from the two compared images. However, the drawback of image processing based mapping in general is that it has the possibility to fail if there is a large variation in pixel/voxel intensity [7].

3 Concept of Spatial Description

3.1 Spatial Description Based on Fiducial Points and a Set of Spatial Relations

The proposed approach of mapping involved the concepts of a query region, fiducial points and fiducial lines. A fiducial point is a point in space. A fiducial line is made up by creating a straight line through a pair of fiducial points. A query region is made up of connected multiple single-elements within a closed boundary. Given two images $I1$ and $I2$, mapping one image onto another starts by selecting the same fiducial points in both images. We then describe a query region using spatial relations between the query region with respect to the fiducial points. Two regions from different images are then mapped according to the similarity of their spatial description. For example, if query region X in image $I1$ is described as X is north of fiducial point $P1$ and X is west of fiducial point $P2$, then its equivalent region in image $I2$ must be a region that is located north of fiducial point $P1$ and west of fiducial point $P2$. Figure 2 depicts the framework of the approach. The image processing based mapping inspires the idea of using fiducial points. Because a fiducial point can become a point of reference for an anatomical location, this paper intends to describe a query region based on these points using spatial relations. The idea of using spatial relations, on the other hand, is inspired by ontology based mapping. Because a spatial relation can describe the location of a region in space, this paper proposes to describe a query region using fiducial points and a set of spatial relations. By the use of spatial relations, this approach works independently of image pixel/voxel intensities.

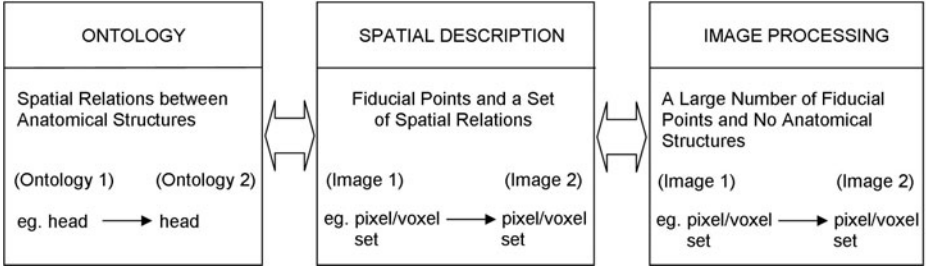


Fig. 2. Ontology-based mappings align images by identifying correspondences among elements of two ontologies based on spatial relations between anatomical structures. The image processing based mappings align images based on equivalent pixel/voxel intensities corresponding the fiducial points

On the other hand, the use of fiducial points allows this approach to work independently of spatial relations between segmented regions. The approach does not intend to include a large number of concepts in spatial relations as that replicates the ontology mapping approach. The entire spatial area of an image should be conceptualized with a small number of fiducial points such that the attempt is not a replicate to the image processing mapping approach. We now summarize the formalism of the approach. We define directional relations as

$$D = \{northOf, eastOf, southOf, westOf\} \quad (1)$$

We describe a query region x in an image as

$$S_Q = \{r(x, f_i) \mid r \in D \text{ and } f_i \in (F_{point} \cup F_{line})\} \quad (2)$$

where S_Q is the spatial description for query region x with respect to a fiducial point $F_{point} = \{p_1, p_2, \dots, p_n\}$ or a fiducial line $F_{line} = \{l_1, l_2, \dots, l_n\}$

Figure 3 depicts two images of mouse embryo with 6 fiducial points and 15 fiducial lines. The simplified description for query region X is described as:

$$'southOf(X, P6P2), eastOf(X, P1P4), northOf(X, P2P5), westOf(X, P1P3)'$$

Note that, in the description, we label a fiducial line according to its pair of fiducial points. The location highlighted in Figure 3(b) denotes the matched location corresponding to the description.

4 Experimental Results

A series of experiments were conducted to demonstrate how fiducial points and a set of spatial relations can be used to describe locations. For the experiments, an image representing the mouse embryo was used and 102 spatial regions were annotated in the image. The image generated 97104 query regions each of size




Fig. 3. Spatial description based on fiducial points and a set of spatial relations maps (a) query region X to (b) result region Y

50x50 squared pixels, 68154 query regions each of size 100x100 squared pixels, 44204 query regions each of size 150x150 squared pixels, and 25254 query regions each of size 200x200 squared pixels. For all query regions of size 50x50 squared pixels, the first query region starts at the top-left corner of the image and is increased every time by one pixel in order to generate the following query region and so on. Query regions of other sizes are also generated by following this step of one pixel. The idea of using query regions is to test the mappings of pixels in a query region of one image to pixels in a region of another image based on fiducial points. The percentage of accuracy is calculated by dividing the total number of pixels in X by the total number of pixels in result region Y , and multiply by 100 (see Figure 3). Figure 4(A) depicts the average percentage of accuracy served by number of fiducial points. Results show that the more fiducial points

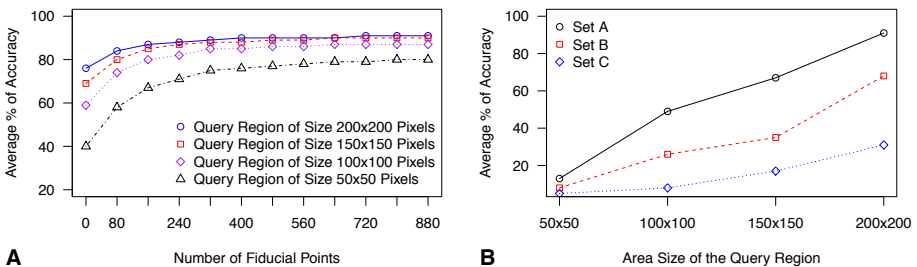


Fig. 4. (A) Average percentage of accuracy served by number of fiducial points. The more fiducial points are included the higher the average percentage of accuracy gets. Moreover, the average percentage of accuracy substantially increases as the query region area size gets larger. (B) Average percentage of accuracy in three different positioning sets of 8 fiducial points served by query region area size. The same number of fiducial point place at different positions produce different average percentage of accuracy.

were included, the more accurate the mapping was. In addition, the mapping accuracy substantially increases as the query region area size gets larger. Fiducial points provide qualitative spatial relations to describe locations. Therefore, the more fiducial points are used the more spatial relations are available to describe locations, which increases the average percentage of accuracy. In general, spatial descriptions will return a location which is either larger or equal to the actual area location. Thus, for cases where spatial descriptions for the corresponding query regions do not return locations that are exactly equal to the actual location, the larger the size of the query region, the more accurate it is to the actual area, by which contribute to much higher accuracy value compared to the smaller one. Figure 4(B) depicts the average percentage of accuracy in three different positioning sets of 8 fiducial points served by query region size. Results show that the same number of fiducial points placed at different positions produce different accuracy. The positions of fiducial points determine spatial relations made available to describe locations. Because the location is determined by spatial descriptions, different positioning set for the same number of fiducial points certainly contributes to different spatial descriptions to describe locations, which produce different average percentages of accuracy. Overall, with the appropriate number of fiducial points used and better selection of fiducial point location, mappings can be improve in terms of accuracy.

5 Discussion

The definition for best match criteria is important in any mapping algorithm. Because anatomical structures exist at different range of scale, arrangement and the position, there is a possibility for an exact copy of location corresponding the query region in one image to be unavailable in another image. The proposed spatial description approach at the current state perform mappings by returning a location that satisfies all spatial relation constraints corresponding to a query region. However, this may not be necessary. Therefore, the google-style matches can be considered. This can be done by specifying a range, for example, allowing for a distance limit from a fiducial line, which will return a location given by the range.

A preliminary experiment has also been conducted to compare the performance of spatial descriptions based on fiducial points and a set of spatial relations with the following approaches: (1) spatial description based on spatial relationships between segmented regions (2) spatial description based on fiducial points and a set of spatial relations, integrated with spatial relations between segmented regions. Experimental results verified that mapping using spatial description based on spatial relationships between segmented regions managed to produce better accuracy compared to spatial description based on fiducial points and a set of spatial relations. However, this result cannot be used to benchmark the overall mapping performance produced by spatial description based on fiducial points. Depending on better selection of fiducial point locations or by increasing the number of fiducial points used, the mapping accuracy can be further

increased. Furthermore, experimental results have verified that the approach of mapping using spatial description based on fiducial points and a set of spatial relations, integrated with spatial relations between segmented regions can yield significantly higher mapping accuracy compared to using either approach alone.

6 Conclusion and Future Work

This paper explores spatial description based approach to facilitate data integration across biomedical atlases. The most important feature of our approach is that the spatial description, which is rule-based, can provide the means to facilitate the mapping between images of biomedical atlases. Future work includes research on selection of fiducial points where the combination can give performance, as effective as both ontology-based and image processing algorithm; and analyse the capability of spatial description to facilitate data integration between (1) images (i.e. from biomedical atlases), (2) natural-language description of space (i.e. free text from biomedical literature) (3) database warehouses (i.e. structured database of biomedical facts).

Acknowledgments. This work is partially funded by EU projects CUBIST (FP7-ICT-2010-257403) and RICORDO (FP7-ICT-2009-248502).

References

1. Rosse, C., Mejino, J.L.V.: The Foundational Model of Anatomy Ontology. In: *Anatomy Ontologies for Bioinformatics: Principles and Practise*. Springer, Heidelberg (2008)
2. Bittner, T., Donnelly, M., Goldberg, L.J., Neuhaus, F.: Spatial Representation and Reasoning. In: *Anatomy Ontologies for Bioinformatics: Principles and Practise. Modeling Principles and Methodologies*. Springer, Heidelberg (2008)
3. Boccignone, G., Napoletano, P., Ferraro, M.: Embedding diffusion in variational bayes: A technique for segmenting images. *International Journal of Pattern Recognition and Artificial Intelligence* 22(1), 811–827 (2008)
4. Wyawahare, M.V., Patil, P.M., Abhyankar, H.K.: Image Registration Techniques: An overview. *International Journal of Signal Processing, Image Processing and Pattern Recognition* 2(3), 1–5 (2009)
5. Izard, C., Jedynek, B.: Bayesian Registration for Anatomical Landmark Detection. In: *IEEE International Symposium on Biomedical Imaging*, Arlington, VA (2006)
6. Khaissidi, G., Tairi, H., Aarab, A.: A fast medical image registration using feature points. *ICGST-GVIP Journal* 9(3), 19–24 (2009)
7. Walia, E., Suneja, A.: A Conceptual Study on Image Matching Techniques. *Global Journal of Computer Science and Technology* 10(12), 83–88 (2010)
8. Human-fetal week 10 sagittal plane,
http://php.med.unsw.edu.au/embryology/index.php?title=Embryonic_Development

Author Index

- Aidemark, Jan 85
Alves, Bruno 17
Argyropaidas, Panagiotis 76
Arvanitis, Theodoros N. 105
Askenäs, Linda 85
Atkinson, Martin 178
- Bacon, Liz 146
Bahati, Raphael 1
Belyaeva, Jenya 178
Breu, Ruth 97
Bromuri, Stefano 128
Burger, Albert 196
Bursa, Miroslav 26
- Chen, Liming 9
Chronaki, Catherine 76
Chudacek, Vaclav 26
Coleman, Alfred 162
Collier, Nigel 58, 186
- Davis, Michael 170
Delaney, Brendan 138
de Quincey, Ed 146
Detschew, Vesselin 34
Di Tullio, Eugenio 43
Doan, Son 58, 186
- Eiring, Øystein 51
Ekinci, Okan 34
Elloumi, Lamia 158
- Fernando, Juanita 93
Fowler, David 154
Fuart, Flavio 178
- Gattnar, Eva 34
Grasso, Floriana 43
Guy, Stacey 1, 114
Gwadry-Sridhar, Femida 1, 114
- Hammond, Simon 154
Hermens, Hermie 158
Herselman, Marlien E. 162
Hocking, Avril 146
- Hunter, Ruth F. 170
Huptych, Michal 26
Huser, Martin 26
- Ioanitorescu, Radu 109
- Janku, Petr 26
- Katt, Basel 97
Kee, Frank 170
Kontogiannis, Vasilis 76
Kostkova, Patty 80, 154
Krampf, Johannes 128
- Lagani, Vincenzo 76
Langford, Gavin 105
Lawrenson, John 154
Lhotska, Lenka 26
Lim Choi Keung, Sarah N. 105
Linge, Jens P. 178
Luck, Michael 138
- Mair, Richard 97
Mantero, Jas 178
Miles, Simon 138
Mulvenna, Maurice 9
- O’Gorman, Josephine 146
Ogunsina, Ire 105
- Papíková, Vendula 122
Pawar, Pravin 158
Peters, Nicolette 93
Potass, Dalenca 162
- Quinn, Darren 9
- Ratzki-Leewing, Alexandria 1, 114
Rentea, Victor 109
Ruiz, Juan 128
- Sandmark, Hélène 67
Schabetsberger, Thomas 97
Schumacher, Michael 17, 128

- Serbanati, Luca Dan 109
Slaughter, Laura 51
Smedberg, Åsa 67
Spilka, Jiri 26
Szomszor, Martin 154
- Taweel, Adel 138
Trojer, Thomas 97
Tully, Mark A. 170
Tyler, Edward 105
Tyson, Gareth 138
- van Beijnum, Bert-Jan 158
van der Goot, Erik 178
Van Staa, Tjeerd 138
Vasilateanu, Andrei 109
Vo, Bao-Khanh Ho 58
- Walker, Simon 146
- Zaizi, Nurzi Juana Mohd 196
Zhao, Lei 105
Zvolský, Miroslav 122