

Ana Fred  
Joaquim Filipe  
Hugo Gamboa (Eds.)

Communications in Computer and Information Science

127

# Biomedical Engineering Systems and Technologies

Third International Joint Conference, BIOSTEC 2010  
Valencia, Spain, January 2010  
Revised Selected Papers



Springer



Ana Fred Joaquim Filipe Hugo Gamboa (Eds.)

# Biomedical Engineering Systems and Technologies

Third International Joint Conference, BIOSTEC 2010  
Valencia, Spain, January 20-23, 2010  
Revised Selected Papers

Volume Editors

Ana Fred

IST - Technical University of Lisbon

Av. Rovisco Pais, 1, 1049-001 Lisboa, Portugal

E-mail: afred@lx.it.pt

Joaquim Filipe

Polytechnic Institute of Setúbal

EST Campus, Estefanilha, 2910-761 Setúbal, Portugal

E-mail: joaquim.filipe@estsetubal.ips.pt

Hugo Gamboa

Institute of Telecommunications

Av. Rovisco Pais, 1, 1049-001 Lisboa, Portugal

E-mail: hgamboa@gmail.com

ISSN 1865-0929

e-ISSN 1865-0937

ISBN 978-3-642-18471-0

e-ISBN 978-3-642-18472-7

DOI 10.1007/978-3-642-18472-7

Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2011920572

CR Subject Classification (1998): J.3, H.2.8, F.1, F.2.2, G.1.0, K.4.1, H.3-5

© Springer-Verlag Berlin Heidelberg 2011

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](http://www.springer.com))

# Preface

The present book includes extended and revised versions of a set of selected papers from the Third International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2010), organized by the Institute for Systems and Technologies of Information Control and Communication (INSTICC), technically co-sponsored by the IEEE Engineering in Medicine and Biology Society (EMB), and in cooperation with AAAI, Workflow Management Coalition (WfMC), ACM SIGART, Universidad Politecnica de Valencia and Centro en Red en Ingeniería Biomédica (CRIB).

The purpose of the International Joint Conference on Biomedical Engineering Systems and Technologies is to bring together researchers and practitioners, including engineers, biologists, health professionals and informatics/computer scientists, interested in both theoretical advances and applications of information systems, artificial intelligence, signal processing, electronics and other engineering tools in knowledge areas related to biology and medicine.

BIOSTEC is composed of four co-located conferences; each specializes in one of the aforementioned main knowledge areas, namely:

- HEALTHINF (International Conference on Health Informatics) promotes research and development in the application of information and communication technologies (ICT) to healthcare and medicine in general and to the specialized support to persons with special needs in particular. Databases, networking, graphical interfaces, intelligent decision support systems and specialized programming languages are just a few of the technologies currently used in medical informatics. Mobility and ubiquity in healthcare systems, standardization of technologies and procedures, certification, privacy are some of the issues that medical informatics professionals and the ICT industry in general need to address in order to further promote ICT in healthcare.
- BIODEVICES (International Conference on Biomedical Electronics and Devices) focuses on aspects related to electronics and mechanical engineering, especially equipment and materials inspired by biological systems and/or addressing biological requirements. Monitoring devices, instrumentation sensors and systems, biorobotics, micro-nanotechnologies and biomaterials are some of the technologies addressed at this conference.
- BIOSIGNALS (International Conference on Bio-inspired Systems and Signal Processing) is a forum for those studying and using models and techniques inspired by or applied to biological systems. A diversity of signal types can be found in this area, including image, audio and other biological sources of information. The analysis and use of these signals is a multidisciplinary area including signal processing, pattern recognition and computational intelligence techniques, amongst others.

- BIOINFORMATICS (International Conference on Bioinformatics) focuses on the application of computational systems and information technologies to the field of molecular biology, including for example the use of statistics and algorithms to understanding biological processes and systems, with emphasis on new developments in genome bioinformatics and computational biology.

The joint conference, BIOSTEC, received 410 paper submissions from more than 55 countries in all continents. In all, 46 papers were published and presented as full papers, i.e., completed work (8 pages/30' oral presentation), 113 papers reflecting work-in-progress or position papers were accepted for short presentation, and another 78 contributions were accepted for poster presentation. These numbers, leading to a “full-paper” acceptance ratio of about 11% and a total oral paper presentation acceptance ratio close to 39%, show the intention of preserving a high-quality forum for the next editions of this conference.

The conference included a panel and four invited talks delivered by internationally distinguished speakers, namely: Peter D. Karp, Rui M. C. Ferreira, Tony Cass and Vicente Traver.

We must thank the authors, whose research and development efforts are recorded here. We also thank the keynote speakers for their invaluable contribution and for taking the time to synthesize and prepare their talks. Finally, special thanks to all the members of the INSTICC team, whose collaboration was fundamental for the success of this conference.

June 2010

Ana Fred  
Joaquim Filipe  
Hugo Gamboa

# Organization

## Conference Co-chairs

Ana Fred	Technical University of Lisbon / IT, Portugal
Joaquim Filipe	Polytechnic Institute of Setúbal / INSTICC, Portugal
Hugo Gamboa	CEFITEC / FCT - New University of Lisbon, Portugal

## Organizing Committee

Patrícia Alves	INSTICC, Portugal
Sérgio Brissos	INSTICC, Portugal
Helder Coelhas	INSTICC, Portugal
Vera Coelho	INSTICC, Portugal
Andreia Costa	INSTICC, Portugal
Bruno Encarnação	INSTICC, Portugal
Bárbara Lima	INSTICC, Portugal
Elton Mendes	INSTICC, Portugal
Carla Mota	INSTICC, Portugal
Vitor Pedrosa	INSTICC, Portugal
José Varela	INSTICC, Portugal
Pedro Varela	INSTICC, Portugal

## HEALTHINF Program Committee

Osman Abul, Turkey	Cigdem Gunduz-Demir, Turkey
Adrian Barb, USA	Jin-Kao Hao, France
Bert-Jan van Beijnum, The Netherlands	Wen-Lian Hsu, Taiwan
Egon L. van den Broek, The Netherlands	Andreas Kerren, Sweden
Edward Brown, Canada	Georgios Kontaxakis, Spain
James Cimino, USA	Nicolas Loménie, France
Donald Craig, Canada	Goran Nenadic, UK
Adrie Dumay, The Netherlands	Chris Nugent, UK
Wendy Filsell, UK	George Sakellaropoulos, Greece
Jose Fonseca, Portugal	Niilo Saranummi, Finland
David Greenhalgh, UK	Barry Smith, USA
Tiago Guerreiro, Portugal	Adrian Tkacz, Poland
	Aristides Vagelatos, Greece

## **HEALTHINF Auxiliary Reviewer**

Frans Van der Sluis, The Netherlands

## **BIODEVICES Program Committee**

Oliver Amft, The Netherlands  
Rodrigo Varejão Andreão, Brazil  
Enrique A. Vargas Cabral, Paraguay  
Fernando Cruz, Portugal  
Pedro Pablo Escobar, Argentina  
Marcos Formica, Argentina  
Juan Carlos Garcia Garcia, Spain  
Jongin Hong, UK  
Leonid Hrebien, USA  
Bozena Kaminska, Canada  
Rui Lima, Portugal

Paulo Mendes, Portugal  
Raimes Moraes, Brazil  
Pedro Noritomi, Brazil  
Laura Papaleo, Italy  
Joaquin Roca-Dorda, Spain  
Mario Sarcinelli-Filho, Brazil  
Jorge Vicente Lopes da Silva, Brazil  
Milutin Stanacevic, USA  
Alexandre Terrier, Switzerland  
Peter Yves-Alain, Canada

## **BIODEVICES Auxiliary Reviewers**

Getúlio Igrejas, Portugal

Evandro Salles, Brazil

## **BIOSIGNALS Program Committee**

Oliver Amft, The Netherlands  
Egon L. van den Broek,  
The Netherlands  
Tolga Can, Turkey  
Francis Castanie, France  
Kuo-Sheng Cheng, Taiwan  
Suash Deb, India  
Julien Epps, Australia  
Adam Gaweda, USA  
Cigdem Gunduz-Demir, Turkey  
Jiri Holcik, Czech Republic  
Jonghwa Kim, Germany  
Gunnar W. Klau, The Netherlands  
Georgios Kontaxakis, Spain  
Narayanan Krishnamurthi, USA  
Arjan Kuijper, Germany  
Christos Loizou, Cyprus  
Marco Loog, The Netherlands  
Mahdi Mahfouf, UK  
Martin Middendorf, Germany  
Mihaela Morega, Romania

Tadashi Nakano, Japan  
Hasan Ogul, Turkey  
Ming Ouyang, USA  
Perrine Paul, Ireland  
Marcos Rodrigues, UK  
Jianhua Ruan, USA  
Carlo Sansone, Italy  
Andres Santos, Spain  
Gerald Schaefer, UK  
Dragutin Sevic, Serbia  
Jose Silvestre Silva, Portugal  
Iryna Skrypnyk, Finland  
Nikolai V. Slavine, USA  
Asser Tantawi, USA  
Gianluca Tempesti, UK  
Hua-Nong Ting, Malaysia  
Duygu Tosun, USA  
Bart Vanrumste, Belgium  
Yuanyuan Wang, China  
Didier Wolf, France



**BIOSIGNALS Auxiliary Reviewers**

Miguel Angel Luengo-Oroz, Spain

Emanuela Marasco, Italy

**BIOINFORMATICS Program Committee**

Osman Abul, Turkey

R. Acharya, USA

Tatsuya Akutsu, Japan

Reda Alhajj, Canada

James Bailey, Australia

Carlos Brizuela, Mexico

Egon L. van den Broek,

The Netherlands

Chris Bystroff, USA

Tolga Can, Turkey

Zehra Cataltepe, Turkey

Kun-Mao Chao, Taiwan

Ronald Cornet, The Netherlands

Antoine Danchin, France

Thomas Dandekar, Germany

Ugur Dogrusöz, Turkey

George Eleftherakis, Greece

Burak Erman, Turkey

Max H. Garzon, USA

Preetam Ghosh, USA

Alejandro Giorgetti, Italy

Julian Gough, UK

Attila Gürsoy, Turkey

Bianca Habermann, Germany

Joerg Hakenberg, USA

Hung-Chung Huang, USA

Liang-Tsung Huang, Taiwan

Yuki Kato, Japan

Sami Khuri, USA

Gunnar W. Klau, The Netherlands

Sophia Kossida, Greece

Zoe Lacroix, USA

Sang Yup Lee, Korea, Republic of

Xiaoli Li, Singapore

Stefano Lonardi, USA

Chin Lung Lu, Taiwan

Gabriel Luque, Spain

Elena Marchiori, The Netherlands

John McCall, UK

Burkhard Morgenstern, Germany

Radhakrishnan Nagarajan, USA

Jean-Christophe Nebel, UK

Kay Nieselt, Germany

Matteo Pellegrini, USA

Fabio Rinaldi, Switzerland

Eric Rivals, France

Dietmar Schomburg, Germany

Silvio C.E. Tosatto, Italy

Jean-Philippe Vert, France

Juris Viksna, Latvia

Thomas Wilhelm, UK

Yanbin Yin, USA

Jingkai Yu, China

Erliang Zeng, USA

Louxin Zhang, Singapore

**BIOINFORMATICS Auxiliary Reviewers**

Wen-Chi Chou, USA

Qin Ma, USA

## **Invited Speakers**

Peter D. Karp, Director Bioinformatics Research Group, Artificial Intelligence Center, USA

Rui M. C. Ferreira, National Coordinator for Cardiovascular Diseases, Office of the High Commissioner for Health, Portugal

Tony Cass, Institute of Biomedical Engineering, Imperial College London, UK

Vicente Traver, ITACA, Universidad Politécnica de Valencia, Spain

# Table of Contents

## Invited Paper

Personal Health: The New Paradigm to Make Sustainable the Health Care System . . . . .	3
<i>Vicente Traver and Raquel Faubel</i>	

## Part I: Health Informatics

Registration and Fusion Techniques for Medical Images: Demonstration and Evaluation . . . . .	15
<i>Evanthia Faliagka, George Matsopoulos, Athanasios Tsakalidis, John Tsaknakis, and Giannis Tzimas</i>	
Prediction of Pancreatic Cancer Survival through Automated Selection of Predictive Models . . . . .	29
<i>Stuart Floyd, Carolina Ruiz, Sergio A. Alvarez, Jennifer Tseng, and Giles Whalen</i>	
Web of Active Documents: An Architecture for Flexible Electronic Patient Records . . . . .	44
<i>Federico Cabitza and Iade Gesso</i>	
A Personalization Environment for Multi-version Clinical Guidelines . . .	57
<i>Fabio Grandi</i>	
Dynamic Routing Using Health Information Policy with Apache Camel . . . . .	70
<i>Edward Brown and Jamie Goodyear</i>	
A Personal Health Record System for Emergency Case Management . . . .	83
<i>Vassiliki Koufi, Flora Malamateniou, and George Vassilacopoulos</i>	
A Method for Representing and Querying Temporal Information in OWL . . . . .	97
<i>Martin J. O'Connor and Amar K. Das</i>	
Designing a Telemedical System for Cardiac Exercise Rehabilitation . . .	111
<i>Florian Klompf, Clemens Busch, Karsten Nebe, Andreas Bleiker, and Detlev Willemsen</i>	
A Value and Model Driven Method for Patient Oriented KPI Design in Health Care . . . . .	123
<i>Erik Perjons and Paul Johannesson</i>	

**Part II: Biomedical Electronics and Devices**

Encapsulated Multi-vesicle Assemblies of Programmable Architecture:  
Towards Personalized Healthcare . . . . . 141  
*Maik Hadorn and Peter Eggenberger Hotz*

Adaptive RBF-HMM Bi-Stage Classifier Applied to Brain Computer  
Interface . . . . . 152  
*José Luis Martínez Pérez and Antonio Barrientos Cruz*

Thermal Characteristics of Microcantilever Biosensors . . . . . 166  
*Mohd. Zahid Ansari and Chongdu Cho*

Prosodic Speech Restoration Device: Glottal Excitation Restoration  
Using a Multi-resolution Approach . . . . . 177  
*Olaf Schleusing, Rolf Vetter, Philippe Renevey,  
Jean-Marc Vesin, and Valérie Schweizer*

An Investigation on ALD Thin Film Evanescent Waveguide Sensor for  
Biomedical Application . . . . . 189  
*Agung Purniawan, Paddy French, Gregory Pandraud, and  
Pasqualina M. Sarro*

**Part III: Bio-inspired Systems and Signal Processing**

Probabilistic Patient Monitoring with Multivariate, Multimodal  
Extreme Value Theory . . . . . 199  
*Samuel Hugueny, David A. Clifton, and Lionel Tarassenko*

Weighted SVMs and Feature Relevance Assessment in Supervised  
Heart Beat Classification . . . . . 212  
*Gael de Lannoy, Damien François, Jean Delbeke, and  
Michel Verleysen*

The Investigation and Clinical Significance of Resonance in the Heart  
Rate and Vascular Tone Baroreflexes . . . . . 224  
*Eogeny G. Vaschillo, Bronya Vaschillo, Jennifer F. Buckman,  
Robert J. Pandina, and Marsha E. Bates*

Glottal Inverse Filtering of Speech Based on Homomorphic Prediction:  
A Cepstrum-Based Algorithm not Requiring Prior Detection of Either  
Pitch or Glottal Closure . . . . . 238  
*Rubén Fraile, Malte Kob, Juana M. Gutiérrez-Arriola,  
Nicolás Sáenz-Lechón, J. Ignacio Godino-Llorente, and  
V́ctor Osma-Ruiz*

Texture Discrimination with Artificial Whiskers in the Robot-Rat  
Psikharpax . . . . . 252  
*Steve N’Guyen, Patrick Pirim, and Jean-Arcady Meyer*

Heart Rate Variability and the Acceleration Plethysmogram Signals Measured at Rest .....	266
<i>Mohamed Elgendi, Mirjam Jonkman, and Friso DeBoer</i>	
Emotional Face Perception in Healthy Subjects and Parkinson's Disease: An Effective Connectivity Study .....	278
<i>Elvis Lira da Silva, Gabriela Castellano, João Ricardo Sato, Ellison Fernando Cardoso, and Edson Amaro Jr.</i>	
Biometric Authentication with Electroencephalograms: Evaluation of Its Suitability Using Visual Evoked Potentials .....	290
<i>André Zúquete, Bruno Quintela, and João Paulo Silva Cunha</i>	
<b>Part IV: Bioinformatics</b>	
Reconstructing the Phylogenetic Relationships of the Cyclophyllidean Cestodes: A Case Study Using ITS2 rDNA and Sequence-Structure Alignment .....	309
<i>Veena Tandon, Devendra Kumar Biswal, Pramod Kumar Prasad, and Chenkual Malsawmtluangi</i>	
<i>ReHap</i> : A Framework for Single Individual Haplotyping from Next-Generation Sequencing Data.....	322
<i>F. Geraci and M. Pellegrini</i>	
Inference of Large Phylogenies Using Neighbour-Joining .....	334
<i>Martin Simonsen, Thomas Mailund, and Christian N.S. Pedersen</i>	
Prediction of Minimum Free Energy Structure for Simple Non-standard Pseudoknot .....	345
<i>Thomas K.F. Wong and S.M. Yiu</i>	
Intra- and Inter-Molecular Coevolution: The Case of HIV1 Protease and Reverse Transcriptase .....	356
<i>Patrick Boba, Philipp Weil, Franziska Hoffgaard, and Kay Hamacher</i>	
Structured Output Prediction of Novel Enzyme Function with Reaction Kernels .....	367
<i>Katja Astikainen, Liisa Holm, Esa Pitkänen, Sandor Szedmak, and Juho Rousu</i>	
A Parallel Algorithm for Counting Subgraphs in Complex Networks . . . .	380
<i>Pedro Ribeiro, Fernando Silva, and Luís Lopes</i>	
Efficient Searching for Motifs in DNA Sequences Using Position Weight Matrices .....	394
<i>Nikola Stojanovic</i>	
<b>Author Index</b> .....	407

# **Invited Papers**

# **Personal Health: The New Paradigm to Make Sustainable the Health Care System**

Vicente Traver and Raquel Faubel

Instituto de Aplicaciones de las Tecnologías de la Información y Comunicaciones Avanzadas  
(ITACA), Universidad Politécnica de Valencia  
Camino de Vera sn, ed. 8G, 46022 Valencia, Spain  
vtraver@itaca.upv.es

**Abstract.** New technologies and innovation are leading a new paradigm in health care systems in order to face the growing demand and expectations in promoting and maintaining health and in disease prevention, treatment and care. ICT could have several applications in health scope (care, management, information and training) being an enabler for a personalized health system. RFID system, sensors, internet 2.0, electronic health records and ambient intelligence are some tools for catalyzing changes in health systems towards Personal Health.

**Keywords:** Personal health, e-Health, Telemedicine, ICT.

## **1 Introduction**

In the last few years, health systems are facing a growing demand and expectations in promoting and maintaining health and in disease prevention, treatment and care. In order to confront such demand, new routes must be found to provide a greater number of complex patient-centered services. Experts indicate that a radical transformation of the health services supply process is required [1, 2]. This new approach, called personal health (pHealth), must be focused on health in a broader sense, not just on the treatment, incorporating information and communication technologies (ICT) and recognizing the influence of the patient as a service consumer. This implies a shift of paradigm in which ICT and patient empowerment will be drivers towards a sustainable and patient-centered health systems. Under the personal health, the patient, as a consumer of health services, acts more as a client, and therefore demands the corresponding rights and obligations; care task will be focused on promoting and preventing rather than treatment; finally, disease management will be converted towards a holistic concept of health and life management. In this new model, ICT is considered a tool towards personalized treatments, sustainable health services, better quality of service, closed loop relationship and evidence based medicine.

This chapter aims to approach to pHealth and ICT in the healthcare scope. The latest contributions in ICT as RFID, sensors, ambient intelligence (AmI), social health media - health 2.0- and barriers and initiatives to achieve interoperability between different information systems are other relevant aspects of the application of ICT in health that will be held during this chapter.

## 1.1 Telemedicine and Health

Telemedicine was defined by the World Health Organization (WHO) in 1997, as the delivery of healthcare services, where distance is a critical factor, by all healthcare professionals using information and communication technologies for the exchange of valid information for diagnosis, treatment and prevention of diseases and injuries, research and evaluation, and for the continuing education of healthcare providers, all in the interest of advancing the health of individuals and their communities [3].

According to this concept, telemedicine could have several applications in health scope. It can be useful to inform the population about health issues, for the training and information for healthcare professionals, to support for the continuum of care (management) and, finally, in purely assistance processes either for diagnosis, treatment and care.



**Fig. 1.** Telemedicine applications

e-Health policy in the EU is part of the 'Strategy Europe 2020' with the facing goal of achieving in the European Union, in 2010, an economy based on knowledge, competitive and dynamic, with improvements in employment and social cohesion. This plan of action includes measures to obtain a more efficient supply of services for citizens (e-government, e-health e-learning and e-commerce), intervening authorities, generating demand and promoting the creation of new networks. In 2006, the European Commission published a new strategy to accelerate the growth of the e-health market in Europe, 'Accelerating the Development of the e-health Market in Europe' [4] which deals with the elimination of several legal/regulatory barriers that restrict the development of the e-health.



## 2 New Paradigm for Health Care Systems

### 2.1 Disruptive Innovation

Disruptive innovation is a combination of technological advances and new business models that has the potential to transform the structure, organization, and performance of the healthcare system as a whole. Christensen [5] introduced the theory of disruptive innovation asserting that consumers could produce transformations of the prevailing business models by purchasing new products and services that were more affordable and accessible than those that had been traditionally offered. For the disruptive innovation, two drivers are needed: technological enablers and a disruptive business model that can profitably deliver these routine solutions to customers in affordable and convenient ways.

Technologies drivers for disruptive innovation in healthcare are advances in medical knowledge and advances in ICT. The ICT advances not only support the development of new techniques and devices. They also facilitate the codification, continuous updating, and diffusion of therapies and best-practice care protocols. ICT can make feasible the integration of a patient's health care data to provide real-time medical decision support to professional, as well as to patients and their families. ICT also enables the collection, integration, and analysis of data on the performance of the overall system and supports the use of many advanced systems design, analysis, and governance tools and methods to improve system performance.

Accelerating the development and application of ICT tools and techniques would make possible the design, analysis, and governance of new processes and systems throughout the health care system [6]. Disruptive innovation, applied to healthcare system, implies affordable, simple and network system, comparing with the current healthcare system with complex and high cost alternatives based in an oligarchic system.

### 2.2 Chronic Disease Challenge

Chronic disease management represents one of the greatest challenges for health in Europe. Growth rates of the adult population and aging projections in Spain only serve to highlight the need to implement comprehensive strategies for changing the routines and resources tackling chronic diseases. In the EU, currently 86% of deaths are attributed to chronic illness, accounting for 50-80% of the total expenditure on health, as appropriate. Under these conditions, reports the WHO, those health systems maintaining their current system of disease management will be facing a large problem and they cannot keep the current level of service for the growing number of people under chronic conditions.

Considering this challenge and trying to fit health services and needs, the patients' stratification by risk levels allows an adequate adaptation to the populations' healthcare needs, giving them the most appropriate response. The Kaiser Permanente pyramid [7, 8] has proved to be a very useful method for conceptualizing risk stratified groups of patients with long term conditions. The pyramid models long-term conditions by splitting the population into those who can care for themselves, individuals who need help to manage their diseases, and people who require more intensive

case management. The upper level includes patients who need a more systematic follow up due to the complexity of their clinical situation. It implies 3-5 % of the population with chronic disease. The second level represents 15-25% of chronic patients and includes individuals with high risk of the disease progression, requiring healthcare delivery management. In the first level, we have the population diagnosed with a chronic disease, needing only supportive care to achieve a proper self-management of the disease.

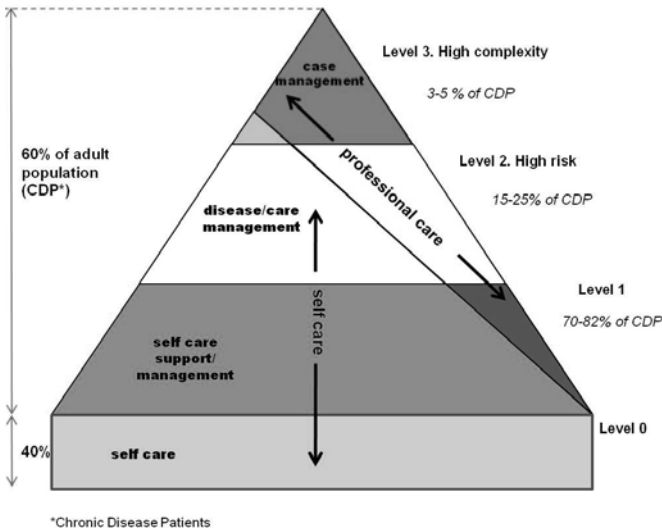


Fig. 2. Kaiser Permanent Pyramid

### 2.3 Personal Health Systems

Personal health systems assist in the provision of continuous, quality controlled, and personalized health services to empowered individuals, regardless of location. In 2002, WHO defined some key elements for this paradigm shift moving towards to the personal health [9] (table 1).

Table 1. Key Elements for pHealth

1.	Identify, score and stratify population
2.	Plan and coordinate care within all health care levels, using case management methodology
3.	Specific guidelines and protocols for each disease
4.	Specific educational disease programmes
5.	Integrated Information Systems, allowing a predefined Balanced ScoreBoard
6.	Align resources and incentives
7.	Evaluate and improve quality, cost and service

First of all, the population must be identified, scored and stratified according to the health risk. For every level, it's needed to plan and coordinate care, using case management methodology. According to the risk level, case management will be applied for patients with high complexity, disease management for the high risk patients and self-care support for the lower level of chronic patients. Specific guidelines and educational programs must be described for each disease. On the other hand, the integration of the information system can facilitate the establishment of a predefined Balanced Scoreboard. In this entire paradigm shift, resources and incentives must be aligned and it's fundamental to evaluate all the process trying to incorporate a continuous improvement system for quality, cost and service.

### **3 Information and Communication Technologies (ICT)**

Some technologies are mature enough to be used extensively, being drivers and enablers for the paradigm shift in healthcare. Identification by radio frequency (RFID), sensors to measure any physiological parameter in a simpler and more reliable way, web 2.0 to facilitate the exchange of information and experience between family and/or patients and health professionals, or electronic medical history are some ICT that can be key elements for the new personal health system.

#### **3.1 RFID**

RFID (Radio Frequency Identification) system is a technology that allows detecting and identifying an object, previously tagged with a smart label, through the information transmitted wirelessly by the tag. A RFID system has two main components. Firstly, a tag, that is attached to the object to be identified. It integrates onboard memory, an integrated circuit or machine state circuit capable to manipulate and store the data and a coupling element to communicate with the reader. It can be included sensors or a power supply (active tags). Secondly, a reader, that typically contains a control unit to manipulate the information and an antenna to communicate with tags, hence reading or writing their memories. The reader connects with a network so it can forward the identity of the tag as well as the information read to a centralized system [10].

RFID enables unique localization, tracking and management identification of any item or living being without contact or direct line of sight. In healthcare it can be applied to ensuring right identification of each patient and his corresponding medical data anywhere in the hospital, guarantying appropriate actions are granted to the right persons (drugs administration, surgical procedures,...) regardless the time and place, identifying lab samples (blood, urine, exudations, biopsies, etc), verifying authenticity and origin of drugs and checking their status and expiry dates while they are moved among different departments and sections, etc [11]. Additionally, inventory control of medical equipment or theft prevention of devices and medical utensils is also susceptible to be accomplished ubiquitously by means of RFID technology.

### 3.2 Sensors

Telemonitoring is a part of telemedicine that includes diagnosis, monitoring, treatment or education of a patient in a remote location. The effective performance of telemonitoring programs requires appropriate infrastructure involving technological innovation and effective communication systems to facilitate the acquisition, transmission, storage, processing and real-time viewing of data, sounds and images. For suppliers, these programs can provide cost savings because they control the evolution of chronic diseases and ensure the efficient use of health resources while maintaining the appropriate and quality care for patients. For consumers (patients), telemonitoring brings a sense of empowerment in their condition management.

Telemonitoring platforms are a group of components consisting of a terminal patient, -which could be several hardware solutions depending on the circumstances of use or application- a set of monitoring devices and software applications that include secure communications with centralized management system. In the market, there are several solutions, each one with its own characteristics. However, all these solutions have limitations of integration into broader services architecture as they are closed and with proprietary protocols.

In the coming years, telemonitoring skills will be improved by increasing diversity and capacity of the sensors (size, transmission capacity, power usage...), improvements in images capturing, telecommunications more affordable, increase in computing power, improvements in consumption (thermal efficiency, battery life ...). Many companies that have shown their ability in this discipline will be involved in this development as Philips Medical Systems, Telehealthcare Viterion, Biotronik, Aerotel Medical Systems, Card Guard AG.

### 3.3 Internet 2.0

The new health system is a patient-centered model [12, 13, 14], where the individual has greater responsibility on their health, either in lifestyle and in decision-making autonomy (patient empowerment). Empowerment implies a greater demand for information -much more relevant for people with chronic diseases- which is increasingly channeled through the Web.

Patients use the web 2.0 [15], to handle personal medical information (personal health records), access to health information and applications, knowledge and information dissemination, and also as a tool of socialization in virtual communities. We must also highlight continued and growing development of alternative devices allowing connection to internet as mobile phones or consumer electronics devices: MP3, game consoles, televisions...

However, it is becoming more complex to find relevant and quality contents due to the increased availability of health information. Several initiatives are being developed in order to avoid the risk of data overload like semi-automatic labeling based on semantic technologies, data mining, collaborative filtering and personalized search engines [16, 17].

### 3.4 Electronic Health Records

Currently, interoperability between different existing health information systems is a great relevance issue and not yet resolved. The most complicated, technically, is to achieve semantic interoperability between electronic health records systems (EHR). Already in 2008 the European Commission sets interoperability as one of the most relevant topics in recommendation COM (2008) 3282. This document provides a set of guidelines for the development and deployment of interoperable EHR systems that allow cross-border patient data exchange, a goal that, from a down-to-earth point of view, could take 20 years to attain. This document was materialized in the EpSOS project (Smart Open Services for European Patients) involving 12 European countries: Austria, Sweden, Czech Republic, Germany, Denmark, France, Greece, Italy, Netherlands, Slovakia, United Kingdom and Spain. EpSOS aims to develop a practical framework for eHealth and ICT infrastructure to guarantee the compatibility of the different health records systems. Thereby, health professionals could access electronically - in their own language - to the essential health data of a patient from another country and pharmacies could dispense prescriptions issued by other EU states.

Some key points must be taking into account in order to perform interoperability. The benefits of sharing patient information must be clear; the reliability of the processes should be increased to improve security and confidentiality of data; organizations must rely on those with whom they share information, and information exchange should not cause conflicts with business or legal interests of the participants. Many agencies are working on the standards specification in order to facilitate interoperability. One of these agencies is Continua Health Alliance [18] this non-profit organization is based on a coalition of health-related industry that was formed with the aim of promoting and fostering interoperability among telemedicine peripheral devices and data systems. It comprises specifications for devices to include its own profile, to define the data according to ISO / IEEE, to define the integration with health records (EHRs), HL7 and XDR profile specification.

### 3.5 Ambient Intelligence

Ambient Intelligence (AmI) [19] suggests a model of interaction between people and their surrounded context-sensitive environment responding adaptively to user needs and habits. Technologies are embedded in daily life elements and present in the environment in a nonintrusive way, invisible but perceptible, accessed via intelligent but simple and natural interfaces. Mobile devices (PDAs and third generation phones), sensors, in-door localization systems, middleware and embedded computing devices on user clothes are some of the devices that can be found in this kind of environment. Integration and interoperability between the different devices are essential and necessary to define a common communications protocol.

## 4 Conclusions

Health systems are undergoing a paradigm shift towards personal health and this new model requires comprehensive action, integrating all the stakeholders of health system. In this new model, ICT are a tool to face up the new challenges in the health

scope. This challenge facing Europe and Spain in the coming years is well represented in the 5 Challenge of the 7th Framework Programme of the European Commission for the 2011-2012: ICT for health, ageing well, inclusion and governance. This strategy focuses on the application of ICT for disease prediction, prevention, minimally invasive treatments, disease management and support for healthy lifestyles as well as technological solutions to prolong independent living and technologies for people with disabilities. A new model of healthcare is required; consequently, pHealth will happen and the only pending question is when. However, the changes necessary for the transformation sought depend not only on technology: are also related to other factors such as the introduction of new models of organization, professional roles, training plans or new stakeholders in the health system [20].

## References

1. Healy, J.C.: Integration and Informatics and Communication Technologies (ICT) in the EU national health systems: status and trends. *Swiss Medical Informatics* 52, 3–5 (2004), <http://www.sgmi-ssim.ch/smi/online/smi-52/index.html> (accessed July, 2010)
2. Committee on Quality of Health Care in America Institute of Medicine: Crossing the quality chasm: a new health system for the 21st century. National Academy Press, Washington (2001), <http://www.nap.edu/books/0309072808/html/index.html> [accessed July, 2010]
3. World Health Organization (WHO): Telehealth and Telemedicine will henceforth be part of the strategy for Health-for-All. Press release. WHO, Geneva (1997)
4. European Commission eHealth Taskforce Report: Accelerating the Development of the eHealth Market in Europe (2007), [http://ec.europa.eu/information\\_society/activities/health/docs/lmi-reportfinal-2007dec.pdf](http://ec.europa.eu/information_society/activities/health/docs/lmi-reportfinal-2007dec.pdf) [accessed July, 2010]
5. Christensen, C.: *The Innovator's Dilemma*. Harvard Business School Press, Cambridge (1997)
6. Christensen, C.M., Grossman, J.H., Hwang, J.: *The Innovator's Prescription: A Disruptive Solution for Health Care*. McGraw-Hill, New York (2009)
7. Wallace, P.J.: Physician involvement in disease management as part of the CCM. *Health Care Financ. Rev.* 27, 19–31 (2005)
8. Feachem, R.G., Sekhri, N.K., White, K.L.: Getting more for their dollar: a comparison of the NHS with California's Kaiser Permanente. *BMJ* 324, 135–141 (2002)
9. World Health Organization: Innovative care for chronic conditions: building blocks for action. WHO, Geneva (2002), <http://publications.paho.org/product.php?productid=837&cat=0&page=1> [accessed July, 2010]
10. Xiao, Y., Shen, X., Sun, B., Cai, L.: Security and privacy in RFID and applications in telemedicine. *IEEE Commun. Mag.* 44, 64–72 (2006)
11. Rashvand, H.F., Salcedo, V.T., Sanchez, E.M., Iliescu, D.: Ubiquitous wireless telemedicine. *IET Communications* 2, 237–254 (2008)
12. Aujoulat, I., d'Hoore, W., Deccache, A.: Patient empowerment in theory and practice: Polysemy or cacophony? *Patient Educ. Couns.* 66, 13–20 (2007)

13. Salmon, P., Hall, G.M.: Patient empowerment and control: a psychological discourse in the service of medicine. *Soc. Sci. Med.* 57, 1969–1980 (2003)
14. Brennan, P., Safran, C.: Empowered Consumers. In: Lewis, D., Eysenbach, G., Kukafka, R., Stavri, P., Jimison, H. (eds.) *Consumer Health Informatics: Informing Consumers and Improving Health Care*. Springer, New York (2005)
15. Fernández-Luque, L., Traver-Salcedo, V.: La nueva sanidad y los nuevos pacientes en la Web 2.0. *Diario Médico* (2010), <http://ht.ly/2akxb> [accessed July, 2010]
16. Eysenbach, G.: What is e-health? *J. Med. Internet. Res.* 18, 2 (2001), <http://www.jmir.org/2001/2/e20/> [accessed July, 2010]
17. Gaudinat, A., Joubert, M., Aymard, S., Falco, L., Boyer, C., Fieschi, M.: WRAPIN: new health search engine generation using UMLS Knowledge Sources for MeSH term extraction from health documentation. *Stud. Health. Technol. Inform.* 107, 356–360 (2004)
18. Continua Certification, Version 1.0. (2009), [http://www.continuaalliance.org/static/cms\\_workspace/Continua\\_Certification\\_Public.pdf](http://www.continuaalliance.org/static/cms_workspace/Continua_Certification_Public.pdf)
19. Aarts, E., Harwig, R., Schuurmans, M.: Ambient Intelligence. In: Denning, P.J. (ed.) *The Invisible Future: The Seamless Integration Of Technology Into Everyday Life*. McGraw-Hill Companies, New York (2001)
20. Pagan, M.: *Mobile and Wireless Systems Beyond 3G: Managing New Business Possibilities*. IRM Press, Hershey (2005)

**Part I**  
**Health Informatics**



# Registration and Fusion Techniques for Medical Images: Demonstration and Evaluation

Evanthia Faliagka<sup>1</sup>, George Matsopoulos<sup>2</sup>, Athanasios Tsakalidis<sup>1</sup>,  
John Tsaknakis<sup>3</sup>, and Giannis Tzimas<sup>3</sup>

<sup>1</sup> Computer Engineering and Informatics Department, University of Patras, Patras, Greece

<sup>2</sup> School of Electrical and Computer Engineering  
National Technical University of Athens, Athens, Greece

<sup>3</sup> Department of Applied Informatics in Administration and Economy  
National Technical University of Messolonghi, Greece

{faliagka, tsaknaki}@ceid.upatras.gr, gmatso@esd.ece.ntua.gr,  
{tsak, tzimas}@cti.gr

**Abstract.** In this work, we present an integrated system for the registration and fusion of medical images, named «dIaGnosis». It is often necessary to align medical images to illustrate the changes between the data retrieved at different times so as to assess the progress of a disease, or to assess the effectiveness of a treatment. The proposed system supports the storage, retrieval, distribution and presentation of medical images from different modalities, such as CT and MRI, in DICOM format. It also supports multiple examinations per patient and uses parallel processing threads to perform the processing of the acquired three-dimensional (3D) images in almost real time. In this paper, the architecture and the working environment of the implemented system are presented in detail, along with a pilot scenario that demonstrates the system in use. Additionally, the registration and fusion algorithms implemented are presented and evaluated, along with the image processing techniques used for the enhancement of medical images. The contribution of the proposed work is multilayered. It provides automatic matching methods based on both segmented surfaces and on different levels of gray, and it improves the alignment process when there is a relative movement and / or distortion of images in the data collected from different imaging systems.

**Keywords:** Medical systems, Image processing, Registration, Fusion.

## 1 Introduction

Medical imaging is a vital component of diagnostic medicine, and it also has a significant role in the areas of surgical planning and radiotherapy [1]. Often, medical images acquired in the clinical track are using different imaging technologies. Integrating these images, which are often complementary in nature, is a challenging problem. The first step in the integration process is bringing the tomographic images

into spatial registration, so that the same anatomical regions coincide, a procedure referred to as registration [2]. After registration, a fusion step is required in order to combine information from different modalities, or from the same modalities at different examination periods [3].

A prominent example where the fusion of registered images maximizes the available diagnostic information is tumor diagnosis and radiotherapy treatment. The Magnetic Resonance (MR) imaging system, the SPECT medical imaging and the Positron Emission Tomography (PET) provide functional information even at very early stages of cancerous tumors, but they do not reliably depict the anatomical characteristics of the tested organs. On the other hand, tomographic imaging techniques such as Computer Tomography (CT) and magnetic (MR) scanners, the ultrasound and X-rays provide anatomical information, but usually determine the existence of a cancer tumor only when it is in a later stage compared to the functional techniques. Thus, the combined use of different modalities that offers complementary clinical information is much more effective, allowing early diagnosis and accurate identification of a cancer tumor and hence the effective planning of the radiotherapy treatment.

It is often necessary to align medical data to illustrate the changes between the data retrieved at different times so as to assess the progress of a disease, or to assess the effectiveness of the treatment. In this case the fusion of data is implemented to illustrate the changes, as in the measurement of bone support for implants using dental radiographs. Moreover, the data registration applies to cases where data from anatomical atlases in conjunction with real clinical data and studies on patient populations are used.

In this work, a global alignment-fusion system of medical data was developed, which was named «dIaGnosis». Comparable software systems for processing and visualization of medical data are also implemented by Philips Medical Systems Inc., Siemens Medical Systems Inc and others. Medical data in commercial systems are represented in DICOM format, which is the prominent medical data protocol. Most commercial software provide semi-automatic and automatic registration options, as well as possibilities for data fusion after registration alignment, either on sections base (2D problem) or on surfaces base (3D problem). The proposed system overbalances the existing registration techniques. Specifically, it provides automatic matching based on both segmented surfaces and on different levels of gray, while algorithms are applied directly to three-dimensional (3D) data. In addition, it allows the application of different geometric transformations, including an elastic transformation to improve the registration when there is movement and / or distortion in the data collection of the patient from different imaging systems. Finally, it allows comparison of registration accuracy for the different techniques based on specific criteria to quantify registration.

This paper is organized as follows. Section 2 outlines the architecture of the system proposed and presents the algorithms used for image preprocessing, registration and fusion. Section 3 describes the working environment of the implemented system. Section 4 presents the baseline scenario where most of the procedures supported by the system are shown. The efficiency of registration techniques was tested during the pilot study on skull patient data collected from CT and MR scanners.

## 2 System Architecture

In Fig. 1 the overall system architecture is depicted. The system consists of five complementary layers-subsystems, managing the registration and fusion of the medical data, as well as the interaction with the final user. According to the proposed architecture, the system consists of the following five subsystems.




Fig. 1. The system architecture

### 2.1 Data Collection and Data Management Subsystem

The subsystem of data collection and management allows the storage, retrieval, distribution and presentation of medical images:

- Using Magneto-optical instrument, and
- Data transfer via network from the diagnostic consoles of the CT and MR scanners, or a workstation where digital medical data are acquired.

The medical data collected are a series of sections from the same patient from different imaging modalities (CT and MR scanners) and correspond to a specific region of the human body. In the pilot version of the system the data correspond to the region of the skull as acquired from both CT and MR scanners.

As mentioned earlier, DICOM format is the prominent international protocol for medical data. Thus, the medical data acquired by the scanners are compatible with this format. The entry, management and export data are in DICOM format too [4].

In this subsystem we implemented a function that reads the header of the DICOM file (DICOM header) and includes automatically the following technical characteristics of the system: the number of sections, the number of pixels per section, a data analysis per section (mm/pixel), the interval sections (mm), the number of bits/pixel and the patient data (patient code, DATE examination, etc.), if available.

The basic technical capabilities of the subsystem include:

- Patient (code) correspondence with the initial data of his/her examination.
- Multiple examinations per patient (through appropriate code).
- Data display with multiple horizontal sections in icon size.
- Data storage after their process (in DICOM or other format) in the hard disk of the computer system.
- Determination of reference data from the user.
- Ability to support multiple data to align common reference data.

## 2.2 Medical Data Preprocessing Subsystem

The data preprocessing is an optional step. It applies to data which are characterized by high levels of noise and the containment is achieved by using the appropriate filters. So, it is usual that before the registration a re-sampling of one or both data sets that have the same discretionary analysis is needed. Thus in the subsystem an appropriate technique for re-sampling is incorporated [5]. The data pre-processing subsystem includes the segmentation technique as developed. In this case, anatomical information is extracted from the two data sets (for example the external surfaces of the skull from both CT and MR scanners), which is then used to perform the registration.

### Pre-Processing Techniques

The acquired 3D data may include noise and/or characterized by heterogeneous background. This noise is undesirable and should be removed, without the loss of significant anatomical information contained in images. For noise reduction, suitable filters are implemented to improve the quality of images, which are applied on section based (two-dimensional problem - 2D) [6].

Specifically, within the subsystem the following filters have been implemented to improve image quality:

- Mean filter: It is a low-pass filter which reduces high-frequency noise in an image.
- Median filter: it is another filter for noise containment.
- Gamma correction: The factor  $\gamma$  determines the function which distributes the values of pixels, according to the intensity of brightness of the screen. The factor  $\gamma$  is equal to one when there is a linear relationship between pixel values and intensity of brightness. Images that appear darker usually require the factor  $\gamma$  have values larger than one, while those which appear bright usually require the factor  $\gamma$  have values smaller than one.
- Histogram Equalization filter: it is a commonly used technique for better visualization of the diagnostic information of an image. In cases where the biological tissue of interest shows rates (different levels of gray), which vary between certain limits in the digital image, the visualization of the tissue is significantly enhanced if the function which corresponds the values of pixels in the image with brightness in screen changes.
- Adjust brightness and contrast: It is one of the most basic functions for image editing. The implementation of this subsystem provides the opportunity to change the brightness and contrast of images by the simple linear transformation:


$$I'(x, y) = aI(x, y) + b \quad (1)$$

Where  $I(x,y)$  is the pixel of the initial image with coordinates  $(x,y)$  and  $I'(x,y)$  is the pixel of the adjusted image.

### Medical Data Segmentation

This technical requirement is particularly important as it allows the automatic segmentation of the human body structure (s) of the same patient as shown in Fig. 2. Within the project, appropriate algorithms were developed and integrated for structures segmentation using:

- Canny edge detector
- 3D region growing operator: According to this method [7], one or more pixels are initially used  $(x_i, y_i, z_i)$  that are considered as seed points of the region  $R_i = \{(x_i, y_i, z_i)\}$ . Then, the gray level for each seed-point is compared with the gray level of each of its neighboring pixels of the region  $R_i$ . The same procedure is repeated for all the neighboring pixels of the pixels that belong to  $R_i$ . The procedure stops when no more pixels can be added to the region  $R_i$ . The result of applying the above algorithm is homogeneous regions of pixels. Areas consisting of a small number of pixels are incorporated into neighboring regions with similar levels of gray.



**Fig. 2.** Example of MRI image segmentation. (a) Original image. (b) External contour using the 3D region growing operator.

### 2.3 Medical Data Alignment Subsystem

In many cases in the current clinical practice it is desirable to combine information provided by two or more imaging modalities or to monitor the development of a treatment based on data collected at different times by the same modality. In particular, when monitoring the development of a treatment, it is very often the imaging anatomical structures displayed in two sets of data that have been collected at different times to be characterized by geometrical movements, revolutions, etc. It is necessary to find an appropriate geometric transformation, which achieves the spatial coincidence of anatomical structures of the two images. This process of finding the transformation is called registration.

The medical data alignment subsystem consists of a set of techniques for 3D registration of brain data on surface based or using the levels of gray (gray-based). Particular attention has been given to the design of the automatic registration techniques. Alternatively, there is the option of manual registration using appropriate surface driving points as selected by the expert.

Within the design of this subsystem three registration techniques were implemented:

- Automatic registration based on surfaces,
- Automatic registration based on gray levels and
- Manual registration.

### **Surface-Based Registration Technique**

This technique is automatic and based on the spatial matching of segmented anatomical structures of data from different imaging modalities [8].

The basic stages of the automatic method for surfaces registration include:

- **Surface Pre-alignment.** The stage of pre-alignment includes the spatial displacement of two triangulated surfaces, so that the centers of mass coincide. Also, a transformation of scale in each axis is done separately, based on the voxel sizes of the two images [9].
- **Geometric transformation application.** The second phase implements an overall geometric transformation. Its parameters are calculated by optimizing a function that quantifies the spatial matching of a triangulated surface of the reference image (computer tomography - CT) and the modified image (MRI - MRI). Four models of geometric transformation in three dimensions are explored and evaluated based on the final results of the registration [10].
- **Matching function definition.** The registration can be seen as the optimization of a Measure of Match - MOM according to the variables of the selected transformation. At the case of surfaces matching an appropriate matching function is the average of the geometric distance between the transformed points of the magnetic scanner data and the corresponding closest points of computer scanner data.

### **Registration Based on Gray Levels**

This data registration technique is based on the automatic spatial identification of data from different imaging systems and is applied on image values directly, without the prior requirement for segmenting common anatomical structures [11].

### **Manual Registration**

In the case of the manual registration method, the expert selects points in the respective sections of the two imaging modalities and the registration of the data is based on the selection of a particular geometric transformation [12]. This method has been developed so that its performance can be compared to the performance of the proposed automatic registration methods.

## **2.4 Medical Data Fusion Subsystem**

Medical data fusions scope is to combine information from different modalities, after the application of the medical image registration process. The fusion subsystem is designed appropriately to allow the composition of anatomical information from the aligned medical data using techniques such as the pseudo-colour scale, logic functions for the diverse overlay of image parts on another image and change the degree of transparency in the overlay of anatomical structures [13].

### **Fusion Techniques**

Within the proposed system, the following techniques for medical data fusion, were developed and applied after registration:

- **Implementation of logical functions for the diverse parts overlay of one image on the other.** Specifically, after the data registration, the anatomical information derived

from data of the CT Scanner overlays on the respective aligned sections of the MR scanner in order to fuse the information from the two imaging systems. This process is mathematically standardized with the logical operator Exclusive Or (XOR), which is implemented as follows:

$$I(x, y) = I_A(x, y)(1 - M(x, y)) + I_B(x, y)M(x, y) \quad (2)$$

where  $I_A$  and  $I_B$  the reference image and the image to align respectively and  $M$  is the mask that has value 1 at the pixels that overlay from both the  $I_B$  to the  $I_A$ . The mask  $M$  may be the segmented structure of interest of the  $I_B$  image that has to be visualized from the reference system of  $I_A$ , or repeated normalized geometric shapes, where the aim is to visually confirm the accuracy of the registration of  $I_B$  relatively to  $I_A$ .

In the pilot study, information from the CT scanner was isolated and was inserted in the aligned data of the magnetic scanner using logic functions (XOR). This fusion method allows the expert – a doctor to assess the accuracy of the registration method, while it gives information on the position of the bones from the computer tomography in comparison with other soft tissues or tumors, as shown in MR.

- Ability to change the degree of transparency ‘ $a$ ’ during the overlay of anatomical structures in order to achieve a combination of information and assess the quality of the registration result. This technique was implemented on the basis of the relationship:

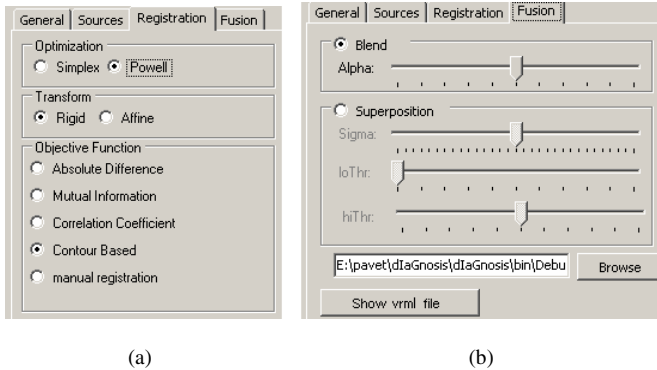
$$I(x, y) = I_A(x, y)(1 - a) + aI_B(x, y) \quad (3)$$

- Fusion of data in 3D is a particularly difficult problem because the extra dimension makes the data display difficult even without the extra complexity of the data fusion. In this work, a simultaneous demonstration of common anatomical structures - surfaces before and after the registration is achieved using the proposed representation techniques in the form of VRML.
- The results are visualized using pseudo-coloring according to the medical system used (e.g. red for the visualization of the anatomical structure of the axial scanner and blue for the magnet scanner), to make clear to the expert the relative position of the two surfaces and the change before and after the registration [14].

## 2.5 Graphical Interface Subsystem

The graphical interface subsystem is an important part of the developed system, as it allows the final user use the necessary functions of the registration software. The subsystem was developed having in mind the following requirements:

- Ease of use and user friendliness,
- Speed enforcement functions and
- Reliable performance of the software’s individual applications



**Fig. 3.** (a) Selection of registration parameters (b) Selection of the fusion parameters and implementation of the fusion settings


The key features of the graphical interface subsystem include:

- **Creation of an appropriate graphical environment:** This feature concerns the design and development of an appropriate interface that offers: a) easy navigation to the software's menus, b) easy access to medical data and c) a clear definition of the integrated techniques – algorithms as shown in Fig. 3.
- **Visualization of the medical information:** An important feature of the software is the ability to visualize data and present the results of the applications and techniques applied in a comprehensible manner. It provides: a) visualization of the original medical data, b) visualization of the data processing results, particularly the results of automatic registration methods and c) visualization of the fusion of information from the registration. In particular it allows simultaneous display of relevant medical data (e.g. CT and MR sections) before and after the registration and presentation of the fusion results. **Quantification of the registration results:** Beyond the visualization of the registration and fusion results another important feature is the presentation of quantitative data. The data are related to registration results based on a) specific success criteria and b) on geometrical differences (displacements and rotations) of the data to align from the reference data.

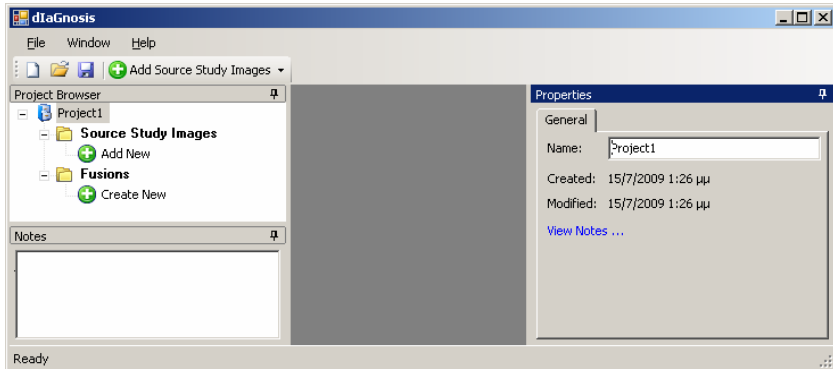
### 3 Working Environment

The user can use the basic components and navigate to the input and output data using a tree structure (Fig. 4). The tree structure starts from the node of the project. The project is the main component of the system. A project consists of source images, processing settings and output images and can be saved and retrieved at will without losing the settings of the user. It is the root of the tree that represents, while the intermediate nodes and leaves of the tree represent individual project data or processing information.





**Fig. 4.** System information data



**Fig. 5.** Starting the application

There are also collection nodes, which group different snapshots of similar information existing in each project. Each node, depending on the information that represents may correspond to a dynamic description, have associated notes, contain interfaces presenting information and have properties that are processed by the user and others. New nodes can be created by the user and added to the project while some existing may be removed.

In Fig. 5 the general working environment of the application is shown. The working environment is dynamic. The user interfaces can be aggregated into tabs, to activate the automatic concealment within the window, to match them all together and more. The user options are saved by closing the application, and retrieved the next time booted. Also some templates of the user interface are created and the user can easily select the one he likes.

The images can be loaded either from an existing list or with the process of surveying examination (Fig. 6).

As the recovery process of the examining image from DICOM files may be slow, the system makes the process to use parallel processing threads. During the information retrieval the system notifies the user about the status of recovery and does not allow access to the node's data.




Fig. 6. Image loading with the process of surveying examination

## 4 Pilot Scenario

In order to have an exhaustive testing of the system a testing scenario was defined. This scenario uses all the processes supported.


Specifically, the system was installed and evaluated by an expert-radiologist on the credibility of the operation and performance of the registration techniques.

Data sets from axial (CT) and magnetic tomography (MRI) from 5 patients from Strahlenklinik of the Stadtische Kliniken Offenbach of Germany were used. The axial tomography data were the «Reference data», while the magnetic tomography data were the «Data to align».

After any registration technique an overlay - fusion of the CT data on the corresponding sections of the MRI data took place. In this way, the expert assessed optically the performance of the registration techniques.

Fig. 7 shows characteristic results of the registration-fusion techniques using real medical data. Based on an analysis of these results we came to the following conclusions:

- The performance of the automatic registration techniques is much better compared to the semi-automatic alignment technique.
- Among the automatic registration techniques based on gray levels, the technique of mutual information has better performance compared to the technique using the correlation coefficient.
- The technique of surface registration is worse compared to the technique using mutual information and is almost equivalent to the technique using the correlation coefficient.



**Fig. 7.** Registration results using (a) Automatic registration using correlation coefficient (b) Automatic surfaces registration (c) Automatic registration using mutual data coefficient (d) Semi-automatic registration technique

#### 4.1 Quantitative Analysis of Registration Results

The four registration techniques implemented within the proposed system were further quantitatively evaluated in terms of accuracy. Towards this direction, five patients were used forming five pairs of sets, each set consisting of CT and MRI head data of the same patient. The accuracy of each registration technique was measured as the mean distance of the centers of all the external markers for each data set before and after registration (in pixels). The centers of the external markers were obtained manually by an experienced radiotherapist. Comparisons on the performance of these registration techniques based on this criterion are shown in Table 1.

From the quantitative result in Table 1 it is shown that all automatic techniques were performed better than the semi-automatic technique. Furthermore, the mutual information registration technique was outperformed from the other two automatic registration techniques. Finally, the surface and the correlation coefficient registration techniques were performed equivalently.

**Table 1.** Registration Techniques comparison

<i>Data Sets</i>	<i>Registration Techniques</i>			
	<i>Automatic Mutual Information</i>	<i>Automatic Correlation Coefficient</i>	<i>Automatic-Surface Registration</i>	<i>Semi-automatic Registration</i>
<i>Set 1</i>	$0.27 \pm 0.01$	$0.59 \pm 0.02$	$0.60 \pm 0.03$	$1.89 \pm 0.37$
<i>Set 2</i>	$0.28 \pm 0.02$	$0.67 \pm 0.03$	$0.63 \pm 0.01$	$1.87 \pm 1.21$
<i>Set 3</i>	$0.32 \pm 0.01$	$0.80 \pm 0.02$	$0.49 \pm 0.07$	$1.33 \pm 0.07$
<i>Set 4</i>	$0.31 \pm 0.01$	$0.61 \pm 0.05$	$0.50 \pm 0.01$	$0.47 \pm 0.00$
<i>Set 5</i>	$0.29 \pm 0.01$	$0.46 \pm 0.02$	$0.43 \pm 0.00$	$0.53 \pm 0.04$

## 4.2 Three-Dimensional Display of Anatomic Structures

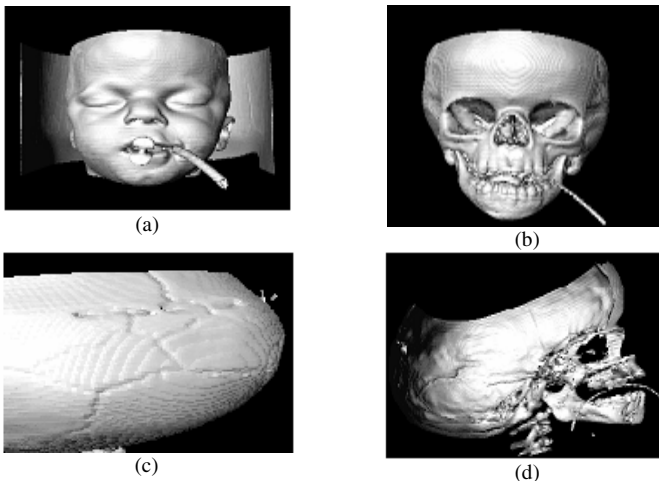
An important factor in the process of medical data fusion is the ability of the system to visualize the results of the registration. Specifically, the system supports the display of the anatomical structures - three-dimensional surfaces - before and after the registration.

The basic method followed to visualize the surface of the CTI and MRI is the technique of Surface Rendering. The surface rendering technique is based on the surface of interest detection based on gray levels. It constructs a set of polygons that is based on neighborly relations between the points forming the surface. The Marching Cubes algorithm is more representative of this category and is used in this application due to the rapid and realistic organs representation such as the skull.

Images were visualized as a 3D surface, using a well known surface rendering algorithm called Marching Cubes [15]. The DICOM images that consist of MRI or CT sectors, are stored as three-dimensional scalar fields, where scalar values represent shades of gray. The algorithm divides the image in imaginary cubes, taking eight neighbor locations at a time and then determines the polygons needed to represent the part of the iso-surface that passes through this cube. The iso-surface is an area of the image with a constant value of gray level (iso-value).

This basic region is called cube and its tops are equal to 1 if the corresponding value of the image gray level is greater than or equal to the iso-value, otherwise 0. The model created in this way is called the topology of the cube. All possible topologies of a cube are 256, as the levels of gray, and in each one of them the iso-surface is pre-triangulated.

In this system, the algorithm was applied to skull binary data from CT and MRI, which will occur after the implementation of the pre-processing and segmentation algorithms. A result of a three-dimensional skull reconstruction of CT and MRI sectors, using the above algorithm is shown in Fig. 8.




**Fig. 8.** Three-dimensional skull reconstruction of CT and MRI images

Furthermore, an indicative visual result of the overlay of the axial and magnetic scanner surfaces is presented using VRML and a surface representation algorithm.

It may be noted that the pre-aligned skin surfaces are different and the area of the axial tomography is external and above the area of the MRI. With the method of surfaces registration the registration between two surfaces is enhanced, as shown by the alternation of the two colors of the surfaces.

In Fig. 9 we can see the overlay of the skin surface of the axial tomography (red color) on the corresponding surface of the magnetic tomography (blue color) for a specific couple using the algorithm for surface representation in VRML format.



**Fig. 9.** (a) Before registration (b) Registration using the surface registration method

## 5 Conclusions

In this paper we have illustrated our registration-fusion system in detail, described the algorithms used and shown the basic scenario of the application's usage. Diagnosis is an integrated environment that facilitates the automatic matching based on both segmented surfaces and on different levels of gray and it allows comparison of registration accuracy for the different techniques based on specific criteria to quantify registration. It also improves the registration in case of movement and / or distortion in the data collection of the patient from different imaging systems.

After the implementation of the system, a number of tests were performed for evaluating the developed registration techniques both qualitatively and quantitatively in order to test the stability and accuracy of the techniques. As for future work, we plan to extend our system by developing further fusion and registration techniques. Additionally, more tests will be conducted to support the efficiency of the implemented system.

## References

1. Maintz, J.B.A., Viergever, M.A.: A survey of medical image registration. *Medical Image Analysis* 2(1), 1–36 (1998)
2. Hajnal, J.V., Hill, D.L.G., Hawkes, D.J. (eds.): *Medical image registration*. CRC Press, Boca Raton (2001)
3. Hawks, D.J., Hill, D.L.G., Bracey, E.C.M.L.: Multi-modal data fusion to combine anatomical and physiological information in the head and heart. In: Reiber, J.H.C., van der Wall, E. (eds.) *International Symposium of Computing in Nuclear Medicine and Cardiac MRI*. Kluwer, The Hague (1991)

4. Digital Imaging and Communications in Medicine (DICOM) Part 1: Introduction and Overview. National Electrical Manufacturers Association (2006)
5. Unser, M., Aldroubi, A., Eden, M.: B-Spline.: Signal processing: Part II – Efficient design and applications. *IEEE Trans. Signal Proc.*, 834-83-1000 (1993)
6. Gonzalez, R.C., Woods, R.E.: *Digital image processing*, 2nd edn. Addison Wesley, Reading
7. Zucker, R.: Region growing: childhood and adolescence. *Comput. Graph. Image Process.* 5, 382–399 (1993)
8. Matsopoulos, G.K., Delibasis, K.K., Mouravliansky, N.A., Asvestas, P.A., Nikita, K.S., Kouloulias, V.E., Uzunoglu, N.K.: CT-MRI Automatic Surface-based Registration Schemes Combining Global and Local Optimization Techniques, *Technology and Health Care. Official Journal of the European Society for Engineering and Medicine* 11(4), 219–232 (2003)
9. Matsopoulos, G.K., Delibasis, K.K., Mouravliansky, N.A., Nikita, K.S.: A Combination of Global and Elastic Transformations for the Automatic Medical Surface Registration. *Scattering Theory and Biomedical Technology: Modelling and Applications*. In: Massalas, C.V., Dassios, G., Kiriaki, K., Fotiades, D., Payatakes, A. (eds.) *Scattering Theory and Biomedical Technology: Modelling and Applications*, pp. 250–264. World Scientific Publishing Co., Inc., River Edge (2000) ISBN 981-02-4391-X
10. Van den Elsen, P.A., Pol, E.J.D., Viergever, M.A.: Medical image matching - a review with classification. *IEEE Engng. Med. Biol.* 12, 26–39 (1993)
11. Kagadis, G.C., Delibasis, K.K., Matsopoulos, G.K., Mouravliansky, N.A., Asvestas, P.A., Nikiforidis, G.C.: A Comparative Study of Surface- and Volume-based Techniques for the Automatic Registration between CT and SPECT Brain Images. *Medical Physics* 29(2), 201–213 (2002)
12. Maurer, C.R., Fitzpatrick, J.M., Wang, M.Y., Galloway, R.L., Maciunas, R.J., Allen, G.G.: Registration of head volume images using implantable fiducial markers. *IEEE Trans. Med. Imag.* 16, 447–462 (1997)
13. Matsopoulos, G.K.: Automatic Correspondence Methods Towards Point-based Medical Image Registration: An Evaluation Study. In: Exarchos, T. (ed.) *Handbook of Research on Advanced Techniques in Diagnostic Imaging and Biomedical Applications*, ch. 26, pp. 407–425. IGI Global publication (2008) (ISBN 978-3-540-78870-6)
14. Gomes, J., Darsa, L., Costa, B., Velho, L.: *Warping and Morphing of Graphical Objects*. In: *Computer Graphics. The Morgan Kaufman Series* (1998)
15. Lorensen, W.E., Cline, H.E.: Marching Cubes: A high resolution 3D surface construction algorithm. *Computer Graphics* 21 (1987)

# Prediction of Pancreatic Cancer Survival through Automated Selection of Predictive Models

Stuart Floyd<sup>1</sup>, Carolina Ruiz<sup>1</sup>, Sergio A. Alvarez<sup>2</sup>, Jennifer Tseng<sup>3</sup>,  
and Giles Whalen<sup>3</sup>

<sup>1</sup> Department of Computer Science, Worcester Polytechnic Institute,  
100 Institute Road, Worcester, MA 01609 U.S.A.

<sup>2</sup> Department of Computer Science, Boston College

140 Commonwealth Avenue, Chestnut Hill, MA 02467 U.S.A.

<sup>3</sup> Department of Surgical Oncology, University of Massachusetts Medical School  
Worcester, MA 01605 U.S.A.

{floyd,ruiz}@cs.wpi.edu, alvarez@cs.bc.edu,  
{tsengj,whaleng}@ummmc.org

**Abstract.** Cancer survival forecasting may be attempted using models constructed through predictive techniques of various kinds, including statistical multivariate regression and machine learning. However, no single such technique provides the best predictive performance in all cases. We present an automated meta-learning approach that learns to predict the best performing technique for each individual patient. The individually selected technique is then used to forecast survival for the given patient. We evaluate the proposed approach over a database of retrospective records of pancreatic cancer surgical resections.

## 1 Introduction

Adenocarcinoma of the pancreas is one of the most lethal of all cancers. As a result of substantial progress in the treatment of this disorder over the past quarter-century, the five year survival rate has doubled, but is still less than 6%, based on the most recent (2010) National Cancer Institute data for cancers diagnosed between 1999 and 2006 ([1], Table 22.8). Fortunately, there are groups of patients for whom the outlook is significantly better. Cancer stage at diagnosis is of particular importance. For example, the survival rate for localized cancers is fully four times the average. The results of specific diagnostic tests and individual patient attributes including age also affect prognosis.

### 1.1 Machine Learning

Machine learning refers to a set of techniques, including decision tree induction, neural and Bayesian network learning, and support-vector machines, in which a predictive model is constructed or learned from data in a semi-automated fashion (e.g., [2]). In supervised learning, which is the sort considered in the present paper, each data instance used for learning (training) consists of two portions: an unlabeled portion, and a categorical or numerical label known as the class or target attribute that is provided by human experts. The object of learning is to predict each data instance's label based on

the instance's unlabeled portion. The result of learning is a model that can be used to make such predictions for new, unlabeled data instances.

Machine learning has been successfully applied to pancreatic cancer detection [3] and to the analysis of proteomics data in pancreatic cancer [4]. Machine learning techniques have also been shown to provide improved prediction of pancreatic cancer patient survival and quality of life when used either instead of, or together with, the traditional technique of logistic regression [5], [6], [7].

The quality of the predictions produced by a given machine learning method varies across patients. In particular, the method that provides the best predictive model for one patient will not necessarily be optimal for another patient. The latter fact suggests that overall predictive performance across all patients could be improved if it were possible to reliably predict, for each patient, what machine learning method will provide the best performance for that particular patient. The selected method can then be used to make predictions for the patient in question. This is the basic idea behind the approach described in the present paper<sup>1</sup>

## 1.2 Classical Meta-learning

Several meta-learning approaches have been developed in machine learning, including bagging, boosting, and stacking (see below). These approaches are also known as ensemble methods because they aggregate the predictions of a collection of machine learning models to construct the final predictive model. Ensemble machine learning methods have previously been applied to cancer [9], [10], [4]. The present paper describes a new ensemble machine learning approach and its application to prognosis in pancreatic cancer.

In bagging [11], the models in the ensemble are typically derived by applying the same machine learning technique (e.g., decision tree induction, or neural network learning) to several different random samples of the dataset over which learning is to take place. The bagging prediction is made by a plurality vote taken among the learned models in the case of categorical classification, and by averaging the models predictions in the case of a numerical target. In boosting [12], a sequence of models is learned, usually by the same learning technique, with each model focusing on data instances that are poorly handled by previous models. The overall boosting prediction is made by weighted voting among the learned models. Stacking [13] allows the use of different machine learning techniques to construct the models over which aggregation is to take place. In this context, the individual models are known as level 0 models. The outputs of the level 0 models are then viewed as inputs to a second layer of learning, known as the level 1 model, the output of which is used for prediction.

## 1.3 Proposed Automated Model Selection Meta-learning Approach

The model selection approach proposed in the present paper is an ensemble meta-learning approach in that it involves learning a collection of models. Our approach is more similar to boosting and stacking than to bagging in its use of the full training dataset to learn the individual models. However, it differs from classical bagging,

<sup>1</sup> This paper is based on our previous paper [8].



boosting, and stacking, and is characterized by, its adoption of a new prediction target. Rather than aiming to predict the original target, say survival, directly, the goal changes in our approach to identifying what learned model is best qualified to make the desired prediction for a given data instance. Once identified, the selected model alone is used to predict the original target.

## 1.4 Plan of the Paper

Section 2 describes the pancreatic cancer patient database that was constructed for our work. Section 3 presents the details of the model selection meta-learning method proposed in the present paper. Section 4 describes the results of an experimental evaluation of model selection meta-learning over pancreatic cancer clinical data.

# 2 Pancreatic Cancer Clinical Datasets

A clinical database was assembled containing retrospective records of 60 patients treated by resection for pancreatic adenocarcinoma at the University of Massachusetts Memorial Hospital in Worcester. Each patient record is described by 190 fields comprising information about preliminary outlook, personal and family medical history, diagnostic tests, tumor pathology, treatment course, surgical proceedings, and length of survival. The attributes are divided into three major categories: 111 pre-operative attributes, 78 peri-operative attributes, and the target attribute. A summary of the categories of attributes and the number of attributes in each category is presented in Table 1 and Table 2.

## 2.1 Pre-operative and All-attributes Predictive Attribute Sets

We consider two different subsets of predictive (non-target) attributes, each of which gives rise to a dataset when the survival attribute is added as a prediction target:

- The subset containing only the 111 pre-operative attributes (yields the Pre-Operative Dataset).
- The full set of all 189 non-target attributes (yields the All-Attributes Dataset).

## 2.2 Prediction Target Attribute

The prediction target (or target attribute) of our analysis is survival time, measured as the number of months between diagnosis and death. All patients considered in the present study have known dates of death, hence the potential statistical issue of data censoring on the right [14], that is, prematurely “cut off” data, does not occur.

**Survival Discretization.** This work views survival prediction as a classification task, which requires a discrete target attribute. We consider the following three alternative binnings of the target attribute:

- 9 month split, resulting in 2 target values: less than 9 months (containing 30 patients), and 9 months or more (30 patients).

**Table 1.** Pre-operative attributes

Category	Number of attributes	Description
Patient	6	Biographical, physicians
Presentation	21	Symptoms at diagnosis
History	27	Health history
Serum	8	Lab test results
Imaging	23	Diagnostic image details
Endoscopy	25	Endoscopy details
Preliminary outlook	1	Physician’s pre-surgical evaluation
Total	111	

**Table 2.** Peri-operative attributes

Category	Number of attributes	Description
Treatment	36	Treatment details
Resection	24	Surgical removal details
Pathology	7	Post-surgical tumor type results
No resection	11	Reasons for tumor non-removal
Total	78	

- 6 month split, resulting in 2 target values: less than 6 months (20 patients), and 6 months or more (40 patients).
- 6 and 12 month splits, resulting in 3 target values: less than 6 months (20 patients), 6 to 12 months (20 patients), and over 12 months (20 patients).

### 2.3 Summary of Datasets Considered

Each dataset considered in our evaluation is determined by a choice of the set of predictive attributes (see [2.1](#)) together with a choice of discretization of the survival target attribute (see [2.2](#)). Taking into account the two possible subsets of predictive attributes and the three possible discretizations of the target attribute, we therefore consider a total of six distinct datasets in our evaluation.

## 3 Learning to Predict the Best Performing Model

The present paper proposes a new meta-learning approach that identifies for each data instance the predictive model that is best suited to handle that instance. We refer to this approach as model selection meta-learning. The motivation behind our approach is the fact that different models may make correct predictions for different subsets of instances. If we could accurately select a correct model for each instance, overall classification performance would be improved.

### 3.1 Model Selection Meta-learning for Prediction

The proposed model selection meta-learner approach uses two levels of classifiers to predict the unknown target class of a set of a previously unseen input instance, as follows. First, the level 1 model predicts which of the level 0 models is expected to perform

best on the given instance. The instance is then presented to the selected level 0 model to make the final prediction for that instance. The prediction process is described in more detailed pseudocode below. Details of the training process, which must be carried out before prediction can take place, are described in section [3.2](#)

**Model Selection Meta-learning Prediction Algorithm Pseudocode.** It is assumed below that both levels of the meta-learner have previously been trained (see section [3.2](#)). For a previously unseen data instance:

1. Run data instance through level 1 model to select which level 0 model to use. We assume that the learned level 0 models produce class probability distributions  $(p_1 \cdots p_k)$  as their outputs for a given input instance  $x$ , with each  $p_j$  an estimate of the conditional probability  $P(\text{target value } j|x)$  that the given instance has target value  $j$  (e.g., the conditional probability of a patient surviving 9 months or more given the patient's data). These numerical outputs provide the basis for the selection of a best model during meta-learning. The selected model is the one that is expected, based on training data, to output the highest posterior probability  $P_{\text{model}}(\text{correct target value for } x|x)$  for the instance  $x$ .
2. Run data instance through level 0 model selected by level 1 model to predict the target value of the instance (e.g., survival time of the patient).

Our model selector meta-learning approach is similar to stacking in that it uses a collection of level 0 machine learning models followed by a level 1 learner. The key difference is in the function of the level 1 meta-learner. Stacking's level 1 classifier combines the target class probability distributions generated by running the unseen instance through each of the level 0 models, while our model selector's level 1 classifier selects which of the level 0 models is expected to output the highest probability for the correct class of the given test instance. Despite this fundamental difference with stacking, we will use the level 0 and level 1 stacking terminology throughout, for convenience.

### 3.2 Training the Model Selection Metalearner

In section [3.1](#) we described how the model selection meta-learner is used to predict the target class of a new instance, assuming that the meta-learner has previously been trained. We now describe how the training is carried out.

**Level 0 Models.** We will denote the original input dataset by  $I_0$ . Individual machine learning or other predictive techniques may be applied to  $I_0$  to construct predictive models based on details of the learning algorithms for the respective techniques. We assume that each trained model outputs a probability distribution over the possible target values. In our case, the input dataset will be one of the pancreatic cancer datasets described in Section [2](#). Hence, each level 0 model will be trained to predict the survival time of patients, given as a probability distribution over the possible survival time values. For instance, if the 6 and 12 month survival splits are used, then given a patient's data, the trained level 0 model will output the probabilities that the patient will survive less than 6 months, between 6 and 12 months, and more than 12 months.

**Training Procedure.** Given choices  $L_1^0 \cdots L_k^0$  and  $L^1$  of learning algorithms at levels 0 and 1 respectively, a two-stage approach is used to train the model selection meta-learner. First, a new dataset  $I_1$  is constructed from the original dataset  $I_0$  using cross-validation, by relabeling each training instance with the name of the level 0 technique  $L_i^0$  that outputs the highest probability for that instance’s correct target class; this level 0 technique is considered to be the best predictor for that instance. In the second stage, a level 1 model is constructed by applying the learning algorithm  $L^1$  over the new dataset  $I_1$ . Finally, level 0 models are constructed over the full dataset  $I_0$  following the respective individual learning algorithms  $L_1^0 \cdots L_k^0$ . Detailed pseudocode follows the example below.

*Example 1.* An example to illustrate the construction of the dataset to train the level 1 model is illustrated in Table 3.

**Table 3.** Generation of the level 1 dataset for model selection

Instance	Actual class	$P_{\text{ANN}(+)}$	$P_{\text{NB}(+)}$	Model selected
$a^{(1)}$	–	.28	.88	ANN
$a^{(2)}$	+	.41	.53	NB
$a^{(3)}$	+	.99	.88	ANN
$a^{(4)}$	–	.97	.89	NB

Each row of Table 3 corresponds to a patient instance from the level 0 dataset  $I_0$ . The instance is described by the input attribute vector  $a^{(i)}$  in the first column, and belongs to either the + class or – class as shown in the second column. The third and fourth columns show the probability rating that the level 0 models, in this case artificial neural networks (ANN) and Naïve Bayes (NB) respectively, predict for the instance belonging to the + class. The rightmost column lists the model that most highly rates the actual class (or, equivalently, least highly rates the incorrect class). For each row, a new instance is added to the level 1 dataset  $I_1$ . The new instance contains the input attribute vector  $a^{(i)}$  together with the model selected (rightmost column in Table 3) as the target class. The resulting level 1 dataset  $I_1$  is shown in Table 4.

**Table 4.** Resulting level 1 dataset  $I_1$  based on Table 3

Instance	Target class
$a^{(1)}$	ANN
$a^{(2)}$	NB
$a^{(3)}$	ANN
$a^{(4)}$	NB

## Model Selection Meta-learning Training Algorithm Pseudocode

### Inputs

- Set  $I_0$  of input data instances (in our case, each input instance corresponds to the data of a pancreatic cancer patient labeled with the patients survival time as a nominal range)
- Set of level 0 machine learning techniques to use, each of which outputs a class probability distribution
- Level 1 machine learning technique to use
- Integer  $n$  (user-selected number of folds in which to split the input dataset)

*Output:* Trained Level 1 model

(A) Construct a new dataset of training instances  $I_1$  (to be used to train the level 1 model) as follows:

1. Initialize  $I_1$  as empty.
2. Randomly split  $I_0$  into  $n$  stratified folds:  $I_0^1, \dots, I_0^n$ .
3. For each fold  $I_0^i$ , with  $1 \leq i \leq n$ :
  - (a) For each level 0 machine learning technique received as input, train a level 0 model using the machine learning technique and the data instances in the union of all the  $n$  folds except for fold  $I_0^i$ .
  - (b) For each data instance  $d$  in fold  $I_0^i$ :
    - i. Run each level 0 model on  $d$ . Each will output a probability distribution of the target values.
    - ii. Among the level 0 models, select one with the highest probability for  $d$ 's correct target value. This correct target value is given in the input training data  $I_0$ .
    - iii. Add instance  $d$  to  $I_1$ , replacing its original target value with the identifier of the level 0 technique selected in the preceding step (see Example [II](#)).

(B) Train the level 1 model using the dataset  $I_1$ .

(C) Rebuild each level 0 model over all training instances in  $I_0$ .

## 4 Evaluation

We present our experimental evaluation of the model selection meta-learning technique proposed in section [3](#) over the pancreatic cancer resection datasets described in section [2](#).

### 4.1 Preprocessing and Predictive Techniques Used

**Attribute and Feature Selection.** We evaluate models built with different machine learning algorithms using subsets of attributes selected by various feature or attribute selection techniques. Previous work in pancreatic cancer [\[4\]](#), [\[6\]](#), [\[7\]](#) has shown that feature selection can improve the prediction performance of classification methods.

In the current paper we investigate the use of the Gain Ratio, Principal Components Analysis (PCA), ReliefF, and Support Vector Machines (SVM) for feature selection. All of these techniques rank order the most important features, allowing the number of features retained to be prescribed. We experimentally seek the optimal feature selection approach for a given machine learning algorithm.

**Machine Learning Techniques for Level 0 and Level 1 Classifiers.** We consider artificial neural networks (ANN), Bayesian networks (BN), decision trees, naïve Bayes networks (NB), and support vector machines (SVM). The first three classification methods above have previously been identified [6] as the most accurate for prediction tasks over pancreatic cancer datasets closely related to those in the present study among a wide range of methods. SVM is included in the present work both as an attribute selection method and as a classification method.

For each dataset, we find the best combination of feature selection and machine learning algorithm. We use ZeroR (majority class classifier) and logistic regression as benchmarks against which to compare the performance of the models constructed.

## 4.2 Experimental Protocol

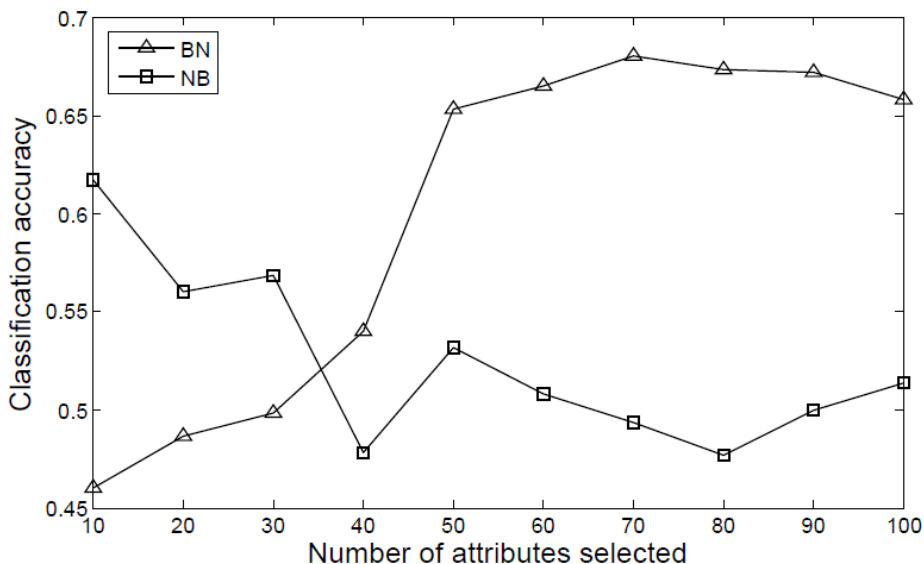
All experiments reported here were carried out using the Weka machine learning toolkit [15]. The classification accuracy for all experiments is calculated as the average value over ten repetitions of 10-fold cross validation, each repetition with a different initial random seed (for a total of 100 runs of each experiment).

For each of the 6 datasets described in Section 2, we apply the following procedure systematically:

1. **Select the Level 0 Classifiers.** We applied each of the machine learning techniques under consideration with and without feature selection to the dataset and recorded the resulting accuracy reported by the 10 repetitions of 10-fold cross validation procedure described above. For each of the machine learning techniques, all of the feature selection approaches were tested with a varying number of attributes to be selected. In most cases, feature selection increased the accuracy of the machine learning methods. Then we selected the top 3 most accurate models among all models: the ones with and the ones without feature selection.
2. **Select the Level 1 Classifier.** Once the top 3 performing level 0 models were identified, we experimentally determined what subset of those 3 top models together with what level 1 machine learning technique would yield the model-selector meta-classifier with the highest predictive accuracy. As above, all machine learning techniques with and without feature selection (and allowing the size of the selected attribute set to vary) were considered for level 1 model construction. Note than in this case, feature selection is applied to the level 1 dataset, not to the original dataset. The model selector meta-classifier with the highest predictive accuracy is reported.

## 4.3 Results and Discussion

We discuss the results of our experimental evaluation, focusing on the pre-operative dataset described by 111 attributes (section 2).

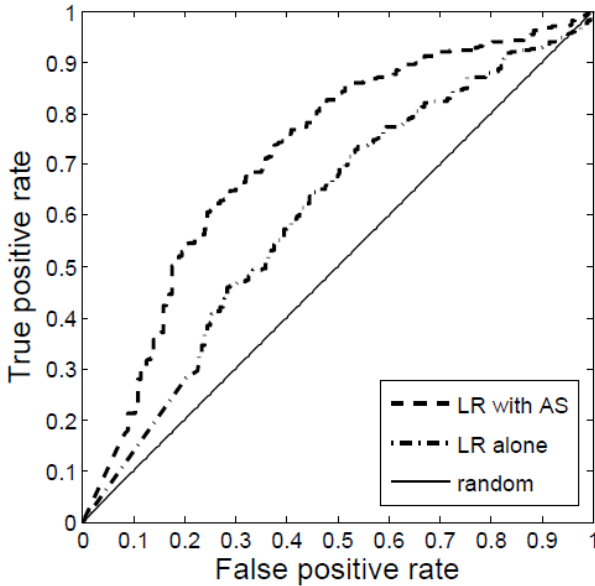


**Fig. 1.** Classification performance of Bayesian techniques for different degrees of attribute selection

**Attribute Selection.** Fig. 1 illustrates two different behaviors that can occur as the number of attributes selected is varied. Only pre-operative predictive attributes are used, with a 6 month survival target split and Support Vector Machine (SVM) attribute selection. The naïve Bayes (NB) classifier clearly benefits from attribute selection here: its classification accuracy is higher when trained over a relatively small number of selected attributes. On the other hand, the Bayes Network (BN) classifier performs best when allowed to operate over a larger set of attributes that has undergone little or no selection. The different results in the two cases may be explained by the fact that the naïve Bayes technique, unlike Bayes Networks, is based on the assumption of conditional independence among the non-class attributes given the class. Attribute selection extracts a smaller set of less-correlated predictive attributes, thus bringing the attribute set closer to satisfying the conditional independence assumption.

The ROC plots in Fig. 2 compare the classification performance of logistic regression, logistic regression with suitably tuned Gain Ratio attribute selection (40 attributes selected), and random prediction of survival time, using only pre-operative predictive attributes and 6 month survival discretization. The plots provide visual confirmation of the fact that the improvement in classification performance due to attribute selection is comparable to the improvement of logistic regression over a purely random prediction of survival time. The corresponding areas under the ROC curves in Fig. 2 are .50, .60, and .72.

**Pre-Operative Dataset, 9 Month Split.** We consider predictive performance when patient survival is discretized into two classes, splitting at nine months. Survival predictions are based on pre-operative attributes only.



**Fig. 2.** Classification performance of logistic regression with and without gain ratio attribute selection

Table 5 shows the classification accuracies of the best performing combinations of a classifier and attribute selection technique for a 9 month survival split, together with the accuracy of the model selection meta-learning approach proposed in the present paper. Among individual predictive techniques, the highest classification accuracies were obtained using Gain Ratio attribute selection in conjunction with either a logistic regression classifier (65.5% accuracy) or an SVM classifier (65.5% accuracy), and ReliefF attribute selection with a Bayesian network classifier (65.3% accuracy). The proposed model selection meta-learning approach slightly outperforms the best individual level 0 classifier methods. In passing, we note that the model selection meta-classifier also outperformed the standard meta-learning techniques of bagging, boosting, and stacking.

**Pre-Operative Dataset, 6 Month Split.** Splitting survival time at 6 months rather than 9 leads to a 2 to 1 class ratio (two-thirds of the patients in this study survived for 6 months or more). This leads to an increase in classification accuracies. Table 6 shows the top three combinations of machine learning classification and feature selection obtained, and the most accurate level 1 classifier constructed over them. Once again our model selection meta-learning method outperformed individual predictive techniques as well as standard meta-learning using bagging, boosting, and stacking.

**Pre-Operative Dataset, 6 and 12 Month Splits.** Classification performance results for the three class dataset obtained by splitting the target attribute at both 6 and 12 months appear in Table 7. The accuracy values are lower in Table 7 than in Table 5 and Table 6 because of the larger number of classes (3 vs. 2). For comparison, randomly guessing the class would lead to an accuracy of approximately 33.3% for this dataset, as the



**Table 5.** Classification accuracy: pre-operative attributes only, nine month split

Machine learning (ML) Technique	Attribute selection (AS)	No. attributes	Accuracy
Best performing ML + AS combinations			
Logistic regression	GainRatio	70	.655
Support vector machine	GainRatio	80	.655
Bayes network	ReliefF	100	.653
Best performing model selection meta-learner			
Level 1: Naïve Bayes	None	111	.673
Level 0: LR, SVM	GainRatio		

**Table 6.** Classification accuracy: pre-operative attributes only, six month split

Machine learning (ML) Technique	Attribute selection (AS)	No. attributes	Accuracy
Best performing ML + AS combinations			
Logistic regression	GainRatio	40	.702
Support vector machine	GainRatio	30	.698
Best performing model selection meta-learner			
Level 1: Logistic regression	PCA	15	.708
Level 0: LR, SVM	GainRatio		

**Table 7.** Classification accuracy: pre-operative attributes only, 6 and 12 month splits

Machine learning (ML) Technique	Attribute selection (AS)	No. attributes	Accuracy
Best performing ML + AS combinations			
Bayes network	ReliefF	20	.527
ANN	GainRatio	50	.518
Support vector machine	ReliefF	80	.485
Best performing model selection meta-learner			
Level 1: Naïve Bayes	None	111	.533
Level 0: ANN	GainRatio		
SVM	ReliefF		

three classes are equally frequent. Model selection meta-learning once again slightly outperforms the individual level 0 models.

**Pre and Peri-Operative Attributes, 6 Month Split.** We discuss here the best model selection meta-classifier model obtained via the approach presented in this paper for the All-Attributes dataset with 6 month survival split, as it illustrates several interesting points. The classification accuracy of this model (75.2%) is significantly greater than that of logistic regression (61.3%), the most widely accepted statistical method in the medical community. As in our other experiments, the model selection meta-classifier is also superior in predictive performance standard meta-learning techniques of bagging (74.5%), boosting (67%), and stacking (72.5%).

**Table 8.** Classification accuracy: pre and peri-operative attributes, 6 month split

$P_{ANN}(> 6 \text{ months})$	$P_{NB}(> 6 \text{ months})$	$P_{Actual}(> 6 \text{ months})$	Correct model(s)
.28	.88	0	ANN
.12	.61	0	ANN
1.00	.15	0	NB
.05	.92	1	NB
.01	.52	1	NB
.41	.53	1	NB
.01	.15	0	Both
.95	.91	1	Both
1.00	.61	1	Both
.99	.55	1	Both
.93	.90	1	Both
.99	.88	1	Both
1.00	.90	1	Both
1.00	.90	0	Neither
.97	.89	0	Neither
.94	.93	0	Neither

*Details of Model Selection.* The meta-classifier constructed by our model selector technique in this particular context combines the models constructed by two top performing level 0 classifiers: Naïve Bayes (using Gain Ratio feature selection) and Artificial Neural Networks (using Gain Ratio feature selection also). A C4.5 decision tree (J4.8 in Weka) coupled with SVM feature selection was used as the level 1 classifier. Next, we will examine the operation of the combined model in more detail.

Table 8 shows the class probability distributions for a small number of selected instances over each of the two level 0 models. The actual target value is also in the table along with a label stating which of the two models (or both) predict this value correctly, or if neither of the models predicts the target value correctly. In 36 out of these 60 instances both models produce the correct classification (31 of which are 6 months or more, and 5 are less than 6 months). In eight instances neither model produces the correct prediction (which is less than 6 months for all eight instances). This leaves 16 instances for which picking the right model would lead to making the correct prediction: in 11 of these naïve Bayes is correct (of which 2 are less than 6 months), and in 5 ANN is correct (of which all are less than 6 months). An interesting observation is that when the artificial neural network and the naïve Bayes model both predict the same target, the artificial neural network is much more certain of its prediction.

*Attributes Selected during Meta-learning.* As mentioned above, SVM feature selection was applied to the level 1 training dataset, reducing the number of attributes from 190 to 70. Remarkably all these 70 selected attributes are pre-operative. We describe below this set of 70 attributes by categories:

*Presentation - Demographic.* (3 attributes selected): Patient’s Height, Weight, and Quality-of-life score (ECOG) at admission.

*Presentation - Initial Symptoms.* (18 attributes selected): Abdominal pain, Back pain, Biliary colic, Clay colored stool, Cholecystitis, Cholangitis, Dysphagia, Fatigue, Indigestion, Jaundice, Nausea, Pruritis, Early Satiety, Vomiting, and Weight Loss.

*Presentation - Presumptive Diagnosis.* (1 attribute selected): Initial diagnosis (e.g., pancreatic tumor, periampullary tumor, etc.).

*Medical History - Comorbidities.* (3 attributes selected): Heart Failure, Ischemic Heart Disease, and Respiratory Diseases.

*Serum Laboratory Tests.* (8 attributes selected): Albumin, Alkaline phosphatase, ALT (alanine transaminase), AST (aspartate aminotransferase), Bilirubin, Amylase, CA19-9 (carbohydrate antigen 19-9), and CEA (carcinoembryonic antigen).

*Diagnostic Imaging - Computed Tomography (CT).* (19 attributes selected): Celiac Artery Involvement, Celiac Nodal Disease, Hepatic Vein Involvement, Inferior Vena Cava Involvement, Lymph node disease or other nodal disease, Node Omission, Portal Vein Involvement, Superior Mesenteric Artery Involvement, Superior Mesenteric Vein Involvement, Tumor Height (cm), Tumor Width (cm), Vascular Omission, and CT Diagnosis.

*Diagnostic Imaging - Endoscopic UltraSound (EUS).* (15 attributes selected): Virtually the same attributes as for CT, and EUS Diagnosis.

*Diagnostic Imaging - Chest X-Rays.* (1 attribute): Chest X-Ray Diagnosis.

*Diagnostic Imaging - Percutaneous Transhepatic Cholangiography (PTC).* (3 attributes selected): Whether stent was used and what type, and PTC diagnosis.

*Learned level 1 model.* The level 1 machine learning technique used here is C4.5 decision trees (J4.8 in Weka). The resulting pruned decision tree appears below. Of the 70 attributes, only 6 are used in the pruned tree: 2 initial symptoms (presentation), including back pain (which was shown to be an important attribute by the Bayesian Nets constructed in other of our experiments) and the occurrence of jaundice; 2 results of diagnostics imaging tests (CT and EUS); and 2 serum lab tests (Bilirubin and Albumin).

```

If patient presents Back Pain
  if CT shows Node Omission
    use Naive Bayes
  else
    if Bilirubin Serum Lab Test <= 0.9
      use Naive Bayes
    else use Artificial Neural Net
else (* patient does not present Back Pain *)
  if patient presents Jaundice
    if EUS shows Vascular Omission
      use Naive Bayes

```

```

else
  if Albumin Serum Lab Test <= 2.4
    use Naive Bayes
  else use Artificial Neural Net
else use Artificial Neural Net

```

## 5 Conclusions

This paper has presented a new approach to combining predictive methods through automated meta-learning, and an evaluation of this technique for the prediction of pancreatic cancer survival using a database of retrospective patient records. The proposed technique, model selection meta-learning, is based on learning which of several available predictive methods can be expected to provide the best results for a given input instance. The motivation for this technique is the fact that different methods sometimes produce conflicting predictions for the same instance. Thus, a system that reliably identifies the best predictor for a given instance will achieve better predictive performance than any of the individual predictors. The experimental evaluation presented in this paper focuses on predicting survival time of pancreatic cancer patients based on attributes such as demographic information, initial symptoms, and diagnostic test results. Individual predictors considered include various machine learning techniques as well as logistic regression. The evaluation results show that the proposed technique of model selection meta-learning produces predictions that are better than those of the individual predictive methods. Also, the proposed technique outperforms the standard meta-learning techniques of bagging, boosting, and stacking in the experiments conducted for this paper. Further work is needed to better establish the magnitude of observed performance differences, and to determine whether any particular machine learning predictors are best suited to being combined through the model selection meta-learning technique introduced in this paper.

## References

1. Altekruse, S.F., Kosary, C.L., Krapcho, M., Neyman, N., Aminou, R., Waldron, W., Ruhl, J., Howlader, N., Tatalovich, Z., Cho, H., Mariotto, A., Eisner, M.P., Lewis, D.R., Cronin, K., Feuer, E.J., Stinchcomb, D.G., Edwards, B.K. (eds.): SEER Cancer Statistics Review, 1975-2007, National Cancer Institute. Bethesda, MD, [http://seer.cancer.gov/csr/1975\\_2007/](http://seer.cancer.gov/csr/1975_2007/); based on November 2009 SEER data submission, posted to the SEER web site 2010
2. Mitchell, T.: *Machine Learning*. McGraw-Hill, New York (1997)
3. Honda, K., Hayashida, Y., Umaki, T., Okusaka, T., Kosuge, T., Kikuchi, S., Endo, M., Tsuchida, A., Aoki, T., Itoi, T., Moriyasu, F., Hirohashi, S., Yamada, T.: Possible detection of pancreatic cancer by plasma protein profiling. *Cancer Res.* 65(22), 10613–10622 (2005)
4. Ge, G., Wong, G.W.: Classification of premalignant pancreatic cancer mass-spectrometry data using decision tree ensembles. *BMC Bioinformatics* 9, 275 (2008)
5. Floyd, S., Alvarez, S.A., Ruiz, C., Hayward, J., Sullivan, M., Tseng, J., Whalen, G.: Improved survival prediction for pancreatic cancer using machine learning and regression, Society for the Surgery of the Alimentary Tract 48th Annual Meeting (SSAT 2007), Washington DC, USA (May 19-23, 2007)

6. Hayward, J., Alvarez, S.A., Ruiz, C., Sullivan, M., Tseng, J., Whalen, G.: Knowledge discovery in clinical performance of cancer patients. In: IEEE International Conference on Bioinformatics and Biomedicine (BIBM 2008), Philadelphia, PA, USA (November 3-5, 2008)
7. Hayward, J., Alvarez, S.A., Ruiz, C., Sullivan, M., Tseng, J., Whalen, G.: Machine Learning of Clinical Performance in a Pancreatic Cancer Database. In: Kim, S. (ed.) special issue on Data Mining Approaches to the Study of Disease Genes and Proteins, *Artificial Intelligence in Medicine* 49(3), 187–195 (2010)
8. Floyd, S., Ruiz, C., Alvarez, S.A., Tseng, J., Whalen, G.: Model Selection Meta-Learning for the Prognosis of Pancreatic Cancer, full paper. In: Proc. Third International Conference on Health Informatics (HEALTHINF 2010), in Conjunction with the Third International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2010), Valencia, Spain, January 20-23, pp. 29–37 (2010)
9. Qu, Y., Adam, B.L., Yasui, Y., Ward, M.D., Cazares, L.H., Schellhammer, P.F., Feng, Z., Semmes, O.J., Wright Jr., G.L.: Boosted decision tree analysis of surface-enhanced laser desorption/ionization mass spectral serum profiles discriminates prostate cancer from non-cancer patients. *Clin. Chem.* 48, 1835–1843 (2002)
10. Bhanot, G., Alexe, G., Venkataraghavan, B., Levine, A.J.: A robust meta-classification strategy for cancer detection from MS data. *Proteomics* 6(2), 592–604 (2006)
11. Breiman, L.: Bagging predictors. *Machine Learning* 24(2), 123–140 (1996)
12. Freund, Y., Schapire, R.E.: A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences* 55(1), 119–139 (1997)
13. Wolpert, D.H.: Stacked generalization. *Neural Networks* 5(2), 241–259 (1992)
14. Lawless, J.F.: *Statistical models and methods for lifetime data*, 2nd edn. Wiley, Chichester (2003)
15. Witten, I.H., Frank, E.: *Data Mining*, 2nd edn. Morgan Kaufmann Publishers, San Francisco (2005)

# Web of Active Documents: An Architecture for Flexible Electronic Patient Records

Federico Cabitza and Iade Gesso

Università degli Studi di Milano-Bicocca  
Viale Sarca 336, 20126 Milano, Italy  
{cabitza, gesso}@disco.unimib.it

**Abstract.** In this paper, we present the WOAD architecture, a design-oriented architecture that allows to build modular and flexible Electronic Patient Record (EPR) based on the metaphor of Active Document. An Active Document is an electronic document that users can easily built by aggregating smaller data modules, called didgets, to mimic their paper-based templates and provide them with proactive behaviors in support of daily practice. After presenting the basic elements of WOAD (i.e., datoms, didgets, templates and mechanisms), we summarize the observational studies that inspired the development of ProDoc, our first WOAD-compliant EPR, and that gave us preliminary user feedbacks for validation. We then illustrate the core implementation details of the WOAD architecture, as it has been deployed in ProDoc.

**Keywords:** WOAD, ProDoc, Datoms, Didgets, Electronic patient record.

## 1 Motivations and Background

In the last five years, our research has focused on the analysis of the Patient Records (either electronic, i.e., EPR or still paper-based) that were used in a number of hospitals of our region and on the design of innovative EPRs that could improve user experience, compliance to hospital and regional healthcare policies, data quality and patient safety. To gain the necessary knowledge on the EPRs used in the considered settings, and to get user feedback on what needs these applications met (or failed to met), we undertook approximately 150 hours of general observations, user shadowing and interviews with practitioners in five departments of three from the main hospitals in Northern Italy. In this ethnographic work, we could recognize most of the unintended shortcomings of ICT reported by other important works (e.g., [12]), especially problems related to workflow inclusion in daily practice and paper persistence. The former issue regarded alterations in work dynamics and ergonomic shortcomings in EPR interfaces that we often saw contributing in eliciting bad emotions and frustration in practitioners; the latter issue regarded the observation that practitioners kept using a sort of parallel paper-based record for utilitarian reasons, as original and informal data entry that is compiled before the electronic counterpart and as pocket-size and foldable proxy of the screenshots of their EPR. In particular, some practitioners we interviewed told us that the precise structure of their paper-based forms was often the outcome of a long-lasting stratification of consolidated work practices, agreements and compromises reached between

clinicians and administrative staff, and local conventions conceived for a more effective and less time-consuming charting. Allegedly, two advantages of traditional forms over electronic ones are lost with the digitization of paper-based patient records: i) high flexibility and easiness in customizing and modifying the template of official paper-based forms; in fact, these forms were usually prints of electronic documents that were created with regular word processors and that, once intended modifications had been accepted by the hospital management, could be modified just in seconds; ii) high flexibility in document use, i.e., in being free to use whatever document of the record at need without being forced to follow any predefined flow of work: in fact, new forms could be created and adopted in daily practice with no strain and without the need to upset usual practice, as could likely happen, e.g., when new forms have to be filled in either to gather new data for clinical research, to comply with new accreditation standards, or new duties about informed consent.

As first results of this long-term project, we conceived the design-oriented framework, WOAD [3], and developed the first WOAD-compliant application, ProDoc; this is a prototypical documental application that we customized for the hospital domain as a highly flexible EPR [4] to address the two requirements above mentioned. In fact, ProDoc allows practitioners to build, customize and use a graphical interface for data entry and retrieval that closely resembles the look and feel of their usual paper-based artifacts so as to mimic the typical interaction with paper forms (see Figure 1). In fact, what in regular EPRs is usually a set of masks to (and views from) the underlying DB, in ProDoc it is a set of persistent documents and forms. Therefore, ProDoc allows users to natively treat and use data in the very terms of the documents they progressively compile. Moreover, ProDoc embeds user-defined active process maps that allow users to get access to any part of the documentation out of any rigid workflow while being aware of the intended flow of activities as it has been defined locally on the basis of practitioners' consensus.

In the next section, we summarize the essential elements of WOAD that underlay the development of ProDoc, i.e., the concept of Active Document, Web of Active Documents and Mechanism. Then, in Section 3 we discuss in more details the WOAD architecture, as it has been deployed in ProDoc, and we illustrate a typical user interaction scenario using ProDoc. The conclusions summarize the main advantages of the WOAD framework in the design of EPRs and outline future lines of research.

## 2 Webs of Active Documents

WOAD is a design-oriented framework grounded on the concept of “active document” (see Figure 2). Each Active Document (AD) can be seen as composed by a “passive” part, i.e., a content container with a specific structure, and an “active” part, i.e., some executable code that provides the passive part with context-aware behaviors. In WOAD, the former part includes the computable definition (i.e., the schema) of modular data structures that are intended to represent a particular aspect of the reality of interest and that, accordingly, we call *datoms* (from ‘documental atoms’); in a single document, a single datom is represented by means of a corresponding didget (from ‘documental widget’) that can be used and reused to build different document templates




Fig. 1. Screenshot from the ProDoc main panel. The prototype supported also PDFs to allow for rich visual annotation of documents.

(where only the didgets’ topological arrangement changes) that share the same groups of data for different purposes and needs. On the other hand, the active part of an AD is composed by a set of (one or more) executable and modular constructs, which we called *mechanisms*; mechanisms are specialized ‘if-then’ statements that are defined over either datoms, didgets or their content and that can be executed by a WOAD-compliant application (like ProDoc) in order to exhibit document-centered behaviors according to current data.

The seminal idea of an “active documental artifact” was first proposed in [5] to refer to data structures capable of assuming an active role in mediating information exchange and coordination among cooperative actors. The most notable research on active documents is the Placeless Document Project developed at PARC [6]. Placeless documents are documents that are managed according to their properties, i.e. sort of metadata that both describe the document’s content and carry the code to implement elementary functionalities of document management (e.g., automatic backup, logging, transmission). In the WOAD framework, we extend this idea by considering any document and form that practitioners are supposed to fill in and consult in their daily practice as parts of an interconnected document system, i.e., what we call a Web of Active Documents, WoAD. WoADs can be highly customized in different domains and application settings to exhibit active behaviors that support users e.g., in keeping track of the patient’s



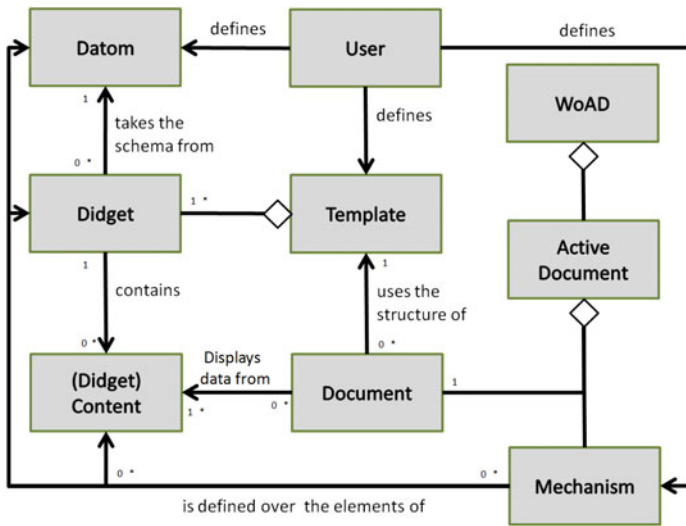


Fig. 2. Relationships between WOAD concepts through an UML entity diagram

illness trajectory, improving the quality of the information exchanged in shift handoffs and patient handovers, enabling activity- and time-related information retrieval and coordinating collaborative tasks (as e.g., order handling). Accordingly in our case study, we conceived of any single form, record or document used to enter and retrieve clinical data (e.g., a therapy prescription sheet) as a single AD that is part of the patient's WoAD, i.e., the patient record as a whole. In what follows, we consider in some more details both the passive part of ADs (i.e., datoms, didgets and templates in Section 2.1) and their active parts (i.e., mechanisms, in Section 2.2).

## 2.1 Datoms, Didgets and Templates

With reference to Figure 3, in WOAD, each document is composed by: i) a set of didgets, each of which is created according to a datom specification and is spatially arranged in a specific document template (see 'Template' in Figure 2); ii) sets of data i.e., the document content, that are associated with the didgets contained in the document's template (see 'Content' in Figure 2).

Datoms are coherent groups of data: at design time, users build them as well-circumscribed sets of form elements, e.g., input fields, iconic elements, buttons, that are gathered together because they relate to either the same abstract data type (e.g., the patient), the same work activity (e.g., drug prescription), or even the same portion of a paper-based artifact, e.g., a table in a chart. Datoms can appear in different positions in one or more documents of a WoAD: when who edits a template puts a datom into a template, she automatically creates a didget for that template in a specific position, i.e., a specific entry point (i.e., a small "form" component) where users can fill in data and where the associated data are displayed every time that users retrieve that specific document from the record.

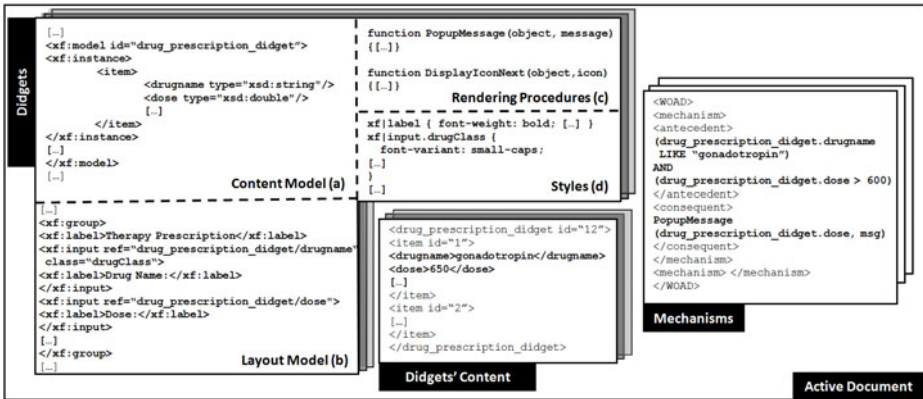


Fig. 3. A graphical representation of the components of an Active Document

The relationship between a datum and its associated didgets is that between a schema specification and the actual data structures that will manage the data at interface level according to that schema. Therefore, a datum specifies what kind of data didgets can store and how to display them: more precisely (see Figure 3), datoms hold a) a content model (i.e., data types, constraints, ranges) that also encompasses a set of quality-oriented functions to check the consistency and completeness of the data inserted in the didgets; b) a layout model (i.e., how data are displayed by didgets according to local “default” styles, see d in Figure 3); c) a set of rendering functions, i.e., executable code that can be interpreted by a client application (e.g., a web browser) to change how the didget’s elements and data are displayed according to the user’s interaction and context.

When who edits a template creates a didget in a certain position of the document, she can specify whether the associated group of elements must appear only once in the document (and exactly there: a so called “single didget”) or, conversely, if users can repeat its data pattern within the didget in tight succession (a so called “multiple didget”). This would allow users to cope with extemporaneous needs for additional room for data that are not predictable at document design time by creating a sort of multi-row table from the simple specification of the first row. For instance, a template of the Anamnesis form can contain a didget to record the examinations previously undertaken by the patient. If this didget has been defined as “multiple”, whenever a physician needs to record more than one examination for a patient, she can add to the Anamnesis form how many rows (i.e., examinations) she needs and all of them will be stored into that didget (as different content items, see Figure 2). Didgets hold all the data that users feed into them: they store and time-stamp these data in a permanent memory, manage metadata that allow to distinguish between (logically) eliminated, provisional and consolidated data and display these latter data in a last-in-first-out fashion. In this way, users can get access to the whole history of data imputation for any single didget. As proof of this concept, in ProDoc we implemented a feature that we informally called “time machine”: in a collapsible section of the main panel, users could find a graphical timeline conveying a visual representation of when relevant events occurred (e.g., an update in a document, the availability of a test report). Users could scroll the timeline along the horizontal axis

by two graduated scales to explore the process history with different time granularity: a scale reporting monthly intervals allowed for quick shifts upon the time axis, while the other one was divided in days and hours to allow for more accurate movements. When a user clicks a point in the timeline, ProDoc gives her access to the documentation and content that was stored in the patient record at that precise time and exactly how it looked. That is, even if the templates of the patient record could change over time (e.g., to comply with a new regional policy that requires to add another checkbox to all forms for the informed consent procedure, or to apply a different classification of patients at the triage), the WOAD platform would deal with it easily thanks to the seamless decoupling at any given time, between clinical content and its topological arrangement in a document, on one hand and, on the other hand, between templates and the didgets they contain. In inserting the didgets, users can also activate (or de-activate) the quality checkers that are associated with the corresponding datoms, as well as to add further elements and areas (i.e., maps, layers) that users can click to get access to groups of templates and their corresponding documents. This functionality allowed us to implement workflow support in ProDoc with four levels of increasing flexibility. On request, ProDoc can display the flow charts of care protocols (the so called clinical pathways [7]) that doctors have wanted to associate to a specific patient. Users can navigate through these active maps (i.e., particular didgets) and, by clicking a specific activity box, they can “filter” the whole patient record to get direct access to all and only the documents associated to that activity. Users could then enact the intended prescriptiveness of the active care protocol by switching between four interaction modes:  $\alpha^1$ ) full action allowed: users can browse and edit the whole patient record as a set of web-based hypertexts and forms;  $\alpha^2$ ) full action allowed with obligation to justify the reason of deviation from the intended process;  $\alpha^3$ ) partly bound action, according to flow relations;  $\alpha^4$ ) full constrained action: charts, records and documents can be accessed only in the intended order in the pathway (e.g., a critical prescription/administration process).

Designers and users can create the datoms and templates they need by means of the Active Document Designer (ADD): this application encompasses the Datom Editor (DE) and the Template Editor (TE), respectively. With the Datom Editor, users can define new datoms in terms of both their content model (e.g., data fields with their type) and the layout model, i.e., their visual aspect. In addition, the DE allows users to define specific styles and rendering procedures (currently defined in JavaScript syntax) that enable advanced features of text formatting and document rendering. For instance, for a certain datom, users can define a rendering procedure that displays a balloon containing some information regarding a specified field (the purpose of these procedures will be clear after reading Section 2.2).

On the other hand, the Template Editor is an editor intended to enable users to create document templates through a graphical interface: they can do so by picking up specific datoms from a palette containing the datoms previously defined with the DE and placing these datoms in the draft template in a what-you-see-is-what-you-get manner. Once datoms have been created and put together as didgets of some template, users can begin use the documents that are built on that template. In fact, a WOAD document is generated by coupling the document template, which is associated with a set of didgets and provides the topological information for their displaying, with the content of its didgets

that is related to the specific document associated to a specific resource, e.g., a patient. The TE also provides a second palette containing the didgets that have already been created in the same WoAD. Users can use this palette in order to reuse the same didgets in different templates. These didgets, which are shared by two or more templates, will display the same content in the documents that, from those templates, are associated to the same resource, in a transparent and strictly synchronized manner.

The Active Document Designer is currently a highly prototypical application that we conceive to allow users create new documents as autonomously as possible according to their ever-changing needs and local conventions [8]. For instance, doctors could want to have only one document for each template be associated with a single patient, as in the case of the Anamnesis form, which reports all the past clinical information through multiple didgets that allow for multiple rows (see above); conversely, doctors could also want to associate many documents based on the same template with the same patient, as in the case of the clinical diary: this is a case where physicians create a new document every day during the patient's stay to organize their clinical annotations in a time-oriented fashion. In both cases, data are recorded as content that is strictly associated with some didget contained in the templates (see Figure 3). In the former case (one document for template with multiple didgets for each patient), doctors can minimize the dispersion of data across different documents and rely on a single "place" to consult; in the latter case (multiple documents for one template for each patient) they can have the didgets' content be partitioned according to some policy (e.g., the stay's length) and each partition be associated with a different document so as to organize the patient record as they were used to in paper-based folders. In either cases, it is noteworthy that the content of all the documents that are based on the same template refers to the same data structures, i.e., the didgets that are contained in the template itself: this guarantees that WOAD-compliant applications can process data as efficiently as in EPRs based on a traditional DBMS, while allowing a higher flexibility in document use as stated above. Currently, datoms and corresponding didgets are represented in XForms<sup>1</sup> syntax and stored in the main component of the WOAD architecture: the Active Document Manager (more details in Section 3). Templates are XML files that store the description of the cartesian displacement of the specific didgets (referenced by IDs) that the document contains.

## 2.2 Mechanisms in the WOAD Framework

As said in Section 2, ADs are coupled to sets of WOAD mechanisms that make them "active" and proactive with respect to their content. Mechanisms are rules that can be defined at level of either i) datom structures, if they are conceived to be valid in general and across the whole WoAD (e.g., syntax checkers, email field validators); ii) didget structures, i.e., when they refer to didgets that are in one or more specific document templates; iii) specific didgets for specific patients, e.g., when a doctor wants to activate either a reminder or a threshold trigger over the body temperature of John Doe only. In all these cases, mechanisms are activated and triggered according to the didgets content. These if-then constructs can be expressed in any rule-based language (for

<sup>1</sup> <http://www.w3.org/TR/xforms/>

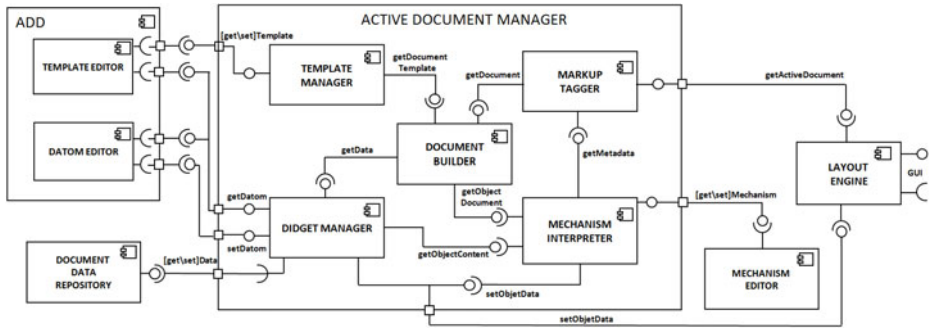
which an interpreter can be integrated to a WOAD architecture) but we also proposed an abstract denotational language to facilitate their definition by users with little or no experience in programming (see the WOAD language presented in [3]). Mechanisms represent conditions over the didgets' content in their *antecedents* (or *if* parts), and trigger simple actions expressed in their *consequents* (or *then* parts) whenever these conditions are met. Mechanisms' antecedents can contain conditions defined over one or more didgets, which are associated with either only one or many templates of the same WoAD, as well as on basic environmental variables, like system time, current users, the patient at hand. Mechanisms are triggered by human interaction with documents: any application behavior for which a programming interface is available can be associated to the mechanism's consequent. Accordingly, we classify mechanisms (to rationalize their design) according to what they do on the document content: we then distinguish between mechanisms that (i) modify the content, e.g., to edit or correct values in data fields; (ii) modify content's attributes and metadata, e.g., timestamps, status flags, urgency attributes; (iii) act on the content, e.g., print (parts of) it, check its quality, validate it, consolidate it (e.g., by digital signature); (iv) transmit the content from one system to another through a communication medium, e.g., by email; (v) route documents and build flows of work, e.g., by allowing users link documents of the same WoAD to each other to easily open a document from another, or by allowing users to open/compile certain (portions of) documents only after that also other (portions of) documents have been compiled (and corresponding tasks performed); (vi) change the content's appearance, e.g., by changing the background color or the font style (in this case, overwriting the local didgets' styles). This last kind of mechanisms act by means of the rendering functions defined at level of didget (see *c* in Figure 3) and have been object of our recent research on how to improve the ways users access and use document content and on how the content can be displayed to make it more "meaningful". In fact, as researchers active in the field of CSCW, we agree with [9] that EPRs should embed specific functionalities to help practitioners be aware of interdependencies between their work and the activities of others to get an understanding of the collaborative context for their own activity. To this aim, in [8] we proposed the concept of Awareness Promoting Information (API), i.e., any annotation, graphical clue, affordance, textual style and indication that could make actors aware of something closely related to the context of reading and writing. The execution of mechanisms can then be seen as a process of API generation, i.e., an operation by which the affordance and appearance of documents and their content is modified, and possibly additional information (e.g., a message) is conveyed to the user in order to make her aware of some condition in the context of document use.

In order to decline the requirements that physicians explicitly expressed in terms of mechanisms, we co-defined with some of their key representatives mechanisms that: could check the correctness of liquid balance values and correct them if necessary (type iii and i); that could mark some values filled in by nurses during a night shift as provisional until the doctor on duty officially double checks them and signs the form (type ii); that could allow practitioners correct a datum without eliminating the previous value for legal concerns (type ii); that could produce official reports printing only parts of the record's content (even distributed across different documents) according to values filled

in in specific fields, e.g., the treatment indication, the kind of procedure (type iii); that could prevent users from opening an operation record if the informed consent and the surgery assessment record have not been duly compiled and signed (type v); that could send a copy of an examination request form to the laboratory as soon as all the necessary fields have been filled in and that could send the discharge letter to the family doctor once the hospital stay has been officially closed (both of type iv); and lastly, that would remind practitioners to compile liquid intakes values at regular intervals and that would check that blood examination requests are compiled within noon and raise due alerts if this is not the case so that results can be returned by the end of the day (both of type vi).

### 3 Interaction with the WOAD Architecture


In this section, we describe the components of the WOAD architecture (see Figure 4), whose conceptual elements have been presented in [3], and illustrate how these interact when users are involved in the basic operations of reading and writing an active document, respectively (see Figure 5). We also provide some details about the current implementation of the WOAD components in ProDoc. For our description, we assume that a template has been created through the ADD and stored in a component of the system that we call *Template Manager* (see below). When a user asks for a certain document from a specific patient record e.g., a drug prescription form, through the GUI of the application (see step 1 in Figure 5), the request is taken by the Layout Engine. This is a standard component that displays active documents, thus allowing the user to interact with them and communicate with the application. Currently, the Layout Engine is any regular Internet browser that supports HTML, CSS and JavaScript full standards. The Layout Engine forwards the request to the *Active Document Manager* (see 2 in Figure 5) which is the main component of any application based on the WOAD architecture. The Active Document Manager builds the (passive) document (see Section 2.1) coupling a template with the content related to the document that has been requested by the user in step 1. In addition, the Active Document Manager provides the data structure on which the application works on the basis of the mechanisms defined by the user (see Section 2.2). To these aims, the Active Document Manager is divided into five sub-components: the *Template Manager*, the *Didget Manager*, the *Document Builder*, the *Mechanism Interpreter* and the *Markup Tagger*. The Template Manager is a component that gives both to the ADD and the other components of the Active Document Manager shared access to templates. The Didget Manager creates and keeps the didgets of a WoAD in its working memory and maintains them synchronized with the *Document Data Repository* (see Figure 4), which is any component that provides data persistence features, e.g., a DBMS. Moreover, the Didget Manager gives both to the ADD and the Document Builder shared access to the datoms' definitions. In our current implementation, the Didget Manager is a Java class that both loads the content associated to the didgets into correspondent sets of objects, and serializes these objects into the Document Data Repository. On the other hand, the Document Builder builds an empty form on the basis of a template and fills in it with the (passive) content associated with the specific document requested by the user. The current Document Builder is a Java class



**Fig. 4.** An UML component diagram of the WOAD architecture underlying the ProDoc implementation

that creates an XForms form by joining the schemas of the didgets (i.e., the datoms) contained in the template and associates the didgets' content with it (see Figure 2). With reference to Figure 5, the Document Builder retrieves the instance of the requested document from its internal memory (step 3), in order to build up it. To perform this retrieval operation, the Document Builder maintains in its internal memory the *document didget list* macrostructure. This is a data structure that contains the whole set of document instances related to a single resource (e.g., a patient); moreover, for each document instance, this data structure maintains the references to its associated didgets and the related didgets' content. On the basis of the data reported in this document instance, the Document Builder retrieves i) the datoms referenced from the Didget Manager (steps 4 and 5), ii) the document template from the Template Manager (steps 6 and 7), and iii) the didgets' content provided by the Document Data Repository through the Didget Manager (steps from 8 to 11). This latter operation involves another macrostructure, the *didget content list*, that is managed by the Didget Manager and contains the whole set of didget's content grouped by didget and related to a single resource. For the sake of simplicity, we have reported the above steps in a sequential fashion (i.e., from 4 to 11), but they can all be executed in parallel. Once the operations of retrieval are done, the Document Builder creates the (passive) document by coupling the document template with datoms, filling the resulting XForms form with the retrieved contents associated to the single didgets, and executing a (convenient) set of transform operations in order to obtain a markup document (e.g., in a XHTML-like<sup>2</sup> syntax) that the Layout Engine can easily render. As soon as the Document Manager gets the document contents, it interacts with the Mechanism Interpreter (step 13) to execute the WOAD mechanisms associated with all the didgets included in the document. To this aim, the Mechanism Interpreter checks the WOAD mechanisms against both the document content and the document structure provided by the Document Manager (steps from 14 to 17), activates the mechanisms whose antecedents are satisfied by either the didgets' content or the document structure, and then selects those to execute according to a resolution strategy based on specificity and currentness [10]. The mechanisms' consequents contain instructions that either modify data or build specific metadata to be associated to

<sup>2</sup> [www.w3.org/Markup/Forms/](http://www.w3.org/Markup/Forms/)



**Fig. 5.** An UML sequence diagram of the typical user interaction with ProDoc

the document (e.g., metadata that prevent data from being modified or metadata that change the appearance of certain values). This association metadata-document is performed by the Markup Tagger; this component receives both the (passive part of) the document that has been created by the Document Builder (step 12), and the metadata that has been produced by the Mechanism Interpreter (step 18), and then it translates the metadata either into appropriate rendering attributes (e.g., stylesheet classes) or into calls to rendering procedures (see c in Figure 3). For instance, if the Mechanisms Interpreter has associated a didget text field with the metadata `<editable>false</editable>`, the Markup Tagger translates it into the HTML attribute “disabled” so that the Layout Engine cannot receive user input for that element. In the current implementation, the Mechanisms Interpreter is based on JBoss Drools<sup>3</sup> and mechanisms are translated into rules that are checked against both the didgets and document objects that are built respectively by the Didget Manager and the Document Builder. The Markup Tagger is a Java class that embeds JavaScript code into the XForms form in order to call the rendering procedures or to modify the style attributes of the XHTML page of the document. The output of the Markup Tagger is then the active document itself: this presents the requested content with the layout specified by the corresponding template and displays additional information according to the metadata. The active document is sent to the Layout Engine (step 19) which finally displays it to the user (step 20).

<sup>3</sup> <http://jboss.org/drools/>



On the way round, when a user modifies an active document (step 21 in Figure 5), the Layout Engine sends the single modifications of the content to the Didget Manager (step 22). As a consequence, the document (i.e., its didgets) will be automatically re-processed by the Mechanisms Interpreter (step 23) that constantly monitors all the WoAD didgets. In this way, every data change can be immediately captured and processed by the application logic that is formalized in the mechanisms that fit the current context best. To this aim, the Mechanisms Interpreter directly interacts with the Didget Manager in order to commit the instructions that update the content (step 24). Finally, the Didget Manager stores the documents' updates into the Document Data Repository for the sake of data persistence (step 25).

## 4 Conclusions and Future Work

Summarizing, in this paper we have illustrated the concept of Active Document within the WOAD framework: an AD is composed by reusable and modular "field nuggets", called didgets, and it is made active by modular interpretable code, called mechanisms, that trigger behaviors according to the context. This modularity and the accentuated separation of information-related and functional needs (addressed by specific didgets and mechanisms in the passive and active part of a document, respectively) is what makes a *web of interconnected ADs* a suitable electronic document platform that can be reused in different work settings and maintained over time to flexibly address ever changing users' needs. To the present moment, the WOAD framework encompasses (i) a conceptual model that represents the main entities and relationships involved in collaborative work mediated by complex document systems; (ii) a denotational language by which to express reactive mechanisms in an abstract and platform-independent way (presented in [3]); (iii) a software component architecture (see Figure 4); (iv) a set of prototypical applications designed to facilitate users in building document templates and active mechanisms.

The first vertical application system to be based on a WOAD architecture is ProDoc: built as a proof-of-concept and prototypical Electronic Patient Record, it has provided first feedback from key users and preliminary validation from the field of work, as reported in [4]. In this paper, we have presented the core implementation choices we undertook for the first deployment of ProDoc, based on JBoss Drools (for the Mechanism Interpreter) and an XForms processor (for the Active Document Manager); we also provided some examples of WOAD mechanisms in the hospital domain to give evidence of the advantage that a modular and rule-based approach can give over more traditional approaches that define functionalities at compile-time and then achieve flexibility at run-time through mere configuration facilities.

Users stressed the requirement of being supported in defining and maintaining their own data structures and associated behaviors without the heavy involvement of ICT practitioners (mainly to reduce costs and times of intervention). For this reason, our research agenda aims to build one single application that could integrate user-friendly functionalities to build active documents also in a visual and intuitive way. Moreover, we plan to further investigate the design-oriented methodology we first proposed in [11] to assist IT practitioners in planning and performing digitization programmes of paper-based patient records in a bottom-up and document-centered fashion. In so doing, we

aim to contribute in the definition of ICT tools and programmes that could significantly minimize the impact in the healthcare domain of the main unintended shortcomings that are currently reported in the specialist literature [2].

**Acknowledgements.** The first author acknowledges the valuable contributions of Giovanni Zorzato and Stefano Corna for the development of ProDoc and the Mechanism Interpreter, respectively.

## References

1. Ash, J.S., Berg, M., Coiera, E.: Some unintended consequences of information technology in health care: The nature of patient care information system-related errors. *Journal of the American Medical Informatics Association* 11, 104–112 (2004)
2. Campbell, E.M., Sittig, D.F., Ash, J.S., et al.: Types of unintended consequences related to computerized provider order entry. *Journal of the American Medical Informatics Association* 13(5), 547–556 (2006)
3. Cabitza, F., Simone, C.: WOAD: A framework to enable the end-user development of coordination oriented functionalities. *Journal of Organizational and End User Computing* 22(2) (2010)
4. Cabitza, F., Simone, C., Zorzato, G.: ProDoc: an electronic patient record to foster process-oriented practices. In: *ECSW 2009: Proceedings of the European Conference on Computer Supported Cooperative Work*, pp. 119–138. Springer, Heidelberg (2009)
5. Divitini, M., Simone, C.: Supporting different dimensions of adaptability in workflow modeling. *Computer Supported Cooperative Work* 9(3), 365–397 (2000)
6. Dourish, P., Edwards, W.K., LaMarca, A., Lamping, J., Petersen, K., Salisbury, M., Terry, D.B., Thornton, J.: Extending document management systems with user-specific active properties. *ACM Transactions on Information Systems* 18(2), 140–170 (2000)
7. Cabitza, F., Sarini, M., Simone, C.: Providing awareness through situated process maps: the hospital care case. In: *GROUP 2007: Proceedings of the 2005 International ACM SIGGROUP Conference on Supporting Group Work*, pp. 41–50. ACM, New York (2007)
8. Cabitza, F., Simone, C., Sarini, M.: Leveraging coordinative conventions to promote collaboration awareness. *Computer Supported Cooperative Work* 18, 301–330 (2009)
9. Pratt, W., Reddy, M.C., McDonald, D.W., Tarczy-Hornoch, P., Gennari, J.H.: Incorporating ideas from computer-supported cooperative work. *Journal of Biomedical Informatics* 37(2), 128–137 (2004)
10. Forgy, C.L.: Rete: A fast algorithm for the many pattern / many object pattern match problem. *Artificial Intelligence* 19(1), 17–37 (1982)
11. Cabitza, F.: Faithful to the earth: Reporting experiences of artifact-centered design in health-care. In: *COOP 2010: Proceedings of the 9th International Conference on the Design of Cooperative Systems*. Springer, Heidelberg (to appear, 2010)

# A Personalization Environment for Multi-version Clinical Guidelines

Fabio Grandi

Dipartimento di Elettronica, Informatica e Sistemistica  
Alma Mater Studiorum – Università di Bologna  
Viale Risorgimento 2, I-40136 Bologna, Italy  
fabio.grandi@unibo.it

**Abstract.** In this work, we introduce a personalization environment for the representation and efficient management of multi-version clinical guidelines. The environment is composed of an XML repository accessed through a personalization engine, which uses temporal perspective, patient profile and context information to reconstruct a guideline version tailored to a specific use case. In particular, we apply and extend to clinical guidelines solutions we previously developed for norm texts in the legal domain, and show how multi-version representation capabilities and personalization query facilities can be added to their management.


**Keywords:** Clinical guidelines, Document retrieval, Temporal database, Ontologies, Personalization, Versioning, XML.

## 1 Introduction

Clinical guidelines are definitions of “best practices” encoding and standardizing clinical procedures for a given disease [18]. The advantages of adopting computer-based guidelines as a support for improving the work of physicians and optimizing hospital activities have been acknowledged by many authors and several computer systems have been developed (see e.g. [8,11]). Clinical guidelines are subject to continuous development and revision by committees of expert physicians and health authorities and, thus, multiple versions coexist as a consequence of the clinical and healthcare activity.

In this paper, we propose to apply to the management of clinical guidelines some techniques we previously developed for norm documents in the legal domain [13,14], which present strong similarities. Hence, we will introduce solutions to model and to provide personalized access to multi-version clinical guidelines, which can be stored both in textual and in executable format in an XML repository [30]. The XML language has already been proposed by many authors and adopted in several research projects (e.g. [3,6,26]) as a suitable means to encode clinical guidelines. Hence, our approach can be considered as a compatible extension of such proposals, to which we aim at adding multi-version representation capabilities and personalization query facilities.

To this end, we will describe how a multi-version XML data model and the prototype system we developed for e-Government applications can be applied to the representation and management of multi-version clinical guidelines. In this way, multiple



**Fig. 1.** Personalized access to multi-version guidelines

temporal perspectives, patient profile and context information can be used by an automated personalization service to build a guideline version tailored to a specific use case (see Fig. 1).

The paper is organized as follows. In Section 2 temporal and semantic versioning of clinical guidelines is introduced with reference to advanced application requirements. Section 3 is devoted to the description of a multidimensional XML data model supporting temporal and semantic versioning of guidelines. In Section 4, a prototype system efficiently implementing the personalization engine sketched in Fig. 1 is briefly described. Conclusions and description of future work will finally be found in Section 5.

## 2 Versioning of Clinical Guidelines

The fast evolution of medical knowledge and the dynamics involved in clinical practice imply the coexistence of multiple *temporal versions* of the clinical guideline documents stored in a repository, since guidelines are continually subject to amendments and modifications. In fact, it is crucial to reconstruct, borrowing the term from the legal field, the *consolidated version* of a guideline as produced by the application of all the modifications it underwent so far, that is the form in which it currently belongs to the state-of-the-art of clinical practice and, thus, must be applied to patients today. However, also past versions are still important, not only for historical reasons: for example, a physician might be called upon to justify his/her actions for a given patient P at a time T on the basis of the clinical guideline versions which were valid at time T and applicable to the pathology of patient P. In other words, temporal concerns are important in the medical domain as they are in the legal domain and, thus, a guideline management system should be able to retrieve or reconstruct on demand any temporal version of a given clinical guideline to meet advanced application requirements.

Moreover, another kind of versioning, which we will call *semantic versioning*, plays a fundamental role, because clinical guidelines or some of their parts have limited

*applicability* with respect, for instance, to the population of patients. In fact, a given guideline (e.g. involving treatment of heart diseases) may contain different recommendations which are not uniformly applicable to the same classes of patients: one general therapy may be non applicable to persons who suffer from some metabolic disorders (e.g. diabetes mellitus) or chronic diseases (e.g. kidney failure) or present some addiction (e.g. cocaine); one first-choice drug may not be given to patients who are already under treatment with possibly interacting drugs (e.g. anticoagulants), or show genetic or acquired hypersensitivity or intolerance to some substances (e.g. patients with enzymatic defects or documented allergies), and so on. Hence, when dealing with a specific patient case, a physician may be interested in finding a *personalized version* of a clinical guideline, that is a version tailored to the patient's health state and anamnesis, only containing recommendations which are safely and effectively applicable to his/her personal case.

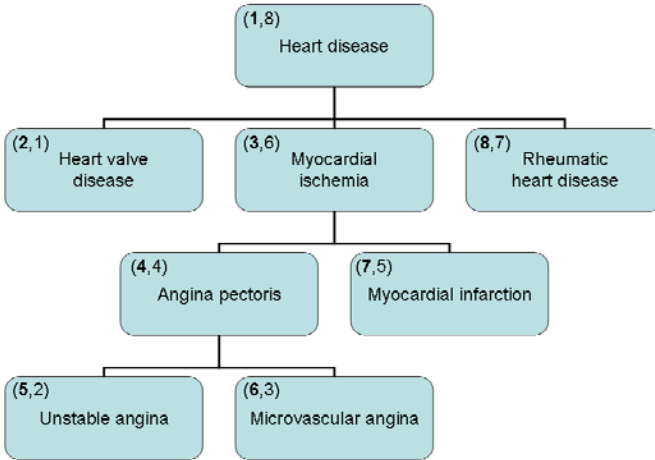
In addition to linking guidelines to classes of patients, semantic versioning can also involve more generic applicability contexts (e.g. hospitals without PET diagnostic equipment, or selected centers taking part to a clinical trial), which might require the application of a particular version of the general guideline, which may also no longer be part of the consolidated state-of-the-art guideline. For instance, consider version v1 of a clinical guideline G which prescribes a biopsy to confirm a cancer diagnosis but has been superseded by a new version v2 which introduces a PET scan for the same cancer diagnosis, making in most cases the biopsy unnecessary. However, in some hospital H which is not equipped with a PET scanner, the right version of G to be followed is v1, although no longer considered valid by the medical community. Therefore, the applicable version of the guideline for context H is  $G(v1)$ , with biopsy as a mandatory diagnostic means. This example also shows how temporal and limited applicability aspects may also interplay in the production and management of versions.

## 2.1 Temporal Versioning

As far as temporal versioning is concerned, several independent time dimensions are involved in the representation and management of clinical guidelines, in particular when we consider an environment also supporting the guideline authoring and approval process. Relevant time dimensions include valid, event, availability, proposal and acceptance times [5,27]. Even considering an environment where only approved guidelines are stored, and retrieved by final users to be consulted or followed, at least two time dimensions are relevant:

**Validity Time.** It is the time the guideline is considered in force by the medical community and, thus, is applied to patients. It has the same semantics of valid time as in temporal databases [20], since it represents the time the guideline actually belongs to the state-of-the-art of clinical practice.

**Efficacy Time.** Borrowing the term from the legal domain, it is the time the guideline can be applied to a concrete case. It usually corresponds to validity, but it might be the case that an obsolete, superseded guideline continues to be applicable to a limited number of cases. While such cases exist, the guideline continues its efficacy though no longer considered in force.



**Fig. 2.** A sample ontology involving a taxonomy of heart diseases, where each class has a name and is associated to a **(pre-order,post-order)** pair

Notice that validity and efficacy time both have the semantics of valid time but represent different and independent valid time notions. Both are necessary to correctly deal with cases as the one in the last described example: the guideline version  $G(v1)$  for the applicability context  $H$  can still be selected today as its efficacy includes current time, although its validity does not. Furthermore, in addition to the time dimensions which model the dynamics of guidelines in the real world, transaction time [20] plays an important role when automatic management of information through computer systems is involved and, thus, should never be neglected, since it allows to execute retro- or pro-active modifications and to keep track of their execution for audit purposes. For example, it might be the case that a physician makes a wrong decision in choosing a drug following the provisions of a guideline retrieved from the system when the returned consolidated version is actually out-of-date; the decision is taken while a modified version of the guideline (e.g. involving the adoption of some more effective and less potentially dangerous drug) is already available but has not been stored in the information system yet. Hence, transaction time is needed to ascertain *a posteriori* that the correct version was stored retroactively and, thus, the physician acted in good faith.

Temporal versioning along multiple time dimensions can be added to documents in an XML repository by making temporal the XML encoding [7], that is introducing timestamps as annotations in the XML document.

## 2.2 Semantic Versioning

Semantic applicability of multi-version resources can be defined with reference to domain ontologies. Ontologies [16,17], which are conceptualizations of a domain into a machine-understandable format, have recently become quite popular with the advent of the semantic web [1], where the introduction of common reference ontologies is necessary to allow information and its interpretation to be shared by both human and automatic agents.

Appropriate applicability of clinical guidelines to individual patients can be defined according to a consensual taxonomy of diseases, like the ICD-10 endorsed by the World Health Organization [19] or the MeSH Section C maintained by the US National Library of Medicine [22]. For instance, consider Fig. 2, which depicts a small portion of a medical ontology representing a classification of principal heart diseases. Notice that, at this stage of the research, we deal with “tree-like” ontologies defined as class taxonomies induced by the IS-A relationship. This will allow us to exploit during query processing the pre-order and post-order properties of trees in order to enumerate the nodes and check ancestor-descendant relationships between the classes; such codes are displayed in the upper left corner of the ontology classes in the Figure, in the form: (**pre-order**,post-order). For instance, the class “Myocardial ischemia” has pre-order “3” which is also its identifier, whereas its post-order is “6”. Before the personalization engine can be used to build a guideline version tailored to a specific patient, the patient must be classified with respect to the disease ontology, on the basis of medical records by means of a suitable reasoning service [13], or through a profile explicitly supplied by the physician. Moreover, additional semantic versioning coordinates, referencing specific domain ontologies, can also be considered to model context-dependent applicability of guidelines.

Hence, in XML resource repositories, reference to ontology concepts (e.g. using class identifiers like those in Fig. 2) can be added to the resource representation and storage as a new versioning coordinate. In this way, applicability annotations can be embedded in the guideline documents to be used by automatic personalization tools. Obviously, also the annotation of clinical guidelines which defines their semantic versioning must be effected by medical domain experts, as part of the guideline drafting and approval process itself. Whenever an ontology definition is changed, temporal versions of the ontology also must be maintained, as the temporal perspectives for navigating the ontology and for searching the guideline repository must be same for consistency reasons. The ontology temporal versioning techniques introduced in [15] can be used to this purpose.

One of the global effects of versioning is an increase in the number or size of the documents to be stored, also depending on the fact that different versions of the same document are stored as separate XML files or are arranged into a single multi-version XML file, owing to a uniform encoding of variant parts within the document structure. The latter solution, which is our choice, is often unavoidable in order to keep the growth of the storage space under control, especially when different versions of the same document may differ by a few nodes only. Personalization, which has shown to be a powerful tool to cope with information overload on the internet [25], can also be particularly effective when used in the management of large XML repositories of versioned documents [13]. In this case, the adoption of personalization techniques can prevent in most cases users to have to go through a huge amount of irrelevant information to find out the right version(s) of the one of interest and, thus, might help to make their search faster and more accurate. Hence, personalization based on semantic versioning may improve the quality of the interaction with the user by further focusing the search on really relevant versions only, which is a desirable feature for clinical guideline management. For example, one of the acknowledged most relevant obstacles

in the use and dissemination of guidelines [4] is the need for adapting them to constraints in local settings (e.g. concerning available hospital resources and practitioners' skills). Management of multi-version guidelines with context-based semantic personalization might help to overcome this problem [9][28]. Other use cases requiring a sort of location-based semantic personalization can also be found: for instance, consider a guideline involving the recommendation of a new drug non yet registered in a given country, or introducing a new protocol only available in selected medical centers participating to an experimental program: the actual contents of the guideline should be changed according to the place where the guideline is retrieved or executed.

From the above discussion and examples, it is evident that the selection and reconstruction of the version(s) of interest for a user become a new challenging problem to support personalization in multi-version guideline retrieval.

### 3 An XML Data Model for Multi-version Guidelines

In this Section, we introduce a multi-version XML document model supporting multiple temporal and semantic versioning coordinates. In doing this, we do not refer to a specific document structure (e.g. defined via a DTD or XML Schema), but we rather introduce a versioning annotation scheme which can be applied to any generic XML resource. In particular, it can be easily adapted to available proposals for the XML encoding of clinical guidelines, including those described in [3][6][26].

We start by formally defining as **version** a piece of text within a guideline document, with a common temporal and semantic pertinence. Owing to the definition, a version can be assigned a timestamp and an applicability annotation to uniquely define its temporal and semantic pertinence. Obviously, different versions of the same object must differ in their temporal and/or semantic pertinence.

For the sake of simplicity, but without loss of generality, we only consider in the examples which follow one time dimension (i.e. validity) and one semantic dimension (i.e. reference to classes in an ontology of diseases like the one in Fig. 2). Let us consider as running example the clinical guideline fragment in Fig. 3 involving recommendations for the treatment of unstable angina patients. The figure displays the text organization, which has a three-level section structure, where section 3.2. has two different versions, namely 3.2(v1) and 3.2(v2), whereas section 3.2(v1).2 has three different versions, namely 3.2(v1).2(v1), 3.2(v1).2(v2) and 3.2(v1).2(v3). The multi-version XML encoding of such guideline fragment is shown in Fig. 4.

In the XML encoding, we use the <version> element to delimit the boundaries of a version within the document. The <valid> and <applies> elements are then used to assign the temporal and semantic pertinence, respectively, to the version which contains them. Validity and applicability properties are inherited by descendant nodes in the XML tree-structure unless locally redefined with a new version definition. Therefore, there is no reason to repeat the valid or applies annotation when the pertinence is not changed from the ancestor version in the XML tree-structure. In general, redefinition may involve only a subset of the versioning dimensions, while the others dimensions are inherited.



**RECOMMENDATIONS**

1. IDENTIFICATION OF PATIENTS  
WITH RISK OF UNSTABLE ANGINA  
...
2. INITIAL EVALUATION AND MANAGEMENT  
...
3. EARLY HOSPITAL CARE
  - 3.1. **Initial Treatment Strategy**  
...
  - 3.2. **Drug Therapy**
    - 3.2(v1). *Anti-Ischemic and Analgesic Therapy*
      - 3.2(v1).1. Therapy with nitrates  
...
      - 3.2(v1).2. Therapy with beta-blockers
        - 3.2(v1).2(v1). ...administration of drug D1...
        - 3.2(v1).2(v2). ...administration of drug D2...
        - 3.2(v1).2(v3). ...administration of drug D3...
    - 3.2(v1).3. Therapy with ACE inhibitors  
...
    - 3.2(v2). *Antiplatelet/Anticoagulant Therapy*  
...
4. CORONARY REVASCULARIZATION  
...
5. LATE HOSPITAL CARE  
...

**Fig. 3.** The structure of a fragment of a sample multi-version clinical guideline

With reference also to Fig. 3 the XML fragment in Fig. 4 shows, within the outermost <recommendations> element, a hierarchical structure based on three levels of sections. The <recommendations> element is composed of one version, which defines its global semantic and temporal pertinence, that is applicable to class C3 in the ontology in Fig. 2 (patients with myocardial ischemia) and valid from 1980 on. It is made of several first-level sections (see also Fig. 3), of which only section 3 is evidenced in the Figure. Such a section, made of only one version to specify applicability to ontology class C4 (patients with angina pectoris), deals with Early Hospital Care. Its temporal pertinence is inherited from the container element.

In general, by means of redefinitions we can introduce, for each part of a document, complex validity and applicability properties including extensions or restrictions with respect to ancestors. For instance, the applicability assignment to section 3 which we just described is a restriction and the attribute *to* is used to this end. Actually, the applicability assigned to the version is the intersection of the *to* value and of the value inherited by the ancestor version (in this case  $C4 \cap C3$ , which equals C4 since it is a subclass of C3). The same applies to second-level section 3.2 (entitled “Drug

---

```

...
<recommendations>
<version number="1">
  <applies to="C3"/>
  <valid from="1980-01-01" to="9999-99-99"/>
  ...
  <section number="3">
  <version number="1">
    <applies to="C4"/>
    <title>Early Hospital Care</title>
    ...
    <section number="2">
    <title>Drug Therapy</title>
    <version number="1">
      <applies to="C5"/>
      <title>Anti-ischemic and Analgesic Therapy</title>
      ...
      <section number="2">
      <title>Therapy with beta-blockers</title>
      <version number="1">
        <valid from="1980-01-01" to="1990-12-31"/>
        ...administration of drug D1...
      </version>
      <version number="2">
        <valid from="1991-01-01" to="1998-12-31"/>
        <valid from="2001-01-01" to="2003-12-31"/>
        ...administration of drug D2...
      </version>
      <version number="3">
        <valid from="1985-01-01" to="9999-99-99"/>
        ...administration of drug D3...
      </version>
      ...
    </section>
    ...
  </version>
  <version number="2">
    <applies also="C7"/>
    <title>Antiplatelet/Anticoagulant Therapy</title>
    ...
  </version>
  ...
</section>
...
</version>
</recommendations>
...

```

---

**Fig. 4.** An XML fragment showing the multi-version encoding of the guideline in Fig. 3

Therapy”), whose first version (entitled “Anti-ischemic and Analgesic Therapy”) applies to class C5 (unstable angina), which is also a restriction, whereas the second version (entitled “Antiplatelet/Anticoagulant Therapy”) is also applicable to class C7 (myocardial infarction), which is an extension indeed. Attribute *also* is used in this case, and the applicability assigned to the version is the union of the *also* value and of

the value inherited by the ancestor version (class C4UC7). In other words, the contents of section 3.2(v2) both apply to angina pectoris and myocardial infarction patients.

The third-level section 3.2(v1).2 entitled “Therapy with Beta-blockers” is made of several versions, each one dealing with the administration of a specific drug and having its own temporal pertinence, whereas the (inherited) applicability is the same (namely C5, unstable angina). In order to derive the validities of the three drugs shown in the Figure, we assume the recommendations underwent the evolution which follows. Drug *D1* was introduced in 1980 and then replaced by the drug *D2* in 1991. However, the use of drug *D2* was suspended from 1999 to 2000, period during which it had been under investigation since suspected of causing adverse reactions. In 2004, due to evidence of long-term adverse effects, *D2* was definitely withdrawn. Drug *D3* has been introduced in 1985. Hence, the resulting history of recommended beta-blockers according to the guideline in Fig. 3 (which will in fact correspond to the answers to a sequence of snapshot queries issued on the multi-version document) is the following:

- from 1980 to 1984: drug *D1*
- from 1985 to 1990: drugs *D1* and *D3*
- from 1991 to 1998: drugs *D2* and *D3*
- from 1999 to 2000: drug *D3*
- from 2001 to 2003: drugs *D2* and *D3*
- from 2004 on: drug *D3*

As for 3.2(v1).2(v2) in the Figure, versions can be assigned multiple intervals as validity: this corresponds to adopt *temporal elements* [10,20], that is disjoint union of intervals, as timestamps.

### 3.1 Manipulation and Retrieval Operations

The multi-version XML data model can be equipped with two basic operators for the management of guideline authoring and maintenance: one devoted to change the textual content of a guideline portion and the other to allow modifications to the temporal and semantic pertinence of a given version. The former can be used for deletion of (a part of) the guideline (abrogation), or the introduction of a new part of the guideline (integration), or the replacement of (a part of) the guideline (substitution). The latter can be used to deal with the time/applicability extension or restriction of (part of) the guideline. Such operators, in order to preserve the well-formedness of the version structure and the inheritance semantics, can be defined in a similar way as the ones defined for multi-temporal norm documents in [14].

Clinical guideline repositories, like the US National Guideline Clearinghouse [23] or the UK National Library of Guidelines [24], are usually managed by traditional information retrieval systems where users are allowed to access their contents by means of keyword-based queries expressing the subjects they are interested in. Adopting a system like the one described in [13] that we developed for norm documents, users are offered the possibility of expressing temporal and semantic specifications for the reconstruction of a consistent version of the retrieved guideline.

In particular, the queries can contain four types of constraints: temporal, structural, textual and applicability. Such constraints are completely orthogonal and allow the users to perform very accurate searches in the XML guideline repository. Let us focus first on

---

```

FOR      $a IN guidelines.xml
WHERE    textConstr ($a//section/title/text(), 'anticoagulant')
AND      tempConstr ('vTime OVERLAPS PERIOD('2007-01-01','2008-12-31')')
AND      applConstr ('C5')
RETURN  $a

```

---

**Fig. 5.** An XQuery-equivalent query executable on a clinical guideline personalization system

the applicability constraint. Consider again the ontology in Fig. 2 and guideline fragment in Fig. 4 for the treatment of John Smith, an “infarctuated” patient (i.e. belonging to class C7), the sample recommendations in Fig. 3 will be selected as pertinent, but only the second version of Section 3.2 will be actually presented as applicable. Furthermore, the applicability constraint can be combined with the other three ones in order to fully support a multi-dimensional retrieval. For instance, a physician (or an health insurance officer) could be interested in all the guidelines ...

- ... which have a section whose title (*structural constraint*) contains the word anti-coagulant (*textual constraint*), ...
- ... which were valid between 2007 and 2008 (*temporal constraint*), ...
- ... and which are applicable to a patient suffering from unstable angina (*applicability constraint*).

More precisely, the system is able to answer queries having the XQuery [32] FLWR syntax in Fig. 5, where `textConstr`, `tempConstr`, and `applConstr` are suitable functions allowing the specification of the textual, temporal and applicability constraints, respectively (the structural constraint is implicit in the XPath expressions used in the XQuery statement).

## 4 Prototype Implementation

The personalization engine in Fig. 1, which is capable to execute queries like the one in Fig. 5, has been implemented as a prototype Multi-version XML Query Processor. The prototype code is written in Java JDK 1.5 and employs ad-hoc data structures (relying on embedded “light” DBMS libraries) and algorithms which allow users to reconstruct on-the-fly the desired personalized version of the XML guideline, satisfying temporal, structural, textual and applicability constraints, by means of a multi-version extension of the *holistic twig join* approach [2]. Guidelines are stored in the XML repository using an indexing scheme based on *multi-version inverted indices*, that is an extension with timestamps and semantic annotations of the indexing solution proposed in [33]. In practice, the query processing algorithm implements the temporal slice operator proposed in [21], to which the processing of semantic constraints has been added, without an appreciable overhead. In fact, thanks to the properties of the adopted pre- and post-order encoding of the ontology classes, applicability constraints can be very efficiently tested during query processing by means of simple comparisons. A detailed presentation of the deployed data structures and holistic join techniques, together with a related work discussion on these topics, can be found in [12].

As a result, we obtain a high overall query processing efficiency mated with low memory requirements. In order to evaluate the performance of the prototype, a specific query benchmark was built and several exploratory experiments were conducted to test the personalization engine behavior under different workloads. The experiments have been effected on a Pentium 4 3Ghz Windows XP Professional workstation, equipped with 1GB RAM and a 160GB EIDE disk with NT file system (NTFS). Test were performed on three XML document collections of increasing size (namely 5,000, 10,000 and 20,000 guidelines, with a total size of 120MB, 240MB and 480MB, respectively). In all collections the guidelines were synthetically generated by means of a suitable tool, which is able to produce XML documents compliant to our multi-version model under different parameter configurations. For each collection, the average, minimum and maximum document size was 24KB, 2KB and 125KB, respectively. Experiments were conducted by submitting queries of five different types, mixing in various ways structural, textual, temporal and applicability constraints.

The system behavior showed a good efficiency in every context, providing a response time (including query analysis, retrieval of the qualifying guideline parts and reconstruction of the result) of a few seconds for most of the queries. Moreover, the selectivity of the query predicates does not impair performances, even when large amounts of documents containing some (typically small) relevant portions have to be retrieved. The system is able to deliver a fast and reliable performance in all cases, since it practically avoids the retrieval of useless document parts. For the same reasons, the main memory requirements of the Multi-version XML Query Processor are quite limited, less than 5% with respect to an approach like the one adopted in [14], where complete documents are retrieved with a traditional XML engine working on structural and textual constraints, and then temporal and applicability constraints are applied using a DOM representation to prune out non-qualifying XML nodes. Notice that this property is very interesting for a system which is likely to run in a highly concurrent multi-user environment, since memory requirements are not crucial for performance. The prototype system also showed a good scalability behavior in every type of query setting, as the computing time for the same query always grows linearly with the number of documents. Full details on performance evaluation can be found in [12][13].

## 5 Conclusions and Future Work

In this paper, we applied to the representation and management of clinical guidelines some techniques we previously developed for norm documents in the legal domain [12][13]. In particular, we introduced solutions to model and to provide personalized access to multi-version guidelines, supporting multiple temporal and semantic versioning coordinates. The proposal involves the definition of a multi-version XML data model and the implementation of a prototype personalization engine.

Preliminary experimental work on query performance, with repositories of syntectic XML documents, showed encouraging results. In particular, the personalization engine proved to be very efficient in a large set of experimental situations and showed excellent scale-up figures with varying load configurations.

We underline that the very same techniques we presented for personalized access to multi-version textual guideline documents can also be applied to the enactment of

workflows implementing multi-version clinical guidelines, provided that workflows are specified using an XML-based definition language, like BPEL [29] or XPDL [31], which can be enriched as well with temporal and semantic annotations in order to define versions [13].

Future work will consider the improvement of the approach to cope with more advanced application requirements (e.g. relaxing of constraint of tree-like ontologies) and the completion of the technological infrastructure required to set up the personalization platform with the design and implementation of auxiliary services (e.g. for automatic patient classification with respect to the disease ontology). Further work will also include the assessment of our developed system in a concrete working environment, with real users and in the presence of a repository of real clinical guidelines.

## References

1. Berners-Lee, T., Hendler, J., Lassila, O.: The Semantic Web. *Scientific American* 284(5), 34–43 (2001)
2. Bruno, N., Koudas, N., Srivastava, D.: Holistic Twig Joins: Optimal XML Pattern Matching. In: *Proc. of SIGMOD 2001*, pp. 310–321 (2001)
3. Buchtela, D., Peleška, J., Veselý, A., Zvárová, J., Zvolský, M.: An XML-based Format for Guideline Interchange and Execution. In: *Proc. of MIE 2008*, pp. 151–156 (2008)
4. Cabana, M.D., Rand, C.S.N.R., Powe, C.S., Wu, A.W., Wilson, M.H., Abboud, P.C., Rubin, H.R.: Why don't Physicians follow Clinical Practice Guidelines? A Framework for Improvement. *Journal of American Medical Association* 282(15), 1458–1465 (1999)
5. Combi, C., Montanari, A.: Data Models with Multiple Temporal Dimensions: Completing the Picture. In: Dittrich, K.R., Geppert, A., Norrie, M.C. (eds.) *CAiSE 2001*. LNCS, vol. 2068, pp. 187–202. Springer, Heidelberg (2001)
6. Dubey, A.K., Chueh, H.C.: An XML-based Format for Guideline Interchange and Execution. In: *Proc. of AMIA 2000*, pp. 205–209 (2000)
7. Dyreson, C.E., Grandi, F.: Temporal XML. In: Özsu, M.T., Liu, L. (eds.) *Encyclopedia of Database Systems*, pp. 3032–3035. Springer, Heidelberg (2010)
8. Fridsma, D.B.: Special Issue on Workflow Management and Clinical Guidelines. *AIM Journal* 22(1), 1–80 (2001)
9. Fridsma, D.B., Gennari, J.H., Musen, M.A.: Making Generic Guidelines Site-specific. In: *Proc. of AMIA 1996*, pp. 597–601 (1996)
10. Gadia, S.K.: A Homogeneous Relational Model and Query Languages for Temporal Databases. *ACM Trans. on Database Systems* 13(3), 418–448 (1988)
11. Gordon, C., Christensen, J.P.: *Health Telematics for Clinical Guidelines and Protocols*. IOS Press, Amsterdam (1995)
12. Grandi, F., Mandreoli, F., Martoglia, R.: Issues in Personalized Access to Multi-version XML Documents. In: Pardede, E. (ed.) *Open and Novel Issues in XML Database Applications*, pp. 199–230. Information Science Reference (2009)
13. Grandi, F., Mandreoli, F., Martoglia, R., Ronchetti, E., Scalas, M.R., Tiberio, P.: Ontology-based Personalization of e-Government Services. In: Mourlas, C., Germanakos, P. (eds.) *Intelligent User Interfaces*, pp. 167–203. Information Science Reference, United Kingdom (2009)
14. Grandi, F., Mandreoli, F., Tiberio, P.: Temporal Modelling and Management of Normative Documents in XML Format. *Data & Knowledge Engineering* 54(3), 327–354 (2005)
15. Grandi, F., Scalas, M.R.: The Valid Ontology: A Simple OWL Temporal Versioning Framework. In: *Proc. of SEMAPRO 2009*, pp. 98–102 (2009)

16. Gruber, T.: Ontology. In: Özsu, M.T., Liu, L. (eds.) *Encyclopedia of Database Systems*, pp. 1963–1965. Springer, Heidelberg (2010)
17. Guarino, N. (ed.): *Formal Ontology in Information Systems*. IOS Press, Amsterdam (1998)
18. Hutchinson, E.A., Baker, R. (eds.): *Making Use of Guidelines in Clinical Practice*. Radcliffe Medical Press (1999)
19. *International Statistical Classification of Diseases and Related Health Problems*. World Health Organization (2010), <http://www.who.int/classifications/icd/en/>
20. Jensen, C.S., Dyreson, C.E., Böhlen, M.H., Clifford, J., Elmasri, R.A., Gadia, S.K., Grandi, F., Hayes, P.J., Jajodia, S., Käfer, W., Kline, N., Lorentzos, N.A., Mitsopoulos, Y., Montanari, A., Nonen, D., Peressi, E., Pernici, B., Roddick, J., Sarda, N.L., Scalas, M.R., Kang, H.-S., Snodgrass, R.T., Soo, M.D., Tansel, A.U., Tiberio, P., Wiederhold, G.: *The Consensus Glossary of Temporal Database Concepts - February 1998 Version*. In: Etzion, O., Jajodia, S., Sripada, S. (eds.) *Dagstuhl Seminar 1997*. LNCS, vol. 1399, pp. 367–405. Springer, Heidelberg (1998)
21. Mandreoli, F., Martoglia, R., Ronchetti, E.: *Supporting Temporal Slicing in XML Databases*. In: Ioannidis, Y., Scholl, M.H., Schmidt, J.W., Matthes, F., Hatzopoulos, M., Böhm, K., Kemper, A., Grust, T., Böhm, C. (eds.) *EDBT 2006*. LNCS, vol. 3896, pp. 295–312. Springer, Heidelberg (2006)
22. *Medical Subject Headings - Section c: Diseases*. US National Library of Medicine (2010), [http://www.nlm.nih.gov/mesh/2009/mesh\\_browser/MeSHtree.C.html](http://www.nlm.nih.gov/mesh/2009/mesh_browser/MeSHtree.C.html)
23. *National Guideline Clearinghouse*. US Agency for Healthcare Research and Quality (2010), <http://www.guideline.gov/>
24. *National Library of Guidelines*. UK National Institute for Health and Clinical Excellence (2010), <http://www.library.nhs.uk/GUIDELINESFINDER/>
25. Riecken, D.: *Personalized Views of Personalization*. *Communications of the ACM* 43(8), 27–28 (2000)
26. Shiffman, R.N., Karras, B.T., Agrawal, A., Chen, R., Marengo, L., Nath, S.: *GEM a Proposal for a More Comprehensive Guideline Document Model using XML*. *Journal of AMIA* 7(5), 488–497 (2000)
27. Terenziani, P., Montani, S., Bottrighi, A., Molino, G., Torchio, M.: *Clinical Guidelines Adaptation: Managing Authoring and Versioning Issues*. In: Miksch, S., Hunter, J., Keravnou, E.T. (eds.) *AIME 2005*. LNCS (LNAI), vol. 3581, pp. 151–155. Springer, Heidelberg (2005)
28. Terenziani, P., Montani, S., Bottrighi, A., Torchio, M., Molino, G., Correndo, G.: *A Context-adaptable Approach to Clinical Guidelines*. In: *Proc. of MEDINFO 2004*, pp. 169–173 (2004)
29. *The Web Services Business Process Execution Language*. WfMC Coalition (2010), <http://www.wfmc.org/standards/docs.htm>
30. *The eXtensible Markup Language home page*. W3C Consortium (2010), <http://www.w3c.org/XML/>
31. *The XML Process Definition Language*. OASIS Organization (2010), [http://www.oasis-open.org/committees/tc\\_home.php?wg\\_abbrev=wsbpel](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsbpel)
32. *The XML Query home page*. W3C Consortium (2010), <http://www.w3c.org/XML/Query>
33. Zhang, C., Naughton, J., DeWitt, D., Luo, Q., Lohman, G.: *On Supporting Containment Queries in Relational Database Management Systems*. In: *Proc. of SIGMOD 2001*, pp. 425–426 (2001)

# Dynamic Routing Using Health Information Policy with Apache Camel

Edward Brown<sup>1</sup> and Jamie Goodyear<sup>2</sup>

<sup>1</sup>Department of Computer Science, Memorial University of Newfoundland  
St. John's, Canada

<sup>2</sup>Progress Software Corporation, 570 Newfoundland Drive, St. John's, Canada  
brown@mun.ca, jamie.goodyear@progress.com

**Abstract.** Apache Camel is an implementation of a messaging framework using the enterprise integration patterns advanced by Hohpe and Woolf. This framework allows the messaging architecture to be configured at run-time, by changing routing rules which determine the disposition of messages transmitted by the system. Our work illustrates an original use of the dynamic routing capability to implement health information policy, by putting the routing rules in the hands of policy administrators rather than network managers. This paper is an extended version of reference [1].

**Keywords:** Policy, Messaging, Architecture, Middleware, Medical context.

## 1 Introduction

This paper deals with information exchange in the context of health care. In particular, policy updates that affect the handling of health information can impose costly revision to network infrastructure and applications. We propose an approach based on the Apache Camel messaging framework, in which routing rules can be updated dynamically. By using dynamic routing rules to implement information policy, rather than simply as a network management or administration capability, the messaging infrastructure becomes responsive to policy changes, without requiring modification to end-user applications.

Section 2 of this paper reviews the background concepts of Message-Oriented Middleware (MOM), Enterprise Integration Patterns (EIPs), Apache Camel as an EIP implementation, and health information policy. Section 3 discusses the use of medical information and policies for dynamic routing of information through a health care information system, and section 4 provides one dynamic routing scenario tested in a prototype system.

## 2 Background

### 2.1 Message-Oriented Middleware (MOM)

Large scale networked applications are not deployed and used in isolation, they are used with other networked applications and services. Integration raises a number of challenges as emphasized by Hohpe and Woolf [2]:



1. Network reliability: Distributed computing involves delays and interruptions. In comparison to using a single computer the distributed application has a larger set of problems to manage.
2. Networks are slow: Data transfer across networks is significantly slower than local method calls. Applications must be designed with slow responses in mind.
3. Applications are different: Integration solutions need to be able to interface and handle data from multiple platforms and data formats.
4. Change is inevitable: Systems change over time, so how they connect with other systems should be able to change to maintain connections. A solution needs to promote loose coupling, otherwise a change at one network node may cause a ripple effect of required changes to other nodes.

Historically, the main approaches to networked applications are:

1. File Transfer: One application writes a file that another application reads. This requires agreement on data format, naming conventions, and on the timing of operations.
2. Shared Database: Many applications share a common database. This reduces the data transfer problem between applications as all data resides on one physical store.
3. Remote Procedure Invocation: An application exposes an interface to another remote application. Communication is done in real time, in a synchronous manner.
4. Messaging: An application publishes a message to a common message queue, and other applications may read the message at a later time. Communication is asynchronous in nature, using an agreed message format.

We adopt messaging as the best approach for an infrastructure that to support frequent change with loosely coupled client applications. A message oriented middleware (MOM) framework provides messaging capabilities for applications and manages message transfer, much as a database is provides data services to applications.

The MOM supports reliable asynchronous application to application communication, (that is, messages are maintained by the queue). Applications communicate by passing data to and from a common queue, such data exchanges constituting "messages". A sender or producer application writes a message to a queue, while a receiver or consumer application receives messages by reading from a queue. The message is effectively a data structure or object, and may embed a number of difference concepts from simple data to events or command requests. Conventionally, a message has two functional components: the header and body. A message header contains information related to its origin, its destination, and other information used by the messaging system, and usually ignored by the application. The message body is a container for the application data payload, generally ignored by the messaging framework.

## 2.2 Enterprise Integration Patterns

A common collection of solution patterns related to the messaging approach are characterized in [2] as Enterprise Integration Patterns (EIPs). These patterns do not directly solve all integration problems, but catalogue solutions to recurring network integration problems. Patterns related to events, channels, services, shared resources, documents, and message management are included in the collection. The two most important patterns for this work are Content Enricher and Dynamic Router.

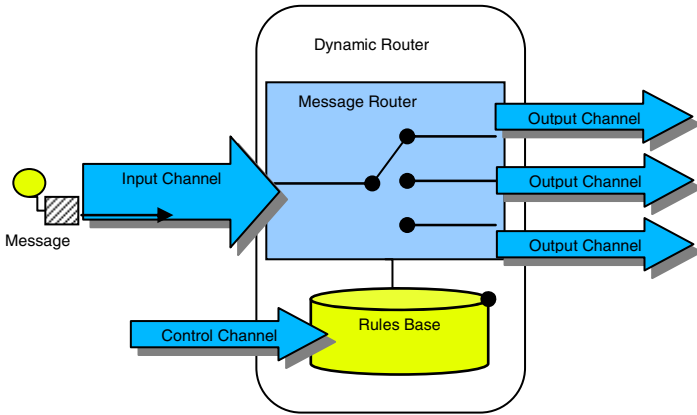


Fig. 1. Dynamic Router EIP after [2]




Fig. 2. Dynamic Router and Content Enrichment EIPs in a client/server configuration

The Dynamic Router EIP is illustrated in figure 1. It routes messages according to a rules base which can be configured manually or automatically. Rules may be changed dynamically without halting any system operations or services. The router makes decisions about where specific messages are to be sent based on the message content and the current rules in its rules base. Typically, information included in the message header are matched with rules in the rule base to determine the appropriate message channel. The message channel itself is effectively a sophisticated queue. Since the router is dynamic, the rules can be modified manually or by appropriately authenticated software. In essence, this paper characterizes the rules in the rules base as the encoding of information management policies.

The dynamic router allows the explicit expression of routing policies in a consistent manner, and at a known and configurable point in the overall system. All applications operate in accordance with the routing rules, since all messaging will be performed through a dynamic router. The rules base frees message distribution from static preplanning, as routes may be added, modified, or removed at run time.

The dynamic router includes a rules engine (not shown). The engine is responsible for accessing, modifying and storing rules for the router. For example, the rules

engine may be directed to modify the rules base by adding or deleting rules, or retrieve the currently stored rules for inspection. Messages relating to administration of the rules base use a special “control” channel to communicate with the engine.

The dynamic router is in common use in financial services software and can even be seen in e-Health information frameworks. [3] However, the rules base is commonly used by network administrators to configure routing to respond to network concerns (such as traffic loads) or application issues (such as preferred client/server pairings). Using the rules base to encode policies based on considerations from the application domain (in this case, health or medicine) suggests putting control of the rules base into the hands of an application domain specialist, such as a policy administrator, rather than in the hands of a network manager.

Typically, the rules engine will match entries in the rules base against information in the message header. If the rules relate to the health domain, then information related to this domain (which we term the *medical context*) should appear in the message header. (There are approaches to analyzing message payload content, but header information, where feasible, is preferred as a conventional and well-structured solution.) Injecting this additional domain information into the message header is the role of the Content Enrichment pattern.

Figure 2 is a conceptual diagram of the relationship between dynamic routing and content enrichment under a client/server configuration. The content enrichment pattern appears as a module in both the client and server environments, where additional information is injected into each message before it encounters the dynamic router. Since the contextual information appears at the point of origin of each message, this is where the content enrichment occurs. (The contextual information is not illustrated on the server side enrichment pattern of figure 2, only for the purpose of abbreviation). Since our objective is to de-couple applications from changes to information management policies, it is important the content enrichment module is independent of the application code. Otherwise, introducing new types of policies with new categories of rules would require extensive application modification.

### 2.3 Apache Camel

Apache Camel is an integration framework based on the Enterprise Integration Patterns advanced by Hohpe and Woolf [2]. It allows the developer to spend less time understanding how each individual application accesses the MOM, and focus on making various systems interoperate through a common framework. The framework is the “glue” for connecting different network nodes together in a seamless fashion.

In its simplest conception Apache Camel is a routing engine builder that allows developers to define their own routing rules for data, determine which data sources are accepted, and determine how to process and transmit data to other destinations. [4] Rather than imposing a fixed data format, Camel provides a collection of high-level abstractions that allow for simplified interactions with systems using the same API (Application Programming Interface) regardless of the protocol or data type the separate applications are using. This is critical to connecting a wide variety of applications that have different development histories.




Fig. 3. Overview of the Camel integration framework from [4]. Reproduced with permission.

Camel comprises routes, processors, and components, all contained within the CamelContext construct. Camel routing engines implement the routes defined using a domain-specific language (DSL). Processors are used to transform and/or manipulate messages en-route, invoking Enterprise Integration Patterns through DSL directives. Components are extension points in Camel used for adding connectivity to other software.

The following Java code constructs a Camel router. It uses the Java Message Service (JMS) [5] as the protocol for the MOM (the usual choice for Camel). JMS promotes loose coupling of applications, provides reliable messaging mechanisms, and allows for asynchronous communication. The DSL appears in the form of a Java package that provides the “from”, “to”, “process” and other DSL directives, as invoked in the code sample to define routing behaviour. Other directives provide implementation of the EIPs.

This router moves messages from JMS queue "A" to JMS queue "B". When a message is placed in queue "A" the router detects the event and executes the coded action to send the message to JMS queue "B". The context indicating why the message was placed in queue "A" or why it is sent to queue "B" is independent of this code and the router construction. Using Camel, the underlying details of processing and transferring data have been abstracted away. For the purpose of illustration, the last ten lines of code simulate messages being placed on queue "A", – since sending messages would normally be part of the application logic. The ten "Test Message: ' + i" messages being sent to queue "A" are routed to queue "B" by Camel, at which point a second Camel routing directive listens for newly received message on queue "B" to be processed, printing the message content to conform the information transfer.

In our demonstration system for dynamic routing using health information policy, Camel provides the EIP implementation that allows dynamic routing required for applying the policy rules. In the previous example, A-B routing was hardcoded for the purpose of a simplified illustration of Camel. Our actual demonstration system uses the DSL directive invoking the EIP dynamic router.

```
public final class CamelExample {

    private CamelExample() {}
    public static void main(String args[]) throws Exception
    {

        // Declare Camel Context
        CamelContext context = new DefaultCamelContext();

        // Configure JMS Broker
        ConnectionFactory connectionFactory =
            new ActiveMQConnectionFactory(
                "vm://localhost?broker.persistent=false");

        // Add JMS Broker to Camel Context
        context.addComponent("test-jms",
            JmsComponent.jmsComponentAutoAcknowledge(
                connectionFactory));

        // Define and add routing rules to Camel Context
        context.addRoutes(new RouteBuilder() {
            public void configure() {
                // Route messages from queue A to queue B
                from("test-jms:queue:A").to("test-jms:queue:B");
                // Setup a listener on queue B that will process
                // messages.
                from("test-jms:queue:B").process(new Processor(){
                    // Print message contents as received
                    public void process(Exchange e) {
                        System.out.println("Received exchange: "
                            + e.getIn());
                    }
                });
            }
        });

        // Camel template - for kicking off exchanges
        ProducerTemplate template =
            context.createProducerTemplate();
        context.start();
        // Send 10 messages to queue A,
        // appending message body text
        for (int i = 0; i < 10; i++) {
            template.sendBody("test-jms:queue:A",
                "Test Message: " + i);
        }
        Thread.sleep(1000);
        context.stop();
    }
}
```

## 2.4 Health Information Policy

Wide adoption of health information applications, such as electronic medical records and electronic prescription services, is predicated on interoperability of these applications – the ability to share and exchange health information when appropriate. A common communication infrastructure for health information applications is therefore critical to advancing the benefits of such systems. One example of such an effort is the Canadian Health Infoway project [6], a government funded initiative which is defining standards for a national infrastructure including a common communication layer, called the Health Information Access Level (HIAL). As an information exchange architecture, the HIAL solution provides for different applications to share and retrieve information stored in health information repositories throughout the system.

There is a tendency to use simple shared database models to govern the complexity of different health information applications: for effective access control, health information resides in a secure repository and is accessed by authenticated data retrieval requests. Security and access control mechanisms can be layered on top of the information access architecture, because the communications model is always a two-party client/server configuration. Under such a model, most security concerns can be addressed by authentication that answers the question: does the party requesting the information have access permission?

Policies regarding the distribution, control and management of information are represented coarsely in such an architecture, by defining or changing access control rights for users or groups of users. Security and privacy solutions and services are deployed when policies are developed using well established security monikers and solutions, such as authentication and identity management. [7]

With such solutions, policies regarding the handling, management and rights regarding the health information of individuals are encoded as information access requests within each software application. Procedures for these policies must be implemented as steps coded in the application software, such as accessing an authentication service, determining the data availability, and making requests to the appropriate data source. In response to a policy change, the applications may need to be re-written and re-deployed, placing additional cost on health care budgets.

Policy changes can be frequent, and involve multiple levels or rules and regulations, from legislation, to institutional policy (for example, within a hospital), all the way down to individual clinical practices. (see [7], at 46) Consider, for example, the case of “reportable” diseases which require notice to a public health authority. The obligation to report such diseases typically falls on the primary care physician. (cf. [8]) The health information system could assist the physician's obligation, by automatically reporting such diagnoses. However, there are variations on regulations, process, and responsible authority between jurisdictions that may have to share health information. These policies may change or be amended fairly quickly, such as in response to a particular epidemic threat. If policy or enforcement entail reprogramming of application code or circumvention of legacy procedures that are embedded in existing application code, then software hinders the implementation of policy changes.

Policy changes below the legislative level are even more frequent. A particular lab or individual may be censured, health service units may be restructured, information workflow may be re-organized, institutional analysis procedures may be

revised: these all have implications for how information is shared and accessed, which we include in our concept of *information policy*.

While we are concerned that information policy may be encoded only in terms of a data access model, there is also a legacy problem related to embedding data policy at the application level. Subsequent re-interpretation of rules or explicit policy changes due to administrative decisions, new regulation, new technologies, innovation in health care strategies or even court decisions may alter obligations or liability. The ability to respond to policy changes is compromised when the application software targets information policy in existence at the time the system was designed. Vendors that design, create and/or implement systems have little incentive to design to accommodate future changes, particularly if they can look forward to being engaged to overhaul their application program each time changes occur.

### 3 Implementing Information Policy

Using the EIPs discussed above, we have designed a prototype system that uses dynamic routing to implement information policy at the communications layer of the system rather than within the individual application software. By relieving application code of the burden of compliance with (changing) information policies, this approach addresses two concerns at once: messaging is no longer restricted to the shared database model, and policy revision becomes an exercise in configuring routing rules at the communications layer, rather than modifying legacy code in the application software.

Routing rules are based on policy relating to the content and purpose of the messages. To apply the EIPs that provide policy-based routing, messages are enriched with the *medical context*. The concept of medical context appears in a variety of applications, including context-aware computing [9], linking of related medical events [10], annotating EHR records with disambiguating context [11], mobile access through user and location context [12], adaptive information for telemedicine communications [13], and hospital applications such as context aware pill containers or hospital beds [14].

Messages enriched with medical context information could be useful for any of these specific applications. Simple applications, such as tracking and logging medical events and data security forensics, are made possible by the mere presence of the medical context information in the messages. However, the work described here is restricted to the advantages of content enrichment with medical context for the purpose of managing information policy, as an application of the dynamic router EIP.

#### 3.1 Routing Based on Medical Context

Figure 4 illustrates a portion of our rules base for the dynamic router EIP under the Apache Camel implementation. At present, the routing rules are simple: messages typically specify their intended destination, and further routing constraints or destinations may be chosen by the dynamic router, by matching properties of the message against the rules base. The matching of rules and distribution of messages to the appropriate queues is provided by the Camel framework; changes to messaging policy are handled by changing the rules base. More complex routing schemes are feasible, but even this simple scheme allows intricate routing behaviours and basic policy

statements to be represented. Messages can be multicast and/or re-directed, so the application layer is not restricted to a simple data access model of communication, even when using this relatively simple set of routing rules.

```
stjohns;Results;dr.farrell.response.queue;tcp://localhost:61616
stjohns;radiology;radiology;tcp://localhost:61612
stjohns;microbio;microbio;tcp://localhost:61620
stjohns;clientreq;Q_RecordRequest;tcp://localhost:61628
stjohns;ReportD;pub.health;tcp://localhost:61624
stjohns;Patient_Record;*meddrid;tcp://localhost:61616
```

**Fig. 4.** Some routing rules as they appear in one test of the policy-oriented rules base

The routing rules supplied to the dynamic router EIP are used to represent information policy. These policy directives (implemented as routing rules) are intended to operate across applications, and may be related to jurisdiction, treatment, nature of the medical event, type of information requested, origin of the request, identity of the individuals or organizations involved, intended use of the information, and other policy related characteristics of the message. This additional information comprises the medical context, injected into message headers by the context enrichment module. Some of this information (such as intended use) is implicit due to the particular application being used. Under our system, however, messages are not routed according to the particular software application, but according to multiple policies that apply regardless of the particular software application in use.

### 3.2 Enriching Messages with Medical Context

The content enrichment EIP (see figure 2) is used to inject medical context information to messages, information that may trigger application of a policy rule by the dynamic router. Information about the source and destination of the message, and the function of the message in terms of data access, updates and requests would already be included in conventional message definition; this information does not need to be added as part of the medical context. Once a message arrives at the dynamic router for delivery, the medical context for that message is matched against the dynamic router's set of rules, to find applicable policy rules that indicate how the message should be routed. The Camel framework provides the dynamic router and the messaging infrastructure, so the new elements are the medical content enrichment and the policy-based routing rules.

Some of the headers deployed in our current prototype are shown in figure 5. When applications access the Camel framework to communicate, the content enricher is called to add the medical context headers to the base messages. As a separate code package or module, the enricher is accessible to all application programs.

The JMS messages used by the Camel framework can be viewed as plaintext data with a separate header section. The information categories shown in figure 5 can be injected by creating a new header section in the JMS message for each desired category. A sample message header after injection in our prototype system is shown in figure 6. The medical context header becomes a part of each JMS message, and is carried and delivered to routers and application code transparently by the JMS



infrastructure. As a consequence of this simple mechanism, new information categories can be trivially added to the system.

**medmsgtype:** Control or Default type. Control messages are sent to the router's control channel.  
**medcommand:** Commands issued to a router.  
**medpolicy:** A policy in the form of routing rules.  
**meddrid:** Healthcare provider identification. Used by router to construct a unique output queue.  
**medptid:** Patient identification.  
**medwhere:** Location of medical event (scoping is currently unstructured).  
**medaction:** Situational information about medical event.

**Fig. 5.** Some categories currently defined to inject medical context information

```
{commandId = 7, responseRequired = true, messageId = ID:jamie-goodyearsmacbook, Local-50300-1240760119906-0:2:1:1:1, originalDestination = null, originalTransactionId = null, producerId = ID:jamie-goodyears-macbook.local- 50300-1240760119906-0:2:1:1, destination = queue://Q_Default, transactionId = null, expiration = 0, timestamp = 1240760120317, arrival = 0, brokerInTime = 1240760120318, brokerOutTime = 1240760120319, correlationId = null, replyTo = null, persistent = true, type = null, priority = 4, groupId = null, groupSequence = 0, targetConsumerId = null, compressed = false, userId = null, content = null, marshalledProperties = org.apache.activemq.util.ByteSequence @f7b44f, dataStructure = null, redeliveryCounter = 0, size = 0, properties = {medptid=555-555-5555, meddrid=dr.farrell.response.queue, medmsgtype= default, medaction=microbio-TBTestReq-sputem-HL7-2.4, medwhere=stjohns}, readOnlyProperties = true, readOnlyBody = true, droppable = false}
```

**Fig. 6.** A JMS message header after injection of medical context information

The dynamic router and the enriched message do not represent new architectural design concepts; in fact, each are “off the shelf” Camel elements based on published and well understood EIPs. Extended message headers are common to all electronic messaging frameworks, not just JMS and the content enrichment EIP. What makes our system different is routing based on medical context information, where the routing configuration is typically considered a network administration or management task limited to technical considerations, such application, network load, client jurisdiction or data protection.

Figure 2 illustrates the combination of the dynamic router and message enrichment. The medical context enrichment module must be installed at the client site where the message originates, as this is where the context is known; it is effectively an add-on to the application. The dynamic router (which may be remote or local) becomes the means for that client to access the communications infrastructure.

## 4 Test Scenario

Our initial implementation effort is framed as a response to systems that seem restricted to a shared database model. For example, CHIRIS (the Canadian Health Infoway Reference Implementation Suite) [15] is intended to simulate the construction of a cross-jurisdictional electronic health record, but is focused on Admit, Dismiss, Transfer (ADT) and EHR viewer applications. Unlike CHIRIS, our design is conceived around medical context and messaging rather than data access and modification. Rather than data access or retrieval, we conceive of messages as parts of a medical narrative, participating in stories about patient care and governed by the directives within a dynamic information policy. Policies are understood as rules about how to treat the messages that constitute the narrative about the patient's health.

The following scenario is taken directly from [1] and traces the operation of our prototype implementation for a particular use case sequence. Specifically, a physician consults with a patient, updates the patient chart, and requisitions a number of lab tests, one of which results in diagnosis of a reportable disease. A more formal description of this use case can be found in [16]. The following description tracks the system operation in terms of the medical context and dynamic router: we are interested in the messaging events, not application software or screen interfaces, but.

Initially, the physician requests a patient record and receives the relevant record. The message transactions are controlled by the router's rules base. The medical context is injected by the software client into a message to the dynamic router. The router, following its rules, will pass the message on to the patient record repository (available as a remote service). The repository site injects its medical context and returns the record to the router as a response message. The response is routed back to the requesting physician. Several rules were involved in this exchange – the absence of any rule to pass the messages would have caused the data retrieval to fail.

In the second portion of the scenario the physician requisitions a chest xray. The rules base contains an entry for xrays, so the request is forwarded to a radiology lab. After the xray procedure, the results are sent back to the requesting physician as a response message with the appropriate medical context injected at the lab site. When this response is routed, a policy rule is matched for updating the patient record, so the lab results are automatically copied by routing a message to the patient record repository as well as the requesting physician. The repository responds by sending a message notifying the primary care physician of the patient record update: in this test scenario, this is the same physician that ordered the lab in the first place. The physician receives two notifications; first the lab results, and second notification of the patient record update. Rules that provide these notifications begin to show the benefits of our proposed architecture: in shared data models, third party notifications occur only if the application developers program them into the application, or they are specified with human intervention. Our design allows notifications to be routed by a simple policy rule change that automatically affects all applications.

To finish the scenario, the physician requisitions a sputum test. The rules base contains an entry for sputum tests, so the request is forwarded to the microbiology lab. The lab results are sent back to the requesting physician, and in our system test we simulate a positive result for a reportable disease. The reportable disease is injected into the message as part of the medical context at the lab before the results

message is sent to the dynamic router for delivery. When the lab message gets to the dynamic router, several rules are matched; the physician is sent a copy, the patient record will be updated, and the public health authority is notified of the reportable disease. The physician will ultimately receive notification of the patient record update and an additional notification from public health that they are aware of the test result. The notification message from public health is also copied to the patient repository according to another routing rule. In this portion of the scenario, the medical context information is used multiple times to provide appropriate messaging. Each time a message is passed, the rules base was consulted to provide directives to where messages should be sent, without any programming of the client applications. Instead, the policy rules base provides a single consistent source of messaging directives.




Fig. 7. All messages in the scenario are sent through a dynamic router

## 5 Conclusions

Our design requires all client applications and services to inject medical context information into their messaging. The test of our prototype dynamic router involved the simulation of different health system deployments, including the lab and record repository services: these were simulated, not actual field deployed systems. In practice, health information applications used at different sites would require a context enrichment module.

A major thrust of this work is to encode information policy enforcement explicitly in the message routing rules, rather than implicitly in the application code. The translation of information management policy into routing rules requires further work: appropriate support tools are needed before manipulation of the routing rules base can be directly managed by a policy administrator instead of a network specialist. We have demonstrated the first step: routing rules that are based on health information policy.

An ancillary consequence of a message-based approach is the flexibility to depart from client/server or shared database model. With the dynamic router possibly redirecting or duplicating messages, additional work will be needed for an appropriate security model. Issues such as routing conflicts or circularity within routing rules will require examination, and new architecture questions arise, such as where to place the dynamic router, what type of redundancy is appropriate, and how intelligent the routing engine needs to be to service the routing rules.

The technical effort in implementing the messaging infrastructure for dynamic routing based on health information policy is negligible: the Apache Camel framework already provides the necessary EIPs and elements. It is simple matter of reconsidering the application of those elements.

**Acknowledgements.** Thanks to Dr. Gerard Farrell and the members of Memorial University of Newfoundland's e-Health Research Unit and Medical Informatics Group for their support of the work herein described.

## References

1. Brown, E., Goodyear, J.: A Flexible Policy Architecture for Medical Information Messaging. In: *BIOSTEC 3<sup>rd</sup> International Joint Conference on Biomedical Engineering Systems and Technologies*, pp. 70–76. INSTICC (2010)
2. Hohpe, G., Woolf, B.: *Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions*. Addison-Wesley Professional, Reading (2003)
3. Krasser, M.: Introduction to the open eHealth integration platform (2009), <http://architects.dzone.com/articles/introduction-open-ehealth>
4. Ibsen, C., Anstey, J., Zbarcea, H.: *Camel in Action, Early Access Edition*. Manning Publications (2009), <http://www.manning.com/ibsen/>
5. Java Message Service, Java Specification Request 914, final release (2003), <http://jcp.org/en/jsr/detail?id=914>
6. EHRS Blueprint, an Inter-operable EHR Framework. Canada Health Infoway Inc., version 1.0 edition (2003), <http://www.infoway-inforoute.ca/>
7. Brown, E., Wareham, H., et al.: *Technology Choices and Privacy Policy in Health Care: Report to the Privacy Commissioner of Canada*. Medical Informatics Group, Memorial University (2007), <http://cpig.cs.mun.ca/TechnologyChoices>
8. *Health Protection and Promotion Act, Revised Statutes of Ontario*, s. 25. Government of Ontario, Canada (1990)
9. Bricon-Souf, N., Newman, C.R.: Context Awareness in Health Care: A review. *International Journal of Medical Informatics* 76(1), 2–12 (2007)
10. De Clercq, E., Bangels, M., France, F.R.: Integration of Electronic Patient Record Context with Message Context. *Studies in Health Technology and Informatics* 107(2), 1028–1032 (2004)
11. Manzoor, S., Ceusters, W.M., Rudnicki, R.: A Middleware Approach to Integrate Referent Tracking in EHR System. In: *American Medical Informatics Association Symposium Proceedings*, pp. 503–507 (2007)
12. Hägglund, M., Scandurra, I., Moström, D., Koch, S.: Bridging the Gap: a Virtual Health Record for Integrated Home Care. *International Journal of Integrated Care* 7, e26 (2007)
13. Doukas, C., Maglogiannis, I., Karpouzis, K.: Context-Aware Medical Content Adaptation Through Semantic Representation and Rules Evaluation. In: *Third International Workshop on Semantic Media Adaptation and Personalization*, Washington, pp. 128–134 (2008)
14. Bardram, J. E.: Applications of Context-Aware Computing in Hospital Work: Examples and Design Principles. In: *ACM Symposium on Applied Computing*, New York, pp. 1574–1579 (2004)
15. CHIRIS User Manual. Canada Health Infoway Inc., version 0.2 edition (2005), <http://sourceforge.net/projects/crrs>
16. Goodyear, J.: *Extending the Health Information Access Layer with a Medical Context Messaging Framework*. Honours Dissertation, Department of Computer Science, Memorial University of Newfoundland (2009)

# A Personal Health Record System for Emergency Case Management

Vassiliki Koufi, Flora Malamateniou, and George Vassilacopoulos

Department of Digital Systems, University of Piraeus  
80, Karaoli & Dimitriou Str., 18534 Piraeus, Greece  
{vassoki, flora, gvass}@unipi.gr

**Abstract.** Recent trends in healthcare delivery have led to a shift towards a patient-centric care model which depends on the evolution of information access and tools beyond Electronic Health Record (EHR) systems controlled by healthcare providers to Personal Health Record (PHR) systems whose architectures are based on the fundamental assumptions that the complete records are centrally stored and that each patient retains authority over access to any portion of his/her record. PHRs offer significant potential to stimulate transformational changes in emergency healthcare delivery since they can provide a complete picture of a person's healthcare record when and where needed. However, a security issue of prominent importance arises which is concerned with the process of granting (revoking) authorization to (from) healthcare professionals without the patient's involvement. This paper presents an approach to automating the authorization propagation process in PHR systems by means of context-aware technology, which is used to regulate user access to data via a fine-grained access control mechanism.

**Keywords:** Personal health records, Information availability, Access control, Emergency, Authorization propagation.

## 1 Introduction

Throughout their lives individuals receive care in different parts of the health care system. This results in patient health data being scattered around disparate and geographically dispersed information systems hosted by different healthcare providers [1], [2]. The lack of interoperability among these systems impedes optimal care as it leads to unavailability of important information regarding patient health status when this is mostly needed (e.g. in case of an emergency).

Recently there has been a remarkable upsurge in activity surrounding the adoption of Personal Health Record (PHR) systems for patients [2]. A PHR is a consumer-centric approach to making comprehensive electronic medical records (EHRs) available at any point of care while fully protecting patient privacy [3]. Unlike traditional EHRs which are based on the 'fetch and show' model, PHRs' architectures are based on the fundamental assumptions that the complete records are held on a central repository and that each patient retains authority over access to any portion of his/her record [3], [4]. Thus an entire class of interoperability is eliminated since the system

of storing and retrieving essential patient data is no longer fragmented. Hence, quality and safety of patient care is enhanced by providing patients and health professionals with relevant and timely information while ensuring protection and confidentiality of personal data.

Providing patients with access to their electronic health records offers great promise to improve patient health and satisfaction with their care, as well as to improve professional and organizational approaches to health care [4]. Although many benefits have been identified, there are many questions about best practices for the implementation of PHR systems [4]. A number of these questions are related to security issues arising in PHR systems.

As any other EHR system, PHR systems require stringent privacy protections to prevent unauthorized access or use [5][6][7]. Most PHR platforms currently deployed (e.g. Microsoft HealthVault [8], ICW LifeSensor [9]) meet these requirements by assigning the patient with the responsibility of granting access to information comprising his/her health record while access to important information (e.g. blood type, allergies etc) is provided to medical staff in case of an emergency by means of an emergency data set. Although this information is valuable while providing first aid to the patient, a more comprehensive view of the his/her health data is required by the medical staff upon arrival to the emergency department of a hospital.

This paper deals with the particular security issue arising in PHR systems which is concerned with the process of granting (revoking) authorization to (from) healthcare professionals without the patient's involvement. This security issue is particularly important in managing emergency cases. To deal with this problem, authorization propagation process is automated by means of context-aware technology, which is used to regulate user access to data via a fine-grained access control mechanism. The latter is a role-based, context-aware access control mechanism that incorporates the advantages of broad, role-based permission assignment and administration across object types, as in role-based access control (RBAC) [10], and yet provides the flexibility for automatically adjusting access permissions on a patient's PHR on the occurrence of unpredictable events (e.g. emergency case).

## 2 Related Work

During the last few years, there has been a growing interest in the utilization of PHR systems as both patients and healthcare organizations realized that their use may entail a number of benefits, such as better access to information, increased patient satisfaction and continuity of care [2], [4]. However, certain barriers to the integration of PHR systems to the clinical practice have been identified, most of them related to security issues [2], [4]. In recognition of these barriers, a number of mechanisms have been developed in an attempt to address several issues mostly regarding access control over the health data comprising a PHR [6], [7], [11]. Some of them are concerned with the provision of access to important healthcare information in case of an emergency.

In Case of Emergency Personal Health Record (icePHR) [12] and My Personal Health Record (myPHR) [13] are applications which, among others, ensure that life saving information is available when mostly needed (i.e. in case of an emergency).

To this end, they provide patients with the ability to upload important health information and then print their own emergency card with information on how to access their own unique, secure web page with this emergency information. However, they don't provide mechanisms for ensuring instant availability of a complete copy of a patient's record to the medical staff treating him/her without the patient's involvement.

The system architecture proposed in this paper utilizes agent technology in an attempt to automate authorization propagation process in cases that a patient is incapable of being involved in this process. To this end, a context-aware access control mechanism has been developed which is triggered when appropriate in order to derive and grant the set of authorizations needed for the treatment of a patient.

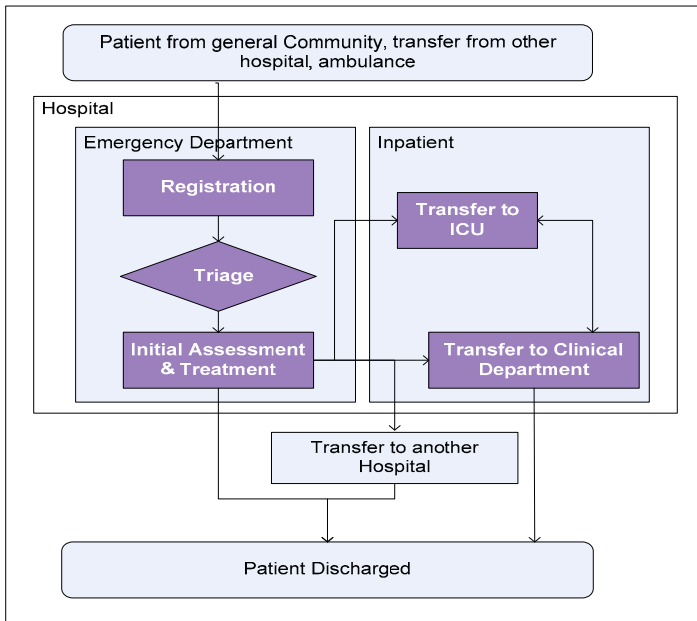
### 3 Motivating Scenario

The basic motivation for this research stems from our involvement in a recent project concerned with designing and implementing a PHR system for the provision of data access at any point of care while fully protecting privacy. This involves providing access to the appropriate people, based on patient wishes, but also granting access to the patient's data in cases where his/her involvement in the authorization propagation process is not feasible. The stringent security needs of the system, where sensitive patient information is used, motivated this work and provided some of the background supportive information for developing the prototype presented in this paper.

Typically, a health district consists of one district general hospital (DGH) and a number of peripheral hospitals and health centers. Suppose a healthcare delivery situation that takes place within a health district where an individual is transferred to a hospital's emergency department (ED). Upon arrival to the ED, the individual is registered as an emergency patient and undergoes a brief triage in order for the nature and severity of his/her illness to be determined. If his/her illness or injury is considered to be serious he/she is seen by a physician more rapidly than the patients with less severe symptoms or injuries. After initial assessment and treatment, the patient is either admitted to the hospital (e.g. to a clinical department or the Intensive Care Unit - ICU), stabilized and transferred to another hospital for various reasons, or discharged [14].

As many emergency department visits are unplanned and urgent, there is a need to ensure that information regarding the longitudinal patient health condition (e.g., problems, allergies, medications, diagnoses, recent procedures, recent laboratory tests) is conveyed to ED physicians automatically upon registration of a patient to an ED. Thus, inefficiencies in care, in the form of redundant testing, care delays, and less-effective treatments prescribed are eliminated and quality of care is enhanced.

Figure 1 shows an indicative high-level view of the patient flow from the time he/she arrives at a hospital's emergency department to the time he/she is discharged. Some of the roles participating in the patient's treatment are physician, nurse, physician assistant (PA), nurse practitioner with specialized training in emergency medicine and in house paramedics and other support staff.



**Fig. 1.** Patient Flow

From an authorization perspective, the following two requirements are of interest here.

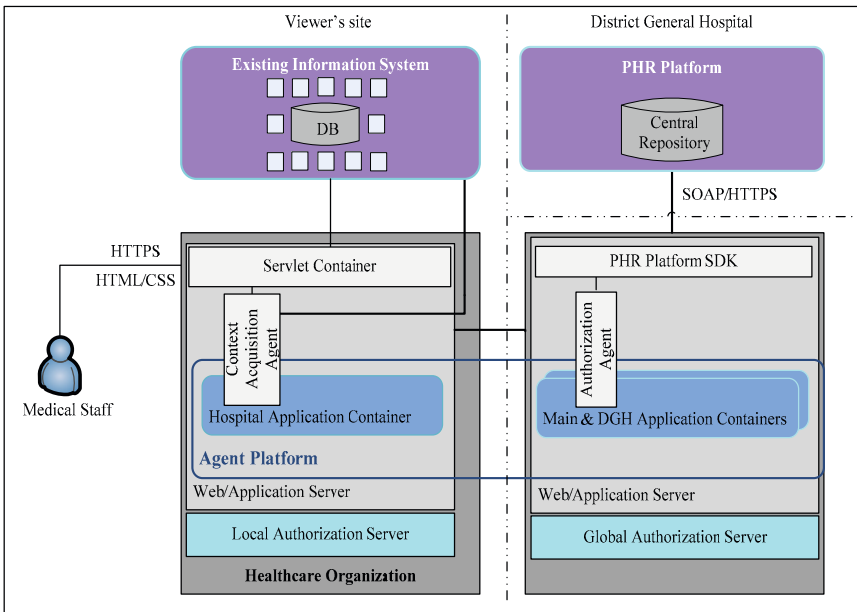
- **Data Access** - A role holder should be allowed to exercise a dynamically determined set of permissions on certain data objects only. For example, a patient's personal physician, if authorized by the patient himself, is allowed to read certain parts of his/her medical record and to update it.
- **Permission Propagation** - Some role holders should receive additional permissions on certain data objects in order to effectively treat the patient but these permissions should be revoked upon patient discharge. For example, in order to form an appropriate care plan, an ED physician should receive the permission to read the complete record of a patient but he/she should not be allowed to retain this permission after the patient has been discharged.

The above requirements suggest that certain data access permissions of the medical staff participating in a patient's treatment may change without the patient's intervention depending on the context (e.g. in the case that an individual is registered as an emergency patient). Moreover, contextual information, such as time and location of attempted access, can influence authorization decisions on certain data objects comprising a patient's PHR. This enables a more flexible and precise access control policy specification that satisfies the least privilege principle by incorporating the advantages of having broad, role-based permissions across data object types, like RBAC, yet enhanced with the ability to simultaneously support the following features: (a) predicate-based access control, limiting user access to specific data objects, and (b) a permission propagation function to specific role holders in certain circumstances.



## 4 System Architecture

The prototype system described here facilitates access to comprehensive patient information which is stored in a central repository. In this environment, a robust security framework is in place in order to ensure that health information follow patients throughout their care in a secure manner and that comprehensive information is made available to appropriate people when this is mostly needed (e.g. in case of an emergency) without the patient's involvement. Figure 2 shows a high-level system architecture, which is described by a three-tier model, comprising of the terminal station used by the medical staff at the department where the patient is being treated (e.g. ED, ICU etc), the PHR platform and the application accessing the PHR platform.



**Fig. 2.** System Architecture

The first tier is the terminal station used by the medical and nursing staff treating the patient. The terminal contains an HTTP(S)-based client, which is the terminal's web browser and provides user interaction with the system.

The second tier of the system architecture is the platform used for the implementation of the PHR system. This supports both patients in actively managing their own health and the medical staff (e.g. physicians) by ensuring the quick and secure availability of a patient's health data such as diagnosis, therapy and prescription data. In such a PHR system access authorization is exclusively granted by the owner (patient) of the record or by a "gatekeeper" he/she assigns (e.g. a relative) [9]. Different read and write permissions can be granted to and be withdrawn from the various users at any time through a terminal station.

The third tier is the application which is distributed among several hosts residing at the DGH and the other healthcare institutions. The infrastructure of this tier consists of the following components:

- **PHR Platform SDK:** It is used for the simple integration of our application into the PHR infrastructure. It provides an Application Programming Interface (API) which can be accessed from JSP/Servlet pages.
- **Agent Platform:** It is the software used for the implementation of the agents which realize the automation of the authorization propagation process in order to support healthcare professionals and frontline staff at the point of care by ensuring instant availability of the complete copy of a patient's medical record.
- **Servlet Container:** It provides a servlet container that hosts and manages the servlets delivering the system functionality. Essentially these servlets provide a web-based front end to the PHR system.
- **Web/Application Server:** It provides the hosting environment to the aforementioned components.


All web transactions are executed under the Secure Socket Layer (SSL) via HTTPS. In addition, security in communication among the agents of the agent platform is ensured by setting up a secure, confidential and mutually authenticated, connection amongst containers of the agent platform by leveraging TLS/SSL support provided by Java [15].

## 5 Security Architecture

The movement towards PHR systems has created new challenges for the sharing of health information in a private and secure manner. In particular, when situations occur where access to medical information is required but patients cannot grant permissions to the medical staff needing the information for treating them, effort should be put in the development and enforcement of a mechanism that automates the authorization propagation process while ensuring privacy and security against unauthorized access to the data.

The number, type and sophistication of tools that protect information in PHR environments are growing at an ever-increasing rate and provide the opportunity to offer health privacy protections beyond those in the paper environment. In many cases, the utilization of role-based access controls is considered as an effective means of limiting access to a patient's information to only those individuals who need it for the patient's treatment.

In our prototype system, a dynamic access control mechanism is incorporated which is based on the role-based access control (RBAC) paradigm and is context-aware. As illustrated in Figure 3, this is described by a two-tier model consisted of a global access control service, residing on a server at the DGH site, and one local access control service, residing at the viewer's site (i.e. any healthcare organization within the health district). Both services use a number of agents for context management.



**Fig. 3.** Security Architecture

The access control mechanism developed is middleware-based and its role is twofold. In particular, it is employed to:

- Grant/revoke authorizations of given subjects to (from) given PHR data objects by taking into account the current context (e.g. medical and nursing staff on duty upon registration of an individual as an emergency patient). In particular, the PHR system where all patient records are stored maintains the authorization data in an authorization base which needs to be updated whenever a patient is registered at (discharged from) a hospital's ED. Thus, the ED medical and nursing staff is granted access to the patient's PHR until the patient is discharged from the hospital's ED. Each member of the staff (i.e. subject) is member of a role. Hence, in order for the relevant authorizations to be determined a set of access control policies are used by means of which role-to-permission assignments are specified.
- Mediate between subjects (healthcare professionals) and objects (PHR data objects) and decide whether access of a given subject to a given object should be permitted or denied according to the context holding at the time of the attempted access (e.g. the time a physician of the ED requests access to a patient's PHR).

In our prototype, users authenticate themselves by using X.509 certificates.

### 5.1 Access Control Policies

In our prototype system, the mapping of roles to the relevant permissions is performed by means of access control policies expressed by using the Core and Hierarchical

RBAC profile of eXtensible Access Control Markup Language (XACML) [16]. These policies are expressed in the form of roles, role hierarchies, privileges and constraints.

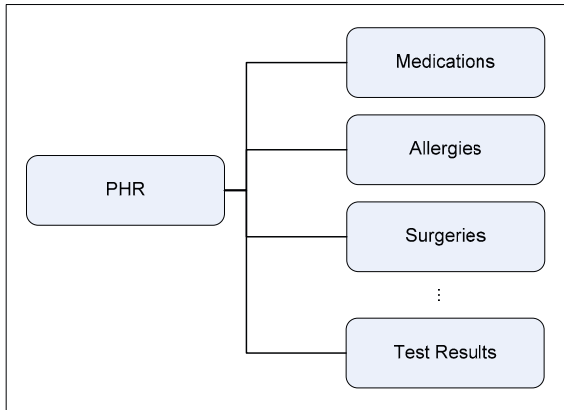
Due to the strict security requirements on medical data comprising a PHR, the specification of access control policies not for the entire record but for its components (i.e. data objects) is of utmost importance. Since the record is organized as a hierarchy, when specifying policies on it the hierarchical resource profile of XACML [16] can be used for the representation of these components. This profile specifies how XACML provides access control for resources that are organized as a hierarchy, such as file systems, XML documents and databases. According to this profile, non-XML data can be represented by a URI of the following form:

```
<scheme>://<authority>/<pathname>
```

where:

- `<scheme>` identifies the namespace of the URI and can be either a protocol (e.g. “ftp”, “http”, “https”) or a file system resource declared as “file”.
- `<authority>` is typically defined by an Internet-based server or a scheme-specific registry of naming authorities, such as DNS, and

`<pathname>` is of the form `<root name>{/<node name>}`. The sequence of `<root name>` and `<node name>` values should correspond to the components in a hierarchical resource.



**Fig. 4.** PHR Data Model

Suppose that the data structure of a PHR is the one illustrated in Figure 4. Then the data object “Allergies” would be represented as follows:

```
https://localhost:8443/PHR/Allergies
```

The policies related to the permissions on data objects a healthcare professional should acquire while treating a patient reside on a server at the DGH site. An excerpt of an access control policy for role “physician” is shown in Figure 5. This is a relatively simple policy that states that an ED physician is authorized to access the complete medical record of each patient he treats. This is specified within the tag `<Resource>` by means of the predicate “all” while the predicate “all” within the tag

<Action> means that the physician has all kinds of permissions on the patient's medical record. Permissions on data objects are dynamically adapted by the constraints imposed by the current context. These are declared within the tag <Condition> and for the role "physician" is whether he/she is requesting patient information using a terminal within the hospital premises.

```

<Resource>
  <ResourceMatch MatchId="&function;string-equal">
    <AttributeValue DataType="&xml;string">all
      </AttributeValue>
    ...
  </ResourceMatch>
</Resource>
...
<Action>
  <ActionMatch MatchId="&function;string-equal">
    <AttributeValue DataType="&xml;string">all</AttributeValue>
    ...
  </ActionMatch>
</Action>
<Condition>
  <Apply FunctionId="&function;string-equal">
    <EnvironmentAttributeDesignator AttributeId="urn:oasis:names:tc:
2.0:environment:
:
      terminal"
        DataType="&xml;string"/>
    <AttributeValue DataType="&xml;string">inPremises</AttributeValue>
  </Apply>
</Condition>

```

Fig. 5. Sample Access Control Policy for Physician

## 5.2 Context Information Management

In our prototype system, the management of context information influencing authorization decisions is performed by a Context Manager. Both the context information model and the Context Manager are described below.

### 5.2.1 Context Information Model

In our prototype system, the contextual information influencing authorization decisions is determined by a pre-defined set of attributes related to:

- the user (e.g. user certificate, user/patient relationship) and
- the environment (e.g. client location and time of attempted access)
- the healthcare provider (e.g. physicians on duty)

For example, the permissions of an ED physician accessing the system via a terminal, are adapted depending on his/her identity (included in his electronic Health Card) as well as the location of the terminal and time of attempted access.

### 5.2.2 Context Manager

Context information is collected by a Context Manager which has been implemented as a multi-agent system. Thus, the Context Manager consists of two kinds of agents:

- **Context Acquisition Agent (CAA):** It is hosted on a server at the site of the health-care organization where the ED belongs and is responsible for the acquisition of the contextual information required for granting authorizations and taking authorization decisions regarding access on the data objects comprising a patient's PHR.
- **Authorization Agent (AA):** It is hosted on a server at the DGH and is responsible for automatically granting (revoking) authorization to (from) healthcare professionals without the patient's involvement. Moreover, it is responsible for managing access to patients' PHRs.

CAA and AA perform the aforementioned tasks in response to a message submitted to them by the corresponding access control service (local or global). Such a message is generated upon each request for access to a patient's PHR. Hence, both CAA and AA implement a cyclic behavior which is continuously running in order to check if a message has been received and process it. This means that the thread of each agent implementing this behavior starts a continuous loop that is extremely CPU consuming. In order to avoid that the agent should execute his behavior (i.e. the number of actions assigned to it) only upon receipt of a new message. To this end, the agent's behaviour is marked as "blocked" so that the agent does not schedule it for execution anymore. When a new message is inserted in the agent's message queue all blocked behaviours becomes available for execution again and the received message is processed.

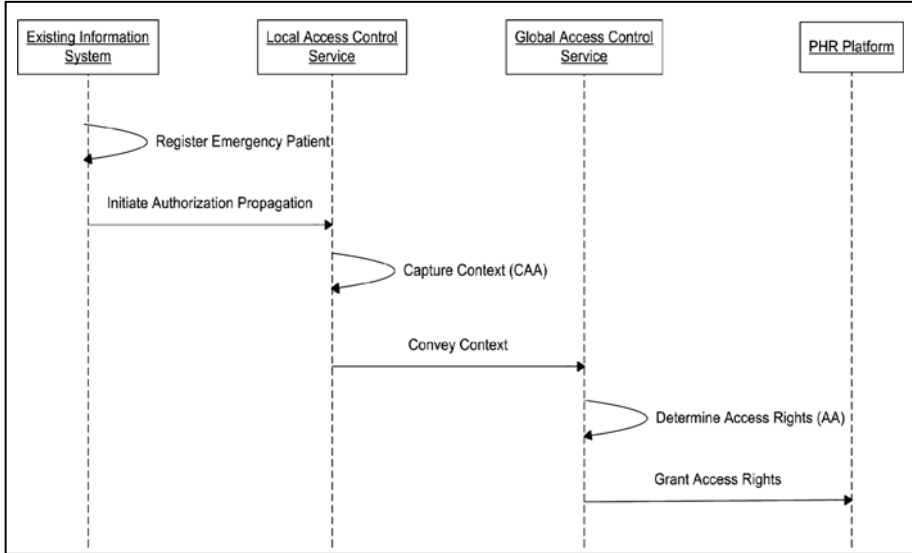
## 6 Implementation Issues

To illustrate the functionality of the proposed architecture, a prototype system has been developed which is based on the case scenario of Section 3.

The prototype implementation of the proposed system and the security services incorporated in it has been developed in a laboratory environment. In our implementation Apache/Tomcat is used as Web/Application Server while the databases used by the existing information systems have been developed using MySQL. The PHR system has been implemented using the ICW Lifesensor Personal Health Record which can store the owner's complete medical information in one convenient and secure location [9]. The patient as owner of the record authorizes health team members or care providers to access their record and assigns specific read and write privileges [9]. ICW Java SDK has been used for the integration of Lifesensor PHR to our application.

The security services used in our application (i.e. the local and global access control services) have been developed using Java Authentication and Authorization Service (JAAS) [17]. These services carry out their tasks by making use of a number of policies and by always taking into account the relevant context which is collected by the Context Manager. The policies declaring access rights on a patient's PHR are expressed by using the Core and Hierarchical profile of eXtensible Access Control Markup Language (XACML) [16]. The Context Manager has been implemented as a

multi-agent system using JADE as a construction and execution environment [15]. JADE is an open-source software framework, aiming at assisting the development and execution of agent-based applications in compliance with the Foundation for Intelligent Physical Agents (FIPA) specifications for interoperable multi-agent systems [18]. Communication between agents, if required, is performed by means of FIPA Agent Communication Language (ACL) [15].



**Fig. 6.** Sequence Diagram for Authorization Propagation Process

Upon arrival to the ED of a hospital, an individual is registered as an emergency patient and the authorization propagation process is triggered in order for the required authorizations to be determined and granted to the medical staff treating him. To this end, the local access control service is invoked which, in conjunction with the local Context Acquisition Agent (CAA), is accessing the local database(s) in order to retrieve the list of the medical staff being on duty at the time. The pieces of information retrieved include starting and ending time of each person's shift. As soon as the information is retrieved, it is communicated to the global access control service which, in conjunction with the Authorization Agent (AA), is determining the corresponding access rights for each person on the list according to a number of XACML policies. The latter are already defined and stored on a server at the DGH site. Finally, the deducted authorizations for each member of the medical staff are granted to him/her by means of the ICW SDK's HealthRecordManager which essentially represents the access to a given personal health record.

The sequence of actions involved in the authorization propagation process is illustrated in Figure 6, by means of a sequence diagram.

After the authorizations have been granted, the nurses and physicians of the ED authenticate themselves in order to gain access to this patient's full medical record by using their credentials (X.509 certificate stored in his electronic Health Card - eHC).

Each access request is handled by the corresponding local access control service which is using CAA to acquire the context holding at the time of the attempted access and forwards the request to the global access control service which in cooperation with the AA decides whether access should be granted or denied to the requesting party. If the requesting party has the required privileges a connection to the corresponding PHR is established and the corresponding part of the patient’s record is provided to him/her.

The sequence of actions performed during access control enforcement is illustrated in Figure 7, by means of a sequence diagram.

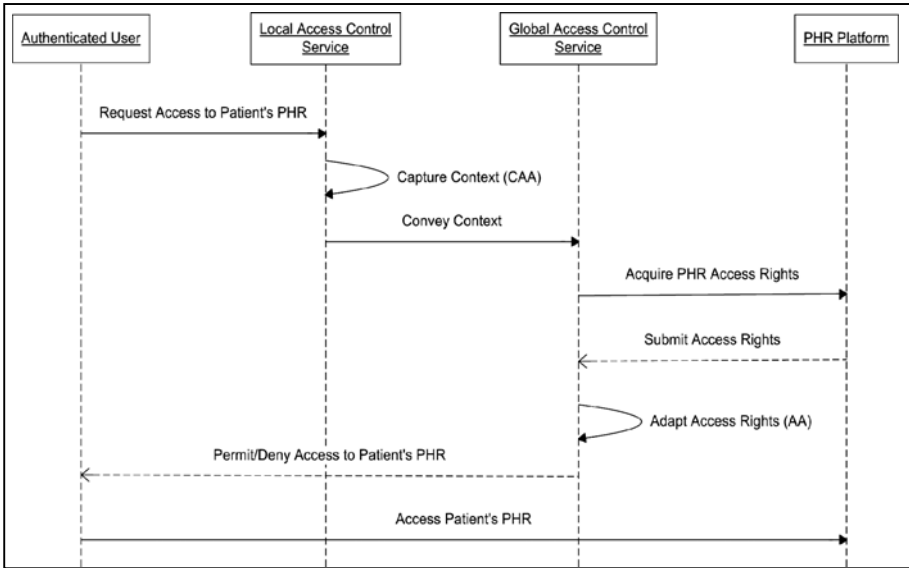


Fig. 7. Sequence Diagram for Access Control Enforcement

After reviewing the patient’s medical record, the ED physician forms the appropriate care plan for the patient under treatment.

## 7 Concluding Remarks

Personal health records can address healthcare information needs as they can provide each person with a complete copy of his medical record. Thus, PHRs constitute a valuable tool for supporting the continuity of care and consequently the quality, access and efficiency of health care delivery. As PHR systems grow in popularity, it is important that they be managed and maintained responsibly without hindering accessibility to important information in cases that it is mostly needed (e.g. emergency cases). Hence, apart from the security and privacy controls which are common to any electronic health record system, in PHR systems a suitable mechanism should be in place that will automate the authorization propagation process without the patients’



involvement. The prototype system presented in this paper deals with this security issue. In particular, a mechanism is presented whereby the process of granting (revoking) authorization to (from) healthcare professionals on patients' PHR is performed without the patient's involvement. To this end, context-aware technology is used. Thus, both clinical and administrative patient data are becoming immediately available to people who need it via accessible, secure and highly usable PHRs, fact that constitutes an enabling factor of the patient-centred shared care.

Since Personal Health Records contain a significant amount of sensitive information, an appropriate legislative and regulatory environment for PHR-based applications is required. The last few years, a number of privacy, safety and security standards and regulations have been specified aiming at ensuring the privacy and security of individually identifiable health information. These include Health Insurance Portability and Accountability Act (HIPAA) [18], Personal Health Information Protection Act in Canada [19], the European Data Privacy Protection Act [20] and its state-specific implementations [21] and strong industry standards (e.g., generic information security standard ISO 17799). However, from a consumer's perspective, the lack of federal privacy protection for confidential health information stored by entities that are not covered by the aforementioned acts, such as commercial PHRs and repositories, remains problematic. Moreover, a number of other issues related to the implementation of PHR-based systems like the one proposed in this paper suggest directions for future work. The most important concern the means used for patient authentication as well as the way medical staff is granted access to medical data in cases where patient registration is performed after the patient has received treatment, as is often the case in EDs.

## References

1. Koufi, V., Vassilacopoulos, G.: HDGPortal: A Grid Portal Application for Pervasive Access to Process-Based Healthcare Systems. In: 2nd International Conference in Pervasive Computing Technologies in Healthcare (PervasiveHealth 2008), Tampere, Finland (2008)
2. Tang, P.C., Ash, J.S., Bates, D.W., Overhage, J.M., Sands, D.Z.: Personal health records: definitions, benefits, and strategies for overcoming barriers to adoption. *JAMIA* 13(2), 121–126 (2006)
3. Lauer, G.: Health Record Banks Gaining Traction in Regional Projects (2009), <http://www.ihealthbeat.org/features/2009/health-record-banks-gaining-traction-in-regional-projects.aspx>
4. Wiljer, D., Urowitz, S., Apatu, E., DeLenardo, C., Eysenbach, G., Harth, T., Pai, H., Leonard, K.J.: Patient accessible electronic health records: exploring recommendations for successful implementation strategies. *J. Med. Internet Res.* 10(4) (2008)
5. Yasnoff, W.A.: Electronic Records are Key to Health-care Reform, *BusinessWeek* (2008)
6. Comini, L., Mazzu, M., Scalvini, S.: Security aspects in electronic personal health record: data access and preservation, Digital Prevention Europe, Briefing Paper (2008)
7. Win, K.T., Susilo, W., Mu, Y.: Personal Health Record Systems and Their Security Protection. *J. Med. Syst.* 30, 309–315 (2006)
8. ICW eHealth Framework, Lifesensor (2009), <http://idn.icw-global.com/solutions/lifesensor/lifesensor.html>
9. National Institute of Standards and Technology (NIST), Role Based Access Control (RBAC) and Role Based Security, <http://csrc.nist.gov/groups/SNS/rbac/>

10. Røstad, L., Nytrø, Ø.: Personalized Access Control for a Personally Controlled Health Record. In: 2nd ACM Workshop on Computer Security Architectures (CSAW 2008), Virginia, USA (2008)
11. Metavante: In Case of Emergency Personal Health Record, <https://www.icephr.com/>
12. My Personal Health Record (MyPHR), <http://myphr.ca/>
13. Wikipedia: Emergency Department, [http://en.wikipedia.org/wiki/Emergency\\_department](http://en.wikipedia.org/wiki/Emergency_department)
14. Java Agent Development Framework, <http://jade.tilab.com/>
15. Organization for the Advancement of Structured Information Standards (OASIS): Core and Hierarchical Role Based Access Control (RBAC) Profile of XACML v2.0, [http://docs.oasis-open.org/xacml/2.0/access\\_control-xacml-2.0-rbac-profile1-spec-os.pdf](http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-rbac-profile1-spec-os.pdf)
16. Java Authentication and Authorization Service, <http://java.sun.com/javase/6/docs/technotes/guides/security/jaas/JAASRefGuide.html>
17. U.S. Congress: Health Insurance Portability and Accountability Act (HIPAA) (1996), <http://www.hipaa.org>
18. Personal Information Protection and Electronic Documents Act, Second Session, Thirty-sixth Parliament, 48–49, Elizabeth II, 1999–2000, Statutes of Canada (2000)
19. European Directive 95/46/EC on the protection of individuals with regard to the processing of personal data and on the free movement of such data (1995)
20. Bundesdatenschutzgesetz (German Federal Data Protection Act), Bundesgesetzblatt 2001, Nr. 23 (2001)

# A Method for Representing and Querying Temporal Information in OWL

Martin J. O'Connor and Amar K. Das

Stanford Center for Biomedical Informatics Research  
Stanford, CA 94305, U.S.A.

**Abstract.** Ontologies are becoming a core technology for supporting the sharing, integration, and management of information sources in Semantic Web applications. As critical as ontologies have become, ontology languages such as OWL typically provide minimal support for modeling the complex temporal information often contained in these sources. As a result, ontologies often cannot fully express the temporal knowledge needed by many applications, forcing users and developers to develop *ad hoc* solutions. In this paper, we present a methodology and a set of tools for representing and querying temporal information in OWL ontologies. The approach uses a lightweight temporal model to encode the temporal dimension of data. It also uses the OWL-based Semantic Web Rule Language (SWRL) and the SWRL-based OWL query language SQWRL to reason with and query the temporal information represented using our model.

## 1 Introduction

The Semantic Web effort [5] aims to provide languages and tools that specify explicit semantic meaning for data and knowledge shared among knowledge-based applications. In particular, the Ontology Web Language (OWL; [23]) and its associated Semantic Web Rule Language (SWRL; [15]) provide a powerful standardized approach for representing information and reasoning with it. Despite the power of these technologies, they have very limited support for temporal information modeling. OWL, for example, provides no temporal support beyond allowing data values to be typed as basic XML Schema dates, times or durations [38]. SWRL includes operators for manipulating these values, but its operators work at a very low level. There are no standard high-level mechanisms to consistently represent and reason with temporal information in OWL.

Because the temporal dimension is central in many information sources, these shortcomings have significant consequences. Primarily, they restrict the complexity of temporal information that can be represented in application ontologies. In addition, they reduce the possibilities for automated temporal validation of the temporal information that they do have. Crucially, these restrictions also limit the temporal expressivity of deductive rules and queries that can be formulated over ontology-encoded data. Formulating these rules and queries thus requires custom solutions using technologies that may not leverage the formal knowledge representation techniques

provided by ontologies. There is a pressing need for solutions that provide robust knowledge-level mechanisms for representing and reasoning with temporal information in OWL ontologies.

## 2 Background

Historically, the importance of time in many applications has driven the development of custom temporal management solutions. In particular, the centrality of the temporal dimension in biomedical data drove early research in the area. One of the first biomedical systems to address the problem was the Time Oriented Database (TOD; [37]). TOD had a three-dimensional view of clinical data, with time represented explicitly as one of the dimensions. Data relating to individual patient visits, for example, were indexed by patient identifier, clinical parameter type, and visit time. TOD supported a set of basic temporal queries that allowed extraction of data values following certain temporal patterns. The Arden Syntax [16] also supports a basic instant-based temporal representation.

Both TOD and the Arden Syntax model time by associating an *instant timestamp* with particular records. An instant timestamp permits a range of simple temporal questions, such as "Did the patient suffer from shortness of breath before the visit?" or "Did the patient receive Ibuprofen last week?" However, associating an *interval timestamp* with data enables more complex queries. Interval timestamps are composed of a start timestamp and a stop timestamp. Later systems used this type of temporal representation. These systems were typically built to operate with relational database systems and exploited the considerable amount of research on temporal database systems in the 1990s.

This research aimed to address the shortcomings of relational databases for representing temporal information. While relational databases can readily store time values, the relational model provides poor support for storing complex temporal information. A simple instant timestamp is all that the relational model provides, and there is no consistent mechanism for associating the timestamp with non-temporal data. For example, if a database row contains some temporal information, there is no indication as to the relationship between it and the non-temporal data in the row. Does the timestamp refer to the time when the information was recorded, or to when it was known? Other shortcomings include no standard way to indicate a timestamp's granularity. As a result, the relational model provides very limited capabilities for temporal querying.

More than a dozen formal extensions to the relational data model were proposed to address the problem. These approaches ultimately led to the development of a consensus query language, called TSQL2 [30]. TSQL2 supports temporal and non-temporal tables. It also provides a temporal relational algebra that can undertake temporal selection of data, temporal joins based on temporal intersection, and temporal catenation of interval time stamps. The TSQL2 query language is compatible with standard SQL, since it is a strict super set of it. Efforts were made to introduce some TSQL2 temporal features into the SQL3 standard [31], but these efforts did not succeed. No complete implementations of TSQL2 were produced, but several temporal query systems were written using its core features. TSQL2-influenced systems include Chronus [1994] and

its later evolution, Chronus II [24]. Like the earlier TOD system [37], both were developed to meet the temporal query requirements of biomedical applications.

Experience with these systems illustrated that the entire TSQL2 specification is not necessary for developing a powerful temporal query language. A few simple language extensions can provide a large increase in temporal expressivity. The most important lesson learned is that a principled temporal model is key to developing the extensions. This model must enforce a consistent representation of all temporal information in a system. One of the important results of the TSQL2 standardization efforts was the convergence on the *valid-time temporal model* [30]. While numerous models were proposed to represent temporal information in relational databases and in other types of information systems, this model was selected because it coupled simplicity with considerable expressivity.

In the valid-time model, a piece of information—which is often referred to as a *fact*—can be associated with instants or intervals denoting the times that the fact is held to be true. Facts have a value and one or more valid-times. Conceptually, this representation means that every temporal fact is held to be true during the time(s) associated with it. No conclusions can be made about the fact for periods outside of its valid-time. The valid-time model provides a mechanism for standardizing the representation of time-stamped data. When this model is used in a relational system, temporal information is typically attached to all tuples in a temporal table. This approach effectively adds a third dimension to two-dimensional relational tables. The valid-time model is not restricted to use in relational systems, however, and can be used in any information system that requires a consistent representation of temporal information.

The valid-time temporal model has been used to model time in the XML domain, primarily by developing extensions to XPath [2, 39]. Extensions to the XQuery language have also been developed to facilitate temporal querying of XML data described using the valid-time model [11].

The importance of time was recognized early in the Semantic Web effort [7]. A variety of approaches have been proposed to represent temporal information in RDF [22] and OWL [23], the two primary Semantic Web languages. The valid-time model predominates in these approaches. One of the earliest RDF systems was BUSTER [35], an information retrieval system that used a temporal reasoning framework to help answer time-oriented questions. Several extensions to the RDF model have been proposed to facilitate representing temporal information in RDF ontologies [12, 28]. Recent systems have described RDF-based temporal query languages that use these extensions. These systems include T-SPARQL [34], which was developed as an extension to the SPARQL RDF query language, and TOQL [3], a SQL-like temporal query language.

A variety of OWL-based temporal systems have also been developed. Unfortunately, adding a temporal dimension to OWL is not straightforward. OWL's logic-based formalism makes it difficult to model dynamically changing information [18]. Also, OWL does not provide any temporal constructs. As with the relational model, a simple instant timestamp representation is all that it supports. Because OWL supports only binary predicates, relations cannot be directly equipped with a temporal argument. Rather than extending OWL's logical model, researchers have attempted to support temporal representations on top of OWL. For example, OWL-Time [14]

proposes an ontology that provides rich descriptions of temporal instants, intervals, durations, and calendar terms. However, this representation is not lightweight and concerns itself with descriptions of individual data elements, rather than building a temporal model to consistently describe all temporal information in a system.

When developing a temporal model on top of OWL, two approaches are common: *reification* and *fluents*. Reification involves introducing a new object plus a relation to associate an entity with its temporal extent or valid-time. A related approach, known as *fluents* [36], defines a time slice to encode the temporal dimension occupied by an entity. Entities are then connected through their time slices, rather than through source ontology relations. Both approaches can be modeled at the user level without modifications to OWL itself. A disadvantage is that the enforcement of temporal semantics is outside OWL, and only limited OWL reasoning can be used by the temporal reasoning processes. Both approaches also require some rewriting of a source ontology to model the temporal dimension, though this process can be relatively straightforward with reification. An additional disadvantage is that these approaches introduce new relations and objects to model the temporal information associated with an entity. As a result, querying temporal information using these approaches can be cumbersome.

Another approach that does not modify OWL is *versioning* [4]. Here, when an ontology is modified, a new version is created to represent the temporal evolution of the ontology. Most implementations adopt a variety of optimization strategies to ensure that entire copies of the ontology are not generated for each new version. However, irrespective of the optimizations adopted, versioning suffers from significant disadvantages. The primary one is that querying for events occurring during particular time intervals can be computationally expensive. Significant information redundancy is also typical of this approach.

A large amount of research has aimed at extending the description logic model underlying OWL to incorporate time [17, 18, 20]. These approaches involve defining new OWL operators and associated semantics. A variety of temporal description logics have been proposed, though there is still no agreement on a standard approach. There has been comparatively little work related to developing query languages for temporal description logics, with very few systems described in the literature [10].

The systems described here demonstrate the variety of ways in which a temporal model can be represented in OWL. Each one involves a variety of tradeoffs. Description logic-based solutions are the most theoretically attractive, because they can fully exploit OWL reasoning functionality. However, they require modifications to OWL itself. User-level models do not involve modification to OWL, but the enforcement of temporal semantics can thus not be handled by OWL reasoners. However, user-level approaches are attractive because they can be readily implemented. User-level models also involve trade-offs. Many have verbose representations of temporal information, requiring the introduction of several new objects to represent each temporal entity. The models also vary in the amount of ontology rewriting that is necessary. Some require an entire ontology redesign, while others can easily be incorporated with existing ontologies.

Irrespective of the approach adopted, a common shortcoming of existing models is their poor support for temporal querying. Unlike efforts in the relational database field that focused on developing expressive temporal query languages, OWL-based

approaches have not emphasized ease of querying. With many user-level OWL models, for example, concise query expression can be difficult because of complex underlying model representations. The development of practical query languages for temporal description logics is an open research challenge. More advanced querying functionalities, such as grouping and aggregation, are also missing from these systems, further limiting expressivity.

### 3 Temporal Model

We have developed a valid-time temporal model in OWL. The model was encoded at the user-level and adopts a reification-based temporal representation mechanism. We chose this approach because it most easily meets the goal of compatibility with existing OWL-based tools. The resulting model can also be easily shared and used in third-party applications. As this paper shows, this model also provides a foundation for a simple, concise, yet expressive temporal query language. We built it by extending an earlier temporal valid-time model [25] and simplified it so that it could more easily be integrated with existing ontologies. This enhanced model was designed to be lightweight, allowing it to be layered on existing OWL ontologies without requiring that they be significantly rewritten. This model concerns itself with time only and provides a simple approach to adding a temporal dimension to existing entities in domain ontologies.

We developed an ontology in OWL to encode this valid-time temporal model [32, 33]. This paper henceforth uses the prefix `temporal` for entities defined in this ontology. The core class that models an entity that can extend over time is represented by an OWL class called `temporal:Fact`. This class is associated with a property called `temporal:hasValidTimes` that holds the time(s) during which the associated information is held to be true. Values of this property are modeled by a class called `temporal:ValidTime`, which has the subclasses `temporal:ValidInstant` and `temporal:ValidInterval`. These subclasses represent instants and intervals, respectively. `temporal:ValidInstant` is associated with the property `temporal:hasTime`, and `temporal:ValidInterval` is associated with the properties `temporal:hasBeginning` and `temporal:hasFinish`. These three properties are of XML Schema type `xsd:DateTime`. Intervals and instances also have granularities associated with them. This association is modeled by the property `temporal:hasGranularity`, with a range class called `temporal:Granularity`. Specific granularities, such as days and minutes, are modeled as instances of this class.

One possible use of the valid instant and interval classes is to take an existing OWL class and add a user-defined property with a range of one of these two classes to it. The choice of class depends on whether one wishes to model an activity that occurs at a single instant in time or one that takes place over an interval of time. Also, if the activity occurs only once, the association will be represented as an OWL functional property, whereas an activity that may repeat can use a non-functional property. For example, consider the case where a user wishes to add a temporal dimension to a blood pressure measurement that is described using a class called `BloodPressure-Measurement`, which has properties for both the systolic and diastolic values. Blood pressures are typically recorded as instantaneous measurements, so the valid instant

class would be the appropriate property range choice here. By using the valid instant class as the range of a user-defined property associated with the measurement class, all instances of `BloodPressureMeasurement` can now use the `temporal:hasTime` and `temporal:hasGranularity` properties associated with the instant. This will allow them to consistently record temporal information associated with the measurement. Similarly, if a user wishes to work with prescriptions using an existing class called `Prescriptions`, they might choose the valid interval class as the range of a user-defined property associated with the class.

A more useful modeling approach is to directly use the `temporal:Fact` class to represent temporal entities. This class can be made the superclass of an existing OWL class in need of a temporal dimension, thus asserting that instances of that class have a temporal extent. For example, if an investigator wishes to take the blood pressure measurements class cited above and model it as a temporal fact, they can simply make the class a subclass of `temporal:Fact`. Instances of this class will now be able to use the `temporal:hasValidTimes` property to store their valid instants as instances of the `temporal:ValidInstant` class. Similarly, the earlier prescriptions class can be modeled as a temporal entity by making it a subclass of the temporal fact class and using `temporal:ValidInterval` to store the temporal intervals associated with it. The granularities of those instants or intervals can also be modeled with the `temporal:hasGranularity` property associated with the `temporal:ValidTime` superclass.

Representing temporal entities as subclasses of the `temporal:Fact` class can clarify the distinction between the temporal and non temporal entities in an ontology. This temporal representation can also coexist with any existing temporal representations in the ontology, and so does not necessitate modifying the temporal component of existing entities. In most cases, existing temporal information will need to be mapped from the source entities to conform to the format encoded by valid-time instants or intervals. This mapping may be non trivial in some cases, but will ensure a consistent representation of temporal information.

## 4 Temporal Querying

Once all temporal information is represented consistently in an ontology, it can then be manipulated using reusable methods. While OWL itself has no temporal operators for manipulating time values, its associated rule language SWRL [15] provides a small set. However, the operators in this set are very basic, and provide simple instant-based comparisons only.

### 4.1 Basic Temporal Queries: Allen's Operators

SWRL provides a mechanism for creating user-defined libraries of custom methods—called *built-ins*—and using them in rules. We have used this mechanism to define a library of methods that implement Allen's [1] interval-based temporal operators. Our library provides built-ins implementing the entire set of the Allen operators. It also supports operations on basic XML Schema temporal types, such as `xsd:date`, `xsd:dateTime`, and `xsd:duration`. Operators to perform granularity conversion and duration calculations at varying granularities are also provided.



The following rule illustrates the use of a built-in defined by this library called `temporal:before`, which can be used to see if one valid time is before another. This rule classifies patients as trial-eligible if they completed any DDI drug therapy before 1999. In this rule, a patient has a property called `hasTreatment` which has a range class that is a subclass of the `temporal:Fact` class and holds a list of valid-time intervals for each treatment.

```
Patient(?p) ^ hasTreatment(?p, ?tr) ^ hasDrug(?tr, DDI) ^
temporal:hasValidTime(?tr, ?trVT) ^ temporal:hasTime(?trVT, ?t)
temporal:before(?t, "1999") → TrialEligible(?p)
```

The library also has a native understanding of the valid-time temporal model and supports an array of temporal operations on intervals defined using the classes in our model. It can thus be used to directly reason about valid time instants and intervals. As a result, basic XSD temporal data values do not have to be extracted from valid time objects. The previous rule can thus be shortened to:

```
Patient(?p) ^ hasTreatment(?p, ?t) ^ hasDrug(?t, DDI) ^
temporal:hasValidTime(?t, ?tVT) ^ temporal:before(?tVT, "1999")
→ TrialEligible(?p)
```

If temporal facts are supported by built-ins, even more concise rules can be written. For example, with this support, a more complex rule indicating that patients are trial-eligible if they had drugs prescribed during the first three months of 2008 that were taken for longer than one week, can be written:

```
Patient(?p) ^ hasTreatment(?p, ?t) ^ hasDrug(?t, ?drug) ^
temporal:overlaps(?t, "2008-1", "2008-3") ^
temporal:duration(?d, ?t, temporal:weeks) ^ swrlb:greaterThan(?d, 1)
→ TrialEligible(?p)
```

Here, the `temporal:overlaps` built-in is supplied with a temporal fact and two literal date values. Internally, it extracts the intervals associated with the temporal fact `t`, constructs an interval from the two supplied dates, and then performs the temporal comparison of the intervals.

The temporal built-ins in our library can take any combination of temporal facts, valid-time instants, valid-time intervals, or XSD date or datetime literal values. As this paper will show, the ability of built-ins to directly reason with objects defined by the temporal model can significantly reduce the number of clauses required to express temporal criteria, resulting in considerably more concise rules. It can also free users from concerns about the low-level representation details of the temporal model.

In addition to being able to write temporal rules, the ability to write temporal queries on an ontology is also desirable. A SWRL-based query language called the Semantic Query-Enhanced Web Rule Language [26] provides such support. Using built-ins, SQWRL defines a set of SQL-like query operators that that can be used to construct retrieval specifications for information stored in an OWL ontology. These operators are used in the consequent of a SWRL rule to format the information matched by a rule antecedent. This antecedent is effectively treated as a pattern specification for the query. The prefix `sqwrl` is conventionally used for SQWRL built-ins. The core built-in defined by SQWRL is `sqwrl:select`. This built-in takes one or more arguments, which are typically variables used in the antecedent of a rule, and builds an internal table using the arguments as the columns of the table.

For example, the earlier rule to determine trial-eligible patients can be rewritten as a query as follows:

```
Patient(?p) ^ hasTreatment(?p, ?t) ^ hasDrug(?t, DDI) ^
temporal:before(?t, "1999")
→ sqwrl:select(?p)
```

This query will return a table with one column listing all patients who have completed a DDI drug therapy before 1999.

SQWRL queries have access to all built-in libraries available to SWRL and can thus be used to perform temporal queries using the full range of built-ins in the temporal library.

## 4.2 More Advanced Temporal Queries: Grouping and Aggregation

Temporal querying capabilities more advanced than the ones noted above are typically required by many applications. In addition to support for basic interval manipulations, many queries need more complex selection of results. For example, queries such as *List the first three doses of the drug DDI* or *Return the most recent dose of the drug DDI* are common. Because of OWL and SWRL's open world assumption, expressing these types of temporal queries on OWL ontologies can be difficult.

These queries require closure operations for correct formulation. The core SQWRL operators support some degree of closure when querying, and do so without violating OWL's open world assumption. As shown, queries like *List all patients in an ontology who had drug treatments for longer than three months* can be expressed fairly directly. However, queries with more complex closure requirements cannot be expressed using the core operators. For example, the query *List the first three DDI treatments for each patient* is not expressible. Basically, while the core operators support basic closure operations, no further operations can be performed on the results of these operators. Queries with negation or complex aggregation functionality are similarly not expressible using the core operators.

We have extended SQWRL to support the closure operations necessary for these types of temporal queries. Two types of collections are supported: *sets* and *bags*. As might be inferred, sets do not allow duplicate elements, whereas bags do. A built-in called `sqwrl:makeSet` is provided to construct a set. Its basic form is:

```
sqwrl:makeSet(<set>, <element>)
```

The first argument of this set construction operator specifies the set to be constructed. The second specifies the element to be added to the set. This built-in constructs a single set for a particular query and will place the supplied element into the set. If a variable is specified in the element position, then all bindings for that variable in a query will be inserted into the set. A built-in called `sqwrl:makeBag` is provided to construct a bag. Its basic form is:

```
sqwrl:makeBag(<set>, <element>)
```

The scope of each collection is limited to the query that contains it. Collection operators, such as, for example, `sqwrl:size`, can be applied to these collections. The results of these operators can be used by built-ins in the query, thus allowing access to the results of the closure operation. Two new SQWRL clauses are provided to contain

these collection construction and manipulation operators. The *collections construction* clause comes after a standard SWRL pattern specification and is separated from it using the  $\circ$  character. It is followed by a *collections operation* clause that contains built-ins that operate on the collections, and is again separated from it by the  $\circ$  character. In outline, a SQWRL query with collections operators looks as follows:

```

<SWRL Pattern Specification>  $\circ$ 
<Collections Construction Clause>  $\circ$ 
<Collections Operation Clause>
   $\rightarrow$  <Select Clause>

```

The construction clause can only contain SQWRL collection construction operators, such as `sqwrl:makeSet` and `sqwrl:makeBag`. The collection operators clause can contain only collection operators, such as `sqwrl:size`, in addition to built-ins that operate on the results of these operations. It may not contain other SWRL atom types.

Collections support basic closure operations over the entire ontology. However, the earlier query to list the first three DDI treatments for each patient is still not expressible using this approach. Additional collection operators to allow more complex queries that group related entities are also required. This additional expressivity is supplied by collections that are partitioned by a group of arguments. A built-in called `sqwrl:groupBy` provides this functionality. The general form of this grouping is:

```

sqwrl:makeSet(<set>, <element>) ^ sqwrl:groupBy(<set>, <group>)
or
sqwrl:makeBag(<bag>, <element>) ^ sqwrl:groupBy(<bag>, <group>)

```

This group can contain one or more entities. The first argument to the `sqwrl:groupBy` built-in is the collection, and the second and (optional) subsequent arguments are the entities to group by. Only one grouping can be applied to each collection.

Using this mechanism, the construction of a set of treatments for each patient can be written:

```

Patient(?p) ^ hasTreatments(?p, ?t)  $\circ$ 
sqwrl:makeSet(?s, ?t) ^ sqwrl:groupBy(?s, ?p)

```

Here, a new set is built for each patient matched in the query and all treatments for each one are added to that patient's set.

SQWRL also provides standard operators such as `sqwrl:union`, `sqwrl:difference`, `sqwrl:intersection`, and so on. These operators employ standard set semantics and generate sets. Queries can use them to examine the results of two or more closure operations, permitting the creation of queries that are far more complex. Putting elements into collections effectively provides a closure mechanism. Clearly, two-phase processing is required for these queries—a query cannot, say, determine how many elements are in a collection until it has been constructed. As mentioned, the language restricts the atoms that are processed in the second phase to collection built-ins and other built-ins that operate on the results of these collection built-ins. That is, the first phase of query execution is analogous to standard rule execution. The second phase consists of operations on the collections constructed as a result of that execution.

More complex groupings have multiple grouping entities. For example, to make a set of the start times of each patient's treatment, the set construction operator must be supplied with both patient and treatment grouping arguments:

```
Patient(?p) ^ hasTreatment(?p, ?t) ^
temporal:hasValidTime(?t, ?vt) ^ temporal:hasStartTime(?bt, ?start) °
sqwrl:makeSet(?s, ?start) ^ sqwrl:groupBy(?s, ?p, ?t)
```

Here, a set is constructed for each patient and treatment combination, and all the start times for the combination are added to the set.

Ordinal selection or aggregation operators can be applied to a collection if its elements have a natural ordering. These operators include `sqwrl:min`, `sqwrl:max`, `sqwrl:avg`, and so on. For example, a query to return the time of the first treatment for each patient can be written:

```
Patient(?p) ^ hasTreatment(?p, ?t) ^
temporal:hasValidTime(?d, ?vt) ^ temporal:hasStartTime(?vt, ?start) °
sqwrl:makeSet(?s, ?start) ^ sqwrl:groupBy(?s, ?p, ?t) °
sqwrl:min(?first, ?s) ^ temporal:equals(?first, ?start)
→ sqwrl:select(?p, ?start)
```

The result is a list of patients and the time of the first treatment for each one.

### 4.3 More Advanced Temporal Queries: Temporal Collections

Combining SQWRL's collection operators with the temporal valid-time model can provide support for very powerful selection operators on temporal results. Using SQWRL's grouping mechanism, aggregation operators such as `earliest`, `latest`, and so on can be supported. Supporting these operators is a key requirement for a temporal query language [8, 29].

We have further extended the temporal built-in library to support these types of collections. The library now supports the same set of construction and manipulation operators as SQWRL, but allows only temporal facts to be placed in the collections. The library natively understands the interval-based valid-time model underlying the facts placed in collections and provides collection operations on the collections. Additional temporal collection operators, such as `temporal:first`, `temporal:firstN`, `temporal:last`, `temporal:lastN`, `temporal:nth`, and so on, are also provided. Operators applied to these temporal collections consider the interval-based semantics of the entities in the collections. If two collections are merged, for example, intervals belonging to value-equivalent entities are merged. This process is known as *coalescing* [6]. The standard Allen temporal operators can also be applied to collections, facilitating queries such as *Were all DDI prescriptions before all AZT prescriptions?*

Consider, for example, a query to return the first treatment for each patient in an ontology, together with drug and dosage information. Again, assuming that each patient has a treatment property that holds a treatment class containing drug and dosage information which is modeled as a temporal fact using the temporal ontology, the query can be expressed as follows:

```
Patient(?p) ^ hasTreatment(?p, ?tr) ^
hasDrug(?tr, ?drug) ^ hasDose(?tr, ?dose) °
temporal:makeSet(?trs, ?tr) ^ temporal:groupBy(?trs, ?p) °
temporal:first(?ftr, ?trs) ^ temporal:equals(?ftr, ?tr)
→ sqwrl:select(?p, ?tr, ?drug, ?dose)
```

A query to return the first three DDI treatments for each patient, together with dosage information, can be written:

```
Patient(?p) ^ hasTreatment(?p, ?tr) ^
  hasDrug(?tr, DDI) ^ hasDose(?tr, ?dose) °
temporal:makeSet(?trs, ?tr) ^ temporal:groupBy(?trs, ?p) °
temporal:firstN(?f3tr, ?trs, 3) ^ temporal:equals(?f3tr, ?tr) →
  sqwrl:select(?p, DDI, ?dose)
```

Here, the `temporal:firstN` built-in is used to select the first three treatments from a patient's treatment set.

A query to return the most recent DDI treatment for each patient, together with dosage information can be written:

```
Patient(?p) ^ hasTreatment(?p, ?tr) ^
  hasDrug(?tr, DDI) ^ hasDose(?tr, ?dose) °
temporal:makeSet(?trs, ?tr) ^ temporal:groupBy(?trs, ?p) °
temporal:last(?ltr, ?trs) ^ temporal:equals(?ltr, ?tr) →
  sqwrl:select(?p, DDI, ?dose)
```

Here, the `temporal:last` built-in is used to select the most recent treatment from each patient's treatment set.

As can be seen from these examples, natively supporting the valid time-model in collections allows expressive, yet relatively concise temporal queries. With SQWRL's grouping mechanism, queries with advanced aggregation requirements can be expressed directly in the languages. More advanced query functionality can also be supported by combining rules and queries to incrementally generate intermediate results at successively higher levels of abstraction.

## 5 Conclusions

We have created a lightweight, yet expressive temporal model that can be used to encode the temporal dimension of data in OWL ontologies. Our model has been designed to be integrated with existing ontologies with minimal redesign of them. It facilitates the consistent representation of temporal information in ontologies, thus allowing standardized approaches to performing temporal reasoning and temporal queries. We used SWRL and the SWRL-based OWL query language SQWRL to show how knowledge-level temporal rules and queries can be constructed on the information in the ontologies. The primary goal of our approach was to develop a concise, yet expressive query language. To date, OWL-based temporal research efforts tend to emphasize temporal representation, with comparatively little emphasis on methods to query the information in them. A particular failing of existing systems is poor support for queries with aggregates, a failing tends to limit the types of queries than be expressed. This paper shows how we have provided these functionalities by developing temporal extensions to SQWRL and demonstrates how the resulting language provides powerful mechanisms for expressing complex temporal queries. In particular, we show that extending SQWRL with collection operators that can be directly applied to data described using the temporal model provides a high degree of expressivity.

We used an initial version of the temporal valid-time model described here to encode temporal information collected during a national clinical trials project [25]. Other researchers have reported using our model in a hypertension management application to identify patients who satisfy a set of evidence-based criteria for quality improvement potential [21]. We are currently using the updated model with the recent set-based SQWRL extensions to reason with breast cancer image annotation for tumor assessment [19, 27].

A possible shortcoming of our approach is that all temporal information in a source ontology must be transformed to conform to the valid-time model. This process can be time-consuming, and typically requires considerable domain expertise. However, the mapping requirement is not unique to our method. If principled temporal reasoning mechanisms are to be applied to temporal information, some sort of mapping process to regularize the information is nearly always required, irrespective of the final reasoning processes.

An additional possible shortcoming is that complex temporal rules and queries can become difficult to maintain and extend as their numbers increase. We are developing management tools to tackle this problem [13].

The methodologies and tools described in this paper aim to enhance the ability of software developers and investigators to encode critical forms of temporal knowledge in their applications. This knowledge can be represented directly in domain ontologies, facilitating much higher-level analyses than would be possible with lower-level techniques. Ultimately, working at the knowledge level will enable investigators to make better sense of the complex temporal patterns typical in many domains.

**Acknowledgements.** This research was supported in part by grant R01LM009607 from the National Library of Medicine.

## References

1. Allen, J.F.: Maintaining knowledge about temporal intervals. *Communications of the ACM* 26(11) (1983)
2. Amagasa, T., Yoshikawa, M., Uemura, S.: A Data model for temporal XML documents. In: Ibrahim, M., Küng, J., Revell, N. (eds.) *DEXA 2000*. LNCS, vol. 1873, pp. 334–344. Springer, Heidelberg (2000)
3. Baratis, E., Petrakis, E.G.M., Batsakis, S., Maris, N., Papadakis, N.: TOQL: Temporal ontology querying language. In: Mamoulis, N., Seidl, T., Pedersen, T.B., Torp, K., Assent, I. (eds.) *SSTD 2009*. LNCS, vol. 5644, pp. 338–354. Springer, Heidelberg (2009)
4. Bedi, P., Marwah, S.: Versioning OWL ontologies using temporal tags. In: *World Academy of Science, Engineering, and Technology* (March 27, 2007)
5. Berners-Lee, T., Hendler, J., Lassila, O.: *The Semantic Web*. Scientific American (2001)
6. Bohlen, M.H., Snodgrass, R.T., Soo, M.D.: Coalescing in temporal databases. In: *Proceedings of the International Conference on Very Large Databases, Mumbai, India* (1996)
7. Bry, F., Koch, C.: Querying the web reconsidered: design principles for versatile web query languages. *Journal of Semantic Web and Information Systems* (2005)
8. Chatziantoniou, D., Ross, K.A.: Querying multiple features of groups in relational databases. In: *Proceedings of the International Conference on Very Large Databases* (1996)

9. Das, A.K., Musen, M.A.: A temporal query system for protocol-directed decision support. *Methods of Information in Medicine* 33, 358–370 (1994)
10. Frasincar, F., Milea, V., Kaymak, U.: tOWL: integrating time into OWL. In: *Semantic Web Information Management*. Springer, Heidelberg (2010)
11. Gao, C., Snodgrass, R.: Temporal slicing in the evaluation of XML queries. In: *29<sup>th</sup> International Conference on Very Large Databases*, Berlin, Germany (2003)
12. Gutierrez, C., Hurtado, C.A., Vaisman, A.: Introducing time into RDF. *IEEE Transactions on Knowledge and Data Engineering* 19(2), 207–218 (2007)
13. Hassanpour, S., O'Connor, M.J., Das, A.K.: Exploration of SWRL Rule Bases through Visualization, Paraphrasing, and Categorization of Rules. In: Governatori, G., Hall, J., Paschke, A. (eds.) *RuleML 2009*. LNCS, vol. 5858, pp. 246–261. Springer, Heidelberg (2009)
14. Hobbs, J.R., Pan, F.: An ontology of time for the Semantic Web. *ACM Transactions on Asian Language Processing (TALIP): Special issue on Temporal Information Processing* 3(1), 66–85 (2004)
15. Horrocks, I., Patel-Schneider, P.F., Boley, H., Tabet, S., Grosz, B., Dean, M.: SWRL: a Semantic Web rule language combining OWL and RuleML. In: *W3C* (2004)
16. Hripcsak, G., Ludemann, P., Allan Pryor, T., Wigertz, O.B., Clayton, P.: Rationale for the Arden Syntax. *Computers and Biomedical Research* 27, 291–324 (1994)
17. Kim, S.K., Song, M.Y., Kim, C., Yea, S.J., Jang, H.C., Lee, K.C.: Temporal Ontology Language for Representing and Reasoning Interval-Based Temporal Knowledge. In: Domingue, J., Anutariya, C. (eds.) *ASWC 2008*. LNCS, vol. 5367, pp. 31–45. Springer, Heidelberg (2008)
18. Krieger, H.-U.: Where temporal description logics fail: Representing temporally-changing relationships. In: Dengel, A.R., Berns, K., Breuel, T.M., Bomarius, F., Roth-Berghofer, T.R. (eds.) *KI 2008*. LNCS (LNAI), vol. 5243, pp. 249–257. Springer, Heidelberg (2008)
19. Levy, M., O'Connor, M.J., Rubin, D.L.: Semantic reasoning with image annotations for tumor assessment. In: San Francisco, C.A. (ed.) *AMIA Annual Symposium*, San Francisco, CA (2009)
20. Lutz, C., Wolter, F., Zakharyashev, M.: Temporal description logics: a survey. In: *15th International Symposium on Temporal Representation and Reasoning* (2008)
21. Mabotuwana, T., Warren, J.: An ontology-based approach to enhance querying capabilities of general practice medicine for better management of hypertension. *Artificial Intelligence in Medicine* 47(2), 87–103 (2009)
22. Manola, F., Miller, E.: RDF Primer. In: *W3C Recommendation* (2004)
23. McGuinness, D.L., van Harmelen, F.: OWL web ontology language overview. In: *W3C* (2004)
24. O'Connor, M.J., Tu, S.W., Musen, M.A.: The Chronus II temporal database mediator. In: *AMIA Annual Symposium*, San Antonio, TX, pp. 567–571 (2002)
25. O'Connor, M.J., Shankar, R.D., Parrish, D.B., Das, A.K.: Knowledge-data integration for temporal reasoning in a clinical trial system. *International Journal of Medical Informatics* 78(1), S77–S85 (2009)
26. O'Connor, M.J., Das, A.K.: SQWRL: a query language for OWL. In: *OWL: Experiences and Directions (OWLED)*, Fifth International Workshop, Chantilly, VA (2009)
27. O'Connor, M.J., Das, A.K.: Semantic reasoning with XML-based biomedical information models. In: *13th World Congress on Medical Informatics (MedInfo 2010)*, Cape Town, South Africa (2010)
28. Pugliese, A., Udreă, O., Subrahmanian, V.S.: Scaling RDF with time. In: *WWW Conference*, pp. 605–614 (2008)

29. Rafiq, M.I., O'Connor, M.J., Das, A.K.: Computational method for temporal pattern discovery in biomedical genomic databases. In: IEEE Computational Systems Bioinformatics Conference (CSB 2005) (2005)
30. Snodgrass, R.T.: The TSQL2 Temporal Query Language. Kluwer, Boston (1995)
31. Snodgrass, R.T., Böhlen, M.H., Jensen, C.S., Steiner, A.: Transitioning temporal support in TSQL2 to SQL3. In: Etzion, O., Jajodia, S., Sripada, S. (eds.) Dagstuhl Seminar 1997. LNCS, vol. 1399, pp. 150–194. Springer, Heidelberg (1998)
32. SWRL Temporal Built-ins (2010), <http://protege.cim3.net/cgi-bin/wiki.pl?SWRLTemporalBuiltIns>
33. SWRL Temporal Ontology (2010), <http://swrl.stanford.edu/ontologies/built-ins/3.3/temporal.owl>
34. Tappolet, J., Bernstein, A.: Applied Temporal RDF: Efficient Temporal Querying of RDF Data with SPARQL. In: Aroyo, L., Traverso, P., Ciravegna, F., Cimiano, P., Heath, T., Hyvönen, E., Mizoguchi, R., Oren, E., Sabou, M., Simperl, E. (eds.) ESWC 2009. LNCS, vol. 5554, pp. 308–322. Springer, Heidelberg (2009)
35. Visser, U.: Intelligent Information Integration for the Semantic Web. LNCS (LNAI), vol. 3159. Springer, Heidelberg (2004)
36. Welty, C., Fikes, R.: A reusable ontology for fluents in OWL. In: Formal Ontology in Information Systems: Proceedings of the Fourth International Conference (FOIS 2006), pp. 226–336. IOS Press, Amsterdam (2006)
37. Wiederhold, G.: Databases for healthcare. Lecture Notes in Medical Informatics. Springer, Heidelberg (1981)
38. XML Schema (2009), <http://www.w3.org/TR/xmlschema11-1/>
39. Zhang, S., Dyreson, C.E.: Adding Valid Time to XPath. In: Bhalla, S. (ed.) DNIS 2002. LNCS, vol. 2544, pp. 29–42. Springer, Heidelberg (2002)



# Designing a Telemedical System for Cardiac Exercise Rehabilitation

Florian Klompmaker<sup>1</sup>, Clemens Busch<sup>2</sup>, Karsten Nebe<sup>1</sup>,  
Andreas Bleiker<sup>3</sup>, and Detlev Willemsen<sup>2</sup>

<sup>1</sup> University of Paderborn, C-LAB, Fürstenallee 11, 33102 Paderborn, Germany  
<sup>2</sup> Schüchtermann Schiller'sche Kliniken GmbH & Co KG

Ulmenallee 5-11, 49214 Bad Rothenfelde, Germany  
<sup>3</sup> Siemens IT Solutions and Services, C-LAB, Fürstenallee 1133102 Paderborn, Germany  
{Florian.Klompmaker, Karsten Nebe}@c-lab.de,  
{cbusch, dwillemsen}@schuechtermann-klinik.de,  
andreas.bleiker@siemens.com

**Abstract.** Since cardiac diseases cause the major amount of money spent in medical care, IT solutions that support patients in the rehabilitation phase are desirable. Such systems can motivate and guide patients in order to exercise frequently and hence to avoid another cardiac event. Such telemedical applications have possess a very special and heterogeneous user group. Therefore User Centered Design is crucial. We present a system setup and user interface design that was prototypically build for this kind of patients.

**Keywords:** User centered design, User interface design, Rehabilitation, Ergometer, Telemedicine, Training, Supervision.

## 1 Introduction

Current surveys of the IT inter-trade organization Bitkom showed that 59.8% of all Germans beyond the age of 65 would like to use telemedicine to extend living in their familiar environment. We estimate similar values for whole Western Europe and North America. Further on 58% of the interviewed persons indicated that they would make use of alarm systems like tumble sensors, ECG or apnea measurement when they are in need of care. It is shown, that even elderly people do not perceive this observation as a problem but as assistance [1]. Therefore there is the demand on politicians and the public health sector to enable age-based assistance systems for self-determined living at home. Since the required technology is available, assumption of costs for applications like tele-monitoring and tele-homecare has to be provided by health insurances. The health care system can benefit from such systems since they help to reduce or avoid expensive hospitalizations [2,3]. In the past years costs in health care exploded. Thereby cardiac diseases caused the major amount of money spent in medical care.

Consequently intelligent IT solutions that provide opportunities for prevention and secondary prevention of cardiac diseases are crucial. Such systems must come along

with clearly arranged and intuitively operable user interfaces, because patients with cardiac diseases are often elderly people that are not that familiar with information technology. It is a major challenge to design usable telemedicine applications for this target group. To face this challenges a user centered design approach is necessary. Potentially users that are interacting with telemedicine applications often suffer under their current state of health. Therefore the system should guide them carefully. They often have age-related visual impairments and in the setup described later on they are quite far away from the screen. These facts have to be considered when designing user interfaces for such patients.

### **1.1 Rehabilitation after Cardiac Events**

Rehabilitation is well established and highly effective to restore health as well as influence the risk factors after a cardiac event. Yet Studies have shown that one year after the cardiac event and rehabilitation phase II modifiable risk factors have not enhanced, but deteriorated in many cases like blood pressure and weight management [4, 5]. One of the most important preventive strategies is exercise training based on aerobic endurance performance such as cycling, walking, running or indoor ergometer training. Although there is an opportunity to join a phase III rehabilitation exercise group called “Herzgruppe” in Germany only 13 – 40% of all patients attending phase II rehabilitation prolong there secondary prevention in a controlled setting. There are multifarious reasons like a lack of availability near the home, a lack of desire for group exercise training, scheduling conflicts and others for not joining controlled programs [6]. A telemedical setting which enables patients to exercise at their own individual ability and at their own schedule could face the reasons for not attending a controlled setting. The OSAmI-D telemedical setting presented in this article allows a patient to schedule individual training sessions and/or do exercise training on her own with her predefined individual exercise regime.

In Germany, the generation to which most heart patients belong typically lives in two-people households with their partners. Their children already left home, leaving the household with enough room for placing an ergometer or other telemedical devices, e.g. in the rooms formerly used by their children. Especially younger heart patients (50-75) are often socially well-integrated and appreciate social exchange and activity. From their physical condition, they still have the ability to use alternative trainings to the ergometer, e.g. outdoor and social activities such as hiking, biking, travel. They have, however– apart from their heart condition – some physical impairments, e.g. reduced agility, motoric ability and eyesight. When designing telemedical applications, developers have to take the properties of the user group into account.

### **1.2 User Centered Design of Telemedicine Applications**

User Centered Design (UCD) is an established methodology in the software-industry that focuses on the users of a future system and aims to create solutions that fits the users needs, their requirements and supports their tasks and goals. The usability of products gains in importance not only for the users of a system but also for manufacturing organizations. According to Jokela, the advantages for users are far-reaching and include increased productivity, improved quality of work, and increased user

satisfaction. Manufacturers also profit significantly through a reduction of support and training costs [7]. The quality of products ranks among the most important aspects for manufacturers in competitive markets and the software industry is no exception to this. One of the central quality attributes for interactive systems is their usability [8] and the main standardization organizations (IEEE 98, ISO 91) have addressed this parameter for a long time [9]. In recent years more and more software manufacturer consider the usability of their products as a strategic goal due to market pressures. Consequently, an increasing number of software manufacturer are pursuing the goal of integrating usability practices into their software engineering processes [10,11].

In order to provide IT-supported systems for prevention and secondary prevention of cardiac diseases it is important to know the users of the system. In fact the majority of heart patients is over 50 and therefore belongs to a generation who already completed its education before modern ICT technologies emerged and became omnipresent in daily life. For this generation, the depth to which knowledge about basic ICT interaction and infrastructure concepts are acquired strongly depends on the professional and personal background of the individual user. Since heart diseases are not confined to a given socio-economic background but rather can affect everybody, the group of heart patients is very heterogeneous with regard to computer and ICT affinity.

In order to create usable solutions it is necessary to involve users in early stages and during the process of development. UCD adds to this by providing different methods applicable at different stages in the process of development. Examples include contextual and behavioural analysis (in terms of interviews, site-visits, etc) in order to gather the users requirements, needs and habits and to know their environment; User-Tests in order to prove that the solution fits the users needs; etc.

The information technology provided has to be as easy to use as possible. Patients should be able use the system without extra mental and physical overloads - patients have to focus on their physical training.

### **1.3 Technical Requirements of Telemedicine Applications**

Telemedicine applications make high demands on the IT infrastructure they are based on. They have to be failsafe and they must offer a very fast data transmission to allow medical experts to immediately react to abnormalities. Another very important issue is data privacy: Transmitted and stored data has to be cryptographically secured in order to restrain unauthorized people from monitoring this data. Furthermore telemedical applications should offer interfaces that support well-known and standardized data structures and protocols to allow the integration into an existing hospital infrastructure. In the research project OSAmI-D (see section 3) all these issues are addressed by using a flexible and modular architecture based on OSGi and hence on Java. Java has the advantage of being executable on many platforms including mobile devices or set top boxes in a user's home. These are ideal basic conditions for a rehabilitation application that is aimed to support a patient within her everyday life.

Due to legal restrictions medical products – including algorithms – have to be certified before patients may use them. This fact makes it complicated to apply UCD

because real patients cannot test an application in the way they would use it if it was already certified. Therefore early paper or screen mockups may be used together with oral questionnaires to overcome this drawback.

## 2 Related Work

There are no remote telemedical systems for a cardiac exercise rehabilitation regime on the market available yet. Some research projects have addressed the issue. The SAPHIRE-Project [12] has also been situated in the same area as the OSAmI-D project. Here a patient would also have the opportunity to exercise with a bicycle ergometer training, but could only exercise on certain dates, while there is a supervisor controlling the data. SAPHIRE has considered neither an offline training nor a mobile training tutor. Further on no user centered design was applied during the project. Separate from SAPHIRE and OSAmI-D there is a EU funded project called HeartCycle which also address cardiac rehabilitation setting on a bicycle ergometer, but there has not been a demonstrator released yet.

On the prevention market there are several partly highly sophisticated training tutors. Starting with the more advanced POLAR systems which store heart rate and provides training diaries, the currently most advanced system is the concept of the company T2BEAM Technologies AG, Switzerland. It has emerged on the market in January 2009. The product called athlosoft (<http://athlosoft.com>) addresses healthy subjects, who can get online training supervision. The data are captured with a 1 lead ECG and transferred via a GPS enabled mobile to a web application where a trainer can supervise the ongoing training. Compared to OSAmI-D it does not allow the definition of medical constraints controlling a training session as well as a supervising platform does not exist.

Gay and Leijdekkers presented an approach on how sensors for measuring vital data may be used in mobile setups [13] Thereby the ECG data is analyzed automatically and locally in order to provide hints to the user or to alert pre assigned caregivers whenever necessary. Anyway here the data is not transmitted during or after an exercise session but monitored permanently by algorithms. Another project focusing on mobile applications especially for elderly people is presented by Oppermann et al. [14]. Here the application that's also monitoring vital data may also only be controlled with the mobile device itself resulting in small and very simple graphical user interfaces. We can benefit from these results when designing mobile training applications.

However there's no system existing using the different technical approaches of live supervision on an ergometer, offline training with predefined medical constraints and mobile training altogether for patients with cardiac diseases. We think it is necessary to achieve more insights of the user's needs, preferences and habits in this specific use case. Therefore UCD is the best possibility towards developing a successful and usable system.

### 3 Background: OSAmI-D

The European research project OSAmI and the German sub-project OSAmI-D are funded by the German Ministry of Education and Research (BMBF) et al. In the project open services for ambient intelligence based on Java and OSGi are going to be developed. OSAmI-D addresses telemedicine and remote ergometer rehabilitation in particular as an example of use. Since ergometer training is a well-known and effective rehabilitation appliance for patients with cardiac diseases, we think that this scenario has the potential to be realized as a commercial solution one day. The idea of the scenario is that patients who successfully finished rehabilitation phase II in a hospital should be able to continue their training afterwards in order to keep up the rehabilitation process and avoid a new cardiac event. Therefore a patient after rehabilitation phase II may herself install an ergometer at home. This ergometer comes along with several sensors for measuring vital data, an internet connection and a touchscreen in front of the exercising user. During an exercise training data from the sensors is collected, automatically analyzed and transmitted to a supervisor system. Supervisors may adapt the training plan of every patient according to abnormalities recorded in the current and/or previous session. Figure 1 shows a schema of the overall setup as far as the hardware components at the patient's home. The figure also shows the vital data sensors that are connected wireless via Bluetooth to the ergometer: An ECG sensor, a blood pressure sensor and an oxygen saturation (SpO<sub>2</sub>) sensor. Technically, the setup consists of a tablet PC that serves as user interface, internet gateway, collector of sensor data and ergometer controller. The ergometer itself offers access to adjustable values via a network interface. E.g. it allows a changing of the current load in Watts by applying the correct pedaling resistance according to the pedaling frequency of the user.

In the project we differ between three training levels. In *Level 1* there's a one-to-one live supervision between the patient and a supervisor. The training data - including vital data measured by the sensors and ergometer data - as far as audio and video recordings of the patient are directly transmitted to the supervisor. The supervisor analyzes the data online and can change the training schedule immediately or issue instructions to the patient. *Level 2* training is quite similar but here a larger group of patients is observed by one supervisor simultaneously. Audio and video transmission




Fig. 1. The OSAmI-D Setup

is optional and can be manually enabled or disabled by patients or the supervisor. Finally *Level 3* enables offline training for lower- risk patients. These can exercise at any time they like. The data is recorded but not transmitted before the training is finished. A supervisor analyzes the recorded data in post processing within a given time period, e.g. 24 hours, and adapts the training schedule if necessary. The schedule is transmitted to the patient's gateway PC afterwards. The *Level 3* scenario also enables a mobile training. Patients can use simplified versions of the vital data sensors to collect data on their mobile phone or PDA via Bluetooth while exercising outdoors. This provides more flexibility to patients who can perform their preferred sport on an arbitrary place. Even though there's no live supervision the system automatically compares the vital data with threshold values predefined in the training schedule by a supervisor. If some data exceeds a limit the patient is informed about that and she's given a hint of how to react (e.g. "Walk a bit slower"). Hence the responsibility in terms of effectiveness and security lies on the patient's side solely.

## 4 A User Study

This section introduces the user study held in a rehabilitation clinic in Germany in summer 2009. The goal of the study was to analyze the current situation of the patients: Their daily routines in the hospital, their personal ergometer rehabilitation program, their private habits and their skills in the IT area. These insights helped us to design the parts of the OSAmI-D system that are visible to the user within a UCD process. In a previews session also four ergometer training supervisors were interviewed [15] which gave us interesting insights and already helped us to find functional and non functional requirements for the overall system architecture and the user interfaces for the supervisors.

### 4.1 Interview Lead through

We interviewed six patients in the rehabilitation clinic, recruited by the clinic staff. All patients have been described as potentially adequate for remote ergometer training supervision. Five of the six patients were around retirement age (65 years in Germany) one was a young adult (approximately 30 years old). All stayed in the rehabilitation clinic with the objective of finishing a three week rehabilitation program including in situ supervised ergometer training three to five times a week. While five patients did already reach their final rehabilitation week the last one was still in the starting week of the training. All of the interviewed persons live really active social lives and are also physically active – most of them stated that they like cycling and are undertaking cycling tours quite often. Even though some of them are beyond retirement age they all seem to be very busy and pointed out to be pinched for time. The patients all have basic computer knowledge but aren't computer experts. This seems to be one reason the clinic staff recruited them as potentially adequate for remote ergometer training supervision. Supervisors said that patients without any computer knowledge wouldn't ever trust in such a system. Therefore they won't ever use it and are from the clinic point of view no potential patients for remote ergometer training supervision. It is an interesting question whether it would be possible to also address users without any IT experience by designing a user interface that is easily

understandable and intuitively operable. However this user group was not the target of the study presented here.

We used an interview-guideline with open questions and performed semi-structured interviews [16]. After asking personal questions and questions about the current in situ training (see next section) we introduced the different training levels and scenario ideas of OSAmI-D to the patients. In doing so we recorded all answers and comments as further user input.

## 4.2 The Current Setting

This section explains the current flow of the ergometer training for patients with cardiac diseases within a rehabilitation clinic. We carefully figured out what's important for them and for the supervisors here and how the several training steps are executed in order to develop the OSAmI-D system according to these workflows.

The in situ training in the early rehabilitation phase is organized in groups of up to 15 patients that are all exercising simultaneously. Therefore they all start a warm-up phase simultaneously, reach the training phase simultaneously and reach and finish the cool down phase of the training simultaneously. The focus of the supervisor here lies on visual control and personal conversations. She predominantly checks the patient's conditions by talking to them in order to detect critical situations. The vital data that is displayed on two large PC displays in a secondary room is not that important and only considered secondarily. Beside the proven medical effectiveness the biggest advantage of an in situ ergometer training is therefore the personal supervision through qualified medical staff. Further on many patients like the training within an exercise group because it enables social interactions with other patients.

However there are some disadvantages of an in situ ergometer training. It is described as very monotonous by the patients because it provides little variety. Comparing this exercise with an outdoor bicycle training it is quite stultifying, e.g. the fresh air, the air flow, the landscape and the sound of nature are missing. During in situ training the participants are looking towards the walls of the exercise room and just listening to quiet music. Further on patients stated that the ergometer training is not challenging but physically very easy. Since patients trust in the supervisors and the training settings they accept this issue but experience the ergometer training in a very passive way and regular feedback from supervisors is missing. Furthermore patients don't even know their personal training setting and schedule in detail. The challenge here is besides designing intuitive user interfaces also to raise new functional requirements towards the underlying system in order to implement new functionalities that motivate patients to exercise periodically and to overcome the monotony.

Therefore we collected the requirements and analyzed the context of use of the OSAmI-D system on the patient's side as a base for the user interface design. Since the training will not take place in situ anymore but at the patient's home we have to think about numerous changes in the training flow and in the interaction modalities between user and computer.

### 4.3 Results

This section introduces the main results of the semi structured interviews by abstracting the main declarations of the patients and analyzing the system requirements from the user's perspective.

Patients appreciate the possibility of performing remote ergometer training supervision as an expanded rehabilitation in general. At least those with basic computer knowledge do not have doubts regarding technology or organization. Since all of the interviewed persons have been physically active before they also do not have an issue with exercising regularly in the future. Nevertheless all patients stated that a long-term indoor ergometer training is not attractive in comparison to alternative sports especially outdoor ones. We think that therefore it is desirable to develop even mobile solutions.

Patients with cardiac diseases have the need for supervision, distraction and confirmation when exercising. Especially they would like to get their vital data visualized ("It would be nice if I knew my own values!"). Knowing these values would enable them to have food for thought and to know that everything is copacetic. Automatic feedback (current sensor values) is therefore very important.

Patients see many advantages in the planned setup: They don't need to arrange fix training dates (at least not when applying *Level 2* or *Level 3* training) and they save time in comparison to an in situ training in the clinic because they would have to travel there. Further on they can define the entertainment program on their own: Reading news, listening to music, watch TV or surf the web. Therefore an entertainment system as part of the patient user interface is desirable. However we don't think that this will overcome the monotonous training in the long term.

## 5 User Interface Design

One of the most definite results we figured out in the study described in the previous section is that patients need to get their personal vital data visualized. This helps to overcome the monotony of in door ergometer training and to compare the results concerning training success and personal threshold values. Therefore the graphical user interface of the OSAmI-D project demonstrator should present the most important data in a clearly arranged way. The interface should further on provide a mechanism to show a calendar containing personal feedback messages from the supervisor. Regular feedback can help to motivate patients to exercise more frequently.

We found out that the user interface of the precursor project had some lack in consistency especially regarding the colors of the graphical elements. Patients stated that they don't understand the applied color-coding of yellow and red buttons. Further on some buttons aren't even needed – e.g. patients would never push a button when they decide to stop the training because of illness but just stop cycling and make further measures. Another inconsistency in the coding was recognized at the also implemented traffic light. Sometimes the yellow color was used to inform the patient about a high heart frequency - an alarm - sometimes it was used to inform the user to continue the training. Hence it was often unclear how dramatic the current alarm or information really is. In OSAmI-D we use traffic light colors to indicate that



everything's fine (green color), there's a warning but the patient may continue the training (yellow) and there's a critical situation and the training will be stopped (red color).

Another issue that was evaluated as distracting was the length of the questionnaires patients have to answer before and after training in *Level 2* and *Level 3*. Many questions were unclear, others may be grouped, others are completely unnecessary (e.g. "Are you feeling good" – A patient that not feels good wouldn't have get on the ergometer to start training). A last thing is that we had to take care of the size of graphical elements and especially fonts when designing the user interface. Patients are often elderly people with visual impairments and they are exercising when interacting with the interface. When considering a mobile training it is therefore especially important to only display the information that is really needed. Designing the interface we also recognized that some button labels are understandable to computer scientists but not to elderly people with less IT knowledge and to persons who don't understand English (the interface language in the study was German but accidentally some English abbreviations were used for button labeling).

## 5.1 Mockups

The implications for remote supervision were used to develop user interface mockups. Here we introduce the most important mockup figures, the overall layout and the idea behind it.

The main menu (See Figure 2) appears after a patient logged in using her username and password. Even though it is not that easy for patients to remember both we are using this login mechanism due to security reasons. The main menu shows the latest supervisor feedback and offers the possibility to start a supervised training either in online mode (*Level 1* or *Level 2*) mode or in offline mode (*Level 3*). All is done via simple single button clicks using a single finger. Since the user interface elements are big (screen size is 17") this is easy for nearly every patient. Everything is clearly arranged and the interface only offers the most important functions.





Fig. 2. User Interface Mockup showing the Main Menu



**Fig. 3.** User Interface Mockup that is presented to a patient while she is exercising

The main menu further on offers possibilities to start the synchronization with a mobile device for performing mobile offline training, to shut down the system and to open the calendar. The calendar shows all past and future training sessions as far as information about changes in the training plan and supervised or not yet supervised sessions in the past.

In the interviews we found out, that the asked people need feedback about what exactly has been pressed on the screen. In most cases it is definite because of the following things that are displayed but not in case of questionnaires. Even so these are essential for a supervisor when analyzing the training in post processing. Therefore we decided to first highlight a selection and then force the user to confirm it by pressing another button. This was well accepted by all test users even if they had to click the screen twice.

Figure 3 shows the user interface while the patient is exercising. It can be seen that we decided to use a traffic light metaphor to display the overall situation, e.g. a green light means that everything is all right and the training should be continued. The yellow light means: Continue the training but additional hints are given (e.g. “Slow down”). When the red light appears the training will be stopped after cooling down. Traffic lights are easy to understand, even by color-blind people. Further on we decided to use easy symbols for the vital data and the ergometer data to display the heart frequency, blood pressure value, current pedaling frequency, load, speed and time. We think that these icons are easy to understand but we have to verify this in the next iteration of the user centered design process.

## 6 Conclusions

We found out interesting insights about the needs and habits of patients with cardiac diseases if they would decide to participate in a remote exercise supervision program. This helped us to design user interface mockups and a first software prototype that provide easy interaction, clarity and overall usability. Patients need flexibility in organizing their free time. E- Health systems should allow them to exercise whenever

they want and wherever they like. We figured out several insights about this kind of medical assistance, the problems that have to be solved and the market potential. Ergometer training for patients with cardiac diseases competes to more attractive physical activities like outdoor sports. It is therefore extremely important to generate and keep up a high motivation of patients.

Regarding the user interface design, the vital data of patients should not longer be hidden but presented to them during training. This enables patients to compare values with previous sessions and allows them to keep an eye on the overall training success. Further on displaying regular feedback from supervisors helps to motivate patients and gives them a feeling of safety. This is a unique feature of a remote supervision system. Patient user interfaces should guide the users fast and intuitively. Patients don't like an overhead of time or frustrating questions that occur every time they start the training.

Even though the user study described above was done in order to acquire information about the user interface design for this very special use case, it also raised functional requirements towards the underlying software. First we found out that videoconferencing is essential in *Level 1* and *Level 2* training since the main information about a patients well being can be taken from his body language (skin color, position, pedaling regularity etc.). Even though in OSAmI-D it was planned to implement videoconferencing at the very beginning because of its technical feasibility we found out that it is very desirable for both user groups – patients and supervisors.

Automatic analysis of the vital data values is another requirement. Some kind of data may be analyzed quite easy (e.g. if the blood pressure value is too high), some analysis is really hard to implement (e.g. detect abnormalities in the ECG curve). Here third party solutions have to be obtained and implemented in the future.

Another functional requirement is the entertainment system. It may help to motivate the patients to exercise periodically and it may help to overcome the monotony of indoor ergometer training. Here also a mobile solution is desirable because it enables patients to exercise outdoor and she's not longer limited to cycling. Therefore the system has to offer possibilities to synchronize mobile phones with the home gateway (in our current setup this is the tablet PC on the ergometer) in order to exchange training schedules and reports.

Patients prefer to have a flexible possibility of making training appointments. Here the calendar functionality is a nice feature that enables an overview of training dates, the different training levels that are available and past training events.

Since the hardware of the ergometer itself (saddle position and material, overall weight) is not changeable (we're using an existing ergometer setup in the project) we could not consider comments regarding these values.

## References

1. Medica: Megatrend Ambient Assisted Living (2008), <http://www.heise.de/newsticker/Medica-Megatrend-Ambient-Assisted-Living-meldung/119306>
2. Stroetmann, K.A., Jones, T., Dobrev, A., Stroetmann, V.N.: An evaluation of the economic impact of ten European e-health applications. *Journal of Telemedicine and Telecare* (2007)

3. Ambient Assisted Living: IT- Technik fürs Alter (2008),  
<http://www.heise.de/newsticker/Ambient-Assisted-Living-IT-Technikfuers-Alter-meldung/137183>
4. EUROASPIRE I and II Group, et al.: Clinical reality of coronary prevention guidelines: a comparison of EUROASPIRE I and II in nine countries. *Lancet* (2001)
5. EUROASPIRE II Study Group: Lifestyle and risk factor management and use of drug therapies in coronary patients from 15 countries; principal results from EUROASPIRE II Euro Heart Survey Programme. *European Heart Journal* (2001)
6. Bjarnason-Wehrens, B., Held, K., Karoff, M.: Herzgruppen in Deutschland – status quo und perspektiven. [Heart groups in Germany-current situation and prospects], *Herz* (2006)
7. Jokela, T.: An Assessment Approach for User-Centred Design Processes. In: *Proceedings of EuroSPI 2001* (2001)
8. Bevan, N.: Quality in Use: Meeting User Needs for Quality. *Journal of System and Software* 49(1) (1999)
9. Granollers, T., Lorès, J., Perdrix, F.: Usability Engineering Process Model. Integration with Software Engineering. In: *Proceedings of HCI International* (2002)
10. Juristo, N., Windl, H., Constantaine, L.: Special Issue on Usability Engineering in Software Development. *IEEE Software* 18(1) (2001)
11. Nebe, K.: Integration von Usability Engineering und Software Engineering: Konformitäts- und Rahmenanforderungen zur Bewertung und Definition von Softwareentwicklungsprozessen, Doctoral Thesis, Shaker Verlag, Aachen (2009)
12. Busch, C., Baumbach, C., Willemsen, D., Nee, O., Gorath, T., Hein, A., Scheffold, T.: Supervised training with wireless monitoring of ECG, blood pressure and oxygen-saturation in cardiac patients. *Journal of Telemedicine and Telecare* 15 (2009)
13. Gay, V., Leijdekkers, P.: A Health Monitoring System Using Smart Phones and Wearable Sensors. *International Journal of ARM* 8 (2007)
14. Oppermann, R., Lorenz, R.: Seniorengerechte Gestaltung eines Dienstes zur mobilen Überwachung von Gesundheitsparametern, Bundesministerium für Bildung und Forschung -BMBF-: Ambient AssistedLiving. 1. Deutscher Kongress mit Ausstellung (2008)
15. Klompmaker, F., Nebe, K., Busch, C., Willemsen, D.: Designing Context Aware User Interfaces for Online Exercise Training Supervision. In: *Proceedings of the 2nd International Conference on Human System Interaction* (2009)
16. Wengraf, T.: *Qualitative Research Interviewing: Biographic Narrative and Semi-structured Method*. Sage, London (2001)

# A Value and Model Driven Method for Patient Oriented KPI Design in Health Care

Erik Perjons and Paul Johannesson

Department of Computer and Systems Sciences, Stockholm University  
Forum 100, SE-164 40, Kista, Stockholm, Sweden  
{perjons,pajo}@dsv.su.se

**Abstract.** In all societies, large resources are spent on health care. However, service quality is still unsatisfactory in large parts of the health care sector. One reason for this state of affairs is the complexity of health care organisations, which make them difficult to govern. Furthermore, due to the high degree of specialisation, health care organisations often need to cooperate, which adds to the complexity. One possible way to improve the cooperation and management of health care organisations is the use of performance management, which is a management approach aiming at optimising the performance of an organisation. The main instrument in performance management is the use of Key Performance Indicators (KPI), which measure how well an organisation fulfills its goals. Designing effective KPIs is a complex and time consuming task that requires substantial efforts. Therefore, there is a need for methods that support organisations in designing KPIs. This paper proposes a value and model driven method for identifying patient oriented KPIs.

**Keywords:** Performance management, KPI, Enterprise modelling, Service quality, SERVQUAL.

## 1 Introduction

Large resources are spent on health care in the European welfare society, but a number of problems still remain, including unequal access to health care, large variations in outcomes of treatments, deficiencies in service quality, and inefficient resource use. A main reason behind these problems is the complexity of the health care sector, which makes it problematic to govern. For example, health care needs to fulfill several and often contradictory goals, such as equality, high quality and efficiency. At the same time, a large number of stakeholders need to interact with each other in order to ensure the delivery of high-quality health care. Furthermore, European health care faces a period of potentially profound changes in social attitudes, demographic structure, economic conditions and medical technologies.

In order to better govern the health care sector, there is an increasing interest in performance management, which is a management approach aiming at optimising the performance of an organisation, i.e. optimising its business processes and outcomes [1]. The main activities in performance management are identifying

business goals, developing key performance indicators (KPIs), monitoring and measuring the performance using the KPIs, analysing the results of the measurements, and acting in order to better fulfill the goals of the organisation [2]. Thus, performance management can be an approach for managing efficiency, cost, and quality in any business, including health care.

The main instrument in performance management is the KPI, which we here define as a *property of a phenomenon that can be used to measure the performance of an organisation*. Examples of KPIs are “sales revenue”, “patient satisfaction”, and “patient throughput”.

KPIs operationalise business goals, i.e. vague and high-level goals can be expressed in terms of KPIs, thereby making them concrete and easy to understand [2]. For example, a goal like “Patients should get fast treatment for serious diseases” can be operationalised as “Waiting time for hip replacement surgery should be less than one month during 2010”, which contain a KPI (“Waiting time for hip replacement surgery”) and a target (“less than one month during 2010”). Such a statement can easily be monitored and measured.

KPIs are powerful instruments for governing an organisation, as they are easily understandable, actionable, and can be effectively monitored. However, this also means that the use of inappropriate KPIs can have significant negative effects on the performance of an organisation [2]. For example, using KPIs focusing on quantity instead of quality can divert employees from paying attention to quality in products and services. Thus, it is essential to design and implement effective KPIs that actually improve the performance of an organisation in different aspects, such as efficiency, cost and quality.

Designing effective KPIs is a complex and time consuming task, which requires substantial effort from performance management experts as well as domain experts, such as health care managers, physicians and nurses [3]. In particular, it is difficult to ensure completeness in KPI design, i.e. when a set of potential KPIs have been identified, there is still a risk that relevant and effective KPIs have been omitted. Another issue in KPI design for health care is the fact that KPIs aimed at measuring service quality in many cases cannot be directly understood by patients because of the complex relationship between diagnostic and therapeutic services and their results [3]. This is different from most other industries, where service quality almost always can be directly perceived by service customers.

The goal of the paper is to propose a method for constructing a complete set of patient oriented KPIs, i.e. KPIs from the patient’s perspective. KPIs measuring internal efficiency, such as revenue and cost efficiency, are not addressed in this paper.

The rest of the paper is structured as follows. Section 2 gives an overview of related work, in particular performance management and service quality analysis. Section 3 presents the method that addresses both technical and functional service quality in health care. Section 4 provides a demonstration of the method based on a case study from the eye care health sector. Finally, Section 5 summarises the results of the paper and gives suggestions for future work.

## 2 Related Work

Performance management, i.e. measuring performance of an organisation, has been an important management instrument in business for the past 20 years [4]. In recent years the use of performance management has accelerated, as evidenced by the proliferation of performance management methods, key performance indicators (KPI), and IT support such as business intelligence and performance dashboard systems. Health care is no exception to this trend; performance management has attracted substantial attention among governmental and private funders of health care, health care providers, patient organisations as well as researchers in health care management [4]. This interest has emerged in parallel with the interest for other, and closely related, health care improvement approaches, such as accreditation, service evaluation, quality improvement, external auditing, outcome research and evidence-based medicine [4]. However, there is no clear demarcation line between these approaches and performance management, as different practitioners and researchers use different definitions and concepts.

Measuring performance is about measuring the outcomes and quality of business processes. In health care, there are two main categories of quality: technical service quality, i.e. clinical results of health services, and functional service quality, i.e. patient assessments of the quality of care [3], see also [5]. In general, technical service quality is difficult to understand and measure for both health care providers, health care funders and patients [3]. Therefore, many performance management initiatives in health care have focused on functional service quality.

A main research theme within functional service quality has been to find the right balance of service quality dimensions. As a starting point, a well-known service quality framework, SERVQUAL [6], has frequently been used. SERVQUAL measures the gap between perceived and expected service quality, using five quality dimensions: reliability, assurance, tangibles, empathy and responsiveness. Often used is also a predecessor to SERVQUAL [7], which includes ten quality dimensions. In SERVQUAL these ten dimensions were condensed to five, each with a set of sub dimensions, called items. SERVQUAL was meant to be a service quality framework independent of business area. However, when used in health care, further dimensions are often added, commonly created using focus groups with physicians, nurses, health care managers and patient representatives, see for example [8], [9], and [10]. The relationship between service quality in health care and patient satisfaction is also a complex issue, discussed in several research papers, see, for example [11].

For identifying KPIs, the most common tool is the use of goals and objectives [2]. Usually the goals and objectives of an organisation are structured graphically in goal models, visualising an hierarchy of strategic, tactical and operational goals.

Goal models are examples of enterprise models, while other examples are value models, process models, conceptual models and information models. In general, enterprise models offer graphical representations of the structure, goals, processes, information, resources, people, and constraints of an organisation. They provide compact and graphical descriptions of an organisation and its environment, which makes them ideal for supporting communication between different stakeholders. They can work as a base for creating a common understanding of the organisation, for change management, as well as for IT system design. Process and conceptual models have

been used extensively both in health management research and health care standards and specifications, such as SAMBA, HISA, HL7, CONTSYS. Recently, using value models for analysis in health care have gained increased attention [12], [13]. In this paper, value models and conceptual models are used as the basis for identifying KPIs. To our knowledge, this has not been done in previous work.

### 3 Method for Value Based KPI Identification

The purpose of the method proposed is to assist a KPI designer in creating a set of KPIs that can be used to assess the quality of health care services from a patient point of view. An overview of the method is depicted in Fig. 1.




Fig. 1. A graphical overview of the method for value and model based KPI identification

The first task to be carried out is to create a value model scoping the health care scenario under consideration, thereby determining the most important health care services as well as other services in the scenario. The second task aims at generating KPIs for technical service quality that concern the results of health services and the consequences of health issues. For this task, we introduce a supporting instrument for representing relevant health care services and health issues. This instrument takes the form of a so called standard service and issue model (standard SIM), see Section 3.2, that can be modified and extended to fit the current health care scenario. The third task aims at generating KPIs for functional service quality that concern the manner in which services are delivered to patients. For this purpose, we introduce a supporting



instrument in the form of a so called standard service dimension model (standard SDM), see Section 3.3, based on the predecessor to SERVQUAL [7], which can be modified and extended to fit the current health care scenario. Finally, there is a reviewing and filtering task, i.e. task four, where the KPIs suggested in the previous tasks are evaluated and those deemed most important are selected as candidates for implementation. Below the tasks are described in detail.

### 3.1 Task 1: Develop Value Model

A first task is to delimit the domain for which the KPIs are to be designed. For this purpose, a value model is to be created. The value model shall describe the actors of the value network under consideration, such as health care providers and patient. The value model shall also show the resource transfers between these actors, including transfers of services such as examinations and treatments. The value model will be the basis for determining health care services and other patient related services for which KPIs are to be developed. A simple example of a value model is shown in Fig. 2, which shows two actors, patient and primary health care provider, as well as their resource transfers, eye examination and payment.




Fig. 2. A simple value model

### 3.2 Task 2: Generate Technical KPIs

Technical service quality focuses on technical accuracy and procedures, in particular the quality and effectiveness of the diagnostic and therapeutic intervention processes. In other words, technical service quality is about the results of health services and the consequences of health issues. A number of top level goals for technical service quality are:

1. Health care services should have positive effects on health issues
2. Secondary diseases should be avoided
3. Health issues should be treated by evidence based health care services
4. Health care services should not give rise to unwanted side effects

In order to measure the fulfillment of these goals, we introduce four generic KPI templates, each addressing one of the goals. When instantiated, these templates will result in a set of KPIs.

The generic KPI templates are:

1. Percentage of health care service  $x_1$  that is followed by health care service  $x_2$
2. Percentage of health issue  $y_1$  that is followed by health issue  $y_2$
3. Percentage of health issue  $y_1$  that is followed by health care service  $x_1$
4. Percentage of health care service  $x_1$  that is followed by health issue  $y_1$

In order to instantiate these templates, relevant health care services and health issues need to be chosen. As this choice depends on the health care scenario under consideration, KPI designers need an instrument for documenting and representing the health care services and health issues they decide to include. When they have done so, they can go on to instantiating the KPI templates to arrive at KPIs. We suggest that relevant health care services and health issues should be represented by a simple conceptual model, called a service and issue model (SIM). Such a model consists of two classes, Health Care Service and Health Issue as well as a number of subclasses of these. The subclasses included in a SIM will depend on the health care scenario being addressed and will, therefore, vary from scenario to scenario.




Fig. 3. The standard SIM

In Fig. 3, we suggest a standard SIM, based on notions from CONTSYS [14], that should be applicable in many scenarios with no or minor modifications.

The method for designing KPIs for technical service quality can now be formulated in three steps:

*Step 1: Adapt Standard SIM*

Construct a SIM. This can be done by modifying the standard SIM of Fig. 3. Fig. 4 shows an example of such a SIM, adapted from the standard SIM.




Fig. 4. A SIM, adapted from the standard SIM described in Fig. 3

### Step 2: Apply Generic KPI Templates

For each of the four generic KPI templates (presented above) and every subclass X of Health Care Service in SIM (see Fig. 4) and every subclass Y of Health Issue in SIM (see Fig. 4), introduce a new specific KPI template by replacing every  $x_i$  (in the generic KPI templates) with X and every  $y_i$  (in the generic KPI templates) with Y. Each such template will have the form “Percentage of X that is followed by Y”, where X and Y are classes in SIM.

An example of a specific KPI template, based on KPI template 1 and the SIM in Fig. 4, is:

- Percentage of primary prevention that is followed by secondary prevention

An example of a specific KPI template, based in KPI template 3 and the SIM in Fig. 4, is:

- Percentage of diagnosis that is followed by a medicine treatment

The created specific KPI templates can also by themselves be used as KPIs.

### Step 3: Apply specific KPI templates

For every specific KPI template from Step 2, “Percentage of X that is followed by Y”, replace X and Y with instances of these classes. This replacing can be carried out in two ways:

- a) A domain expert examines every specific KPI template and replaces X and Y with instances based on her knowledge.
- b) A domain expert populates the SIM with instances. A list of all possible KPIs is then generated automatically. Finally the domain expert selects relevant KPIs from the list.

Two examples of KPIs, created from the specific KPI template “Percentage of diagnosis that is followed by medicine treatment”, are:

- Percentage of *flu diagnosis* that is followed by *Tamiflu treatment*
- Percentage of *stroke diagnosis* that is followed by *warfarin treatment*

After having applied these steps, the KPI designer will have arrived at a number of KPIs that focus on the results and consequences of health care services and health care issues. These KPIs are objective in the sense that they do not primarily depend on the assessments of patients.

## 3.3 Task 3: Generate Functional KPIs

While technical service quality focuses on the effects of health care services, functional service quality refers to the manner in which services are delivered to the patient. Functional service quality includes aspects like facilities, hospital food, employee attitudes, responsiveness, and cleanliness. While technical service quality is typically difficult to judge for the individual patient, functional service quality is usually directly visible to the patient. Therefore, functional service quality often has more impact on a patient’s service quality perception than technical service quality.

In contrast to technical quality, functional quality does not depend only on health care services but also on complementary services, like ordering, information and complaint services. We have identified a number of complementary service types, based on an adaptation of the open-EDI phases [15]. Note that one of the open-EDI phases,

i.e. actualisation, has been excluded as it corresponds to the original service, i.e. the service for which the complementary services are identified.

The complementary services identified are the following:

#### *Identification*

An identification service offered by a health care provider is an information service that provides information about the provider's services and helps a patient in identifying and selecting among health care providers and their services.

#### *Negotiation*

A negotiation service offered by a health care provider is an interactive service where the patient and the health care provider negotiate in order to arrive at an agreement on a future health care service including its cost, scheduling, location and personnel.

#### *Pre-actualisation*

A pre-actualisation service offered by a health care provider is an information service where the health care provider informs the patient on adequate preparations for a health care service.

#### *Post-actualisation*

A post-actualisation service offered by a health care provider is either

- a) an information service where the health care provider informs the patient on the result of a health care service
- b) an information service where the health care provider informs the patient on adequate behaviour to be carried out or observed by the patient after the performance of a health care service
- c) a service where the health care provider accepts and addresses complaints on a health care service

Both health care services and their complementary services should be assessed along a number of service dimensions. Thus, a KPI designer is to select a set of service dimensions and for each service dimension, one or several service items. Which service dimensions and items to investigate depend on the health care scenario under consideration. We envisage that in most cases, the service dimensions will be variations or extensions of those in SERVQUAL or its predecessor. Furthermore, for each service item and each type of service (including complementary service types), the KPI designer shall specify how significant the service item is for that type of service. A service item is significant for a service type if it is important for a patient's perception of services of that type and if it is difficult for the health care provider to score well on the item. Only if both these conditions are satisfied, it becomes interesting to measure the item through KPIs.

Summarising, the KPI designer shall define a service dimension model (SDM), consisting of a set of service dimensions; for each service dimension, a number of service items; and for each combination of service item and complementary service type, the significance of the service item for the service type. A SDM can be represented as a matrix, see Table 1.

We do not believe there exists a canonical SDM that is optimal for each health care scenario. However, we envisage that most SDMs will share large parts. Therefore, it is worthwhile to introduce a standard SDM that can serve as a starting point when a KPI

designer develops her own model for a particular scenario. For this purpose, we suggest the standard SDM in Table 1, based on the predecessor to SERVQUAL, where we have indicated significance on a three level scale: zero, one or two + marks. The significance marks of the model are justified as follows:

*Reliability* is less important for identification and negotiation services partially because no promises have been made before these services, thereby making the question of faithful service execution irrelevant. Consistency of performance and prompt attention to defect service are particularly important for actualisation, as these directly influence the health state of patients.

*Responsiveness* is important in all phases, as customers value prompt services and willingness during their entire episodes of care.

*Competence* is particularly important for actualisation, as it directly influences the health state of patients. It is also important for the pre- and post-actualisation phases, where it also has a strong impact on health state. Competence is less important for the identification and negotiation phases, as the competence needed for these activities is relatively easy to achieve.

*Access* is particularly important for identification and negotiation services, as the use of these services potentially may be stressful and take a long time. Good access may reduce stressfulness as well as time consumption for patients.

*Courtesy* is important in all phases, as customers value a positive attitude and privacy during their entire episodes of care.

*Communication* is especially important in pre- and post-actualisation phases, as these involve much patient interaction and may have substantial impact on the health state as well as the feeling of safety of the patient.

*Understanding the customer* is important in all phases, as health issues are complex and dependent on the individual patient. Therefore, customised services are often needed in health care.

*Tangibles* are of interest when services are given at health care facilities by health care personnel, i.e. in the actualization phase, as well as in the pre- and post-actualisation phases. Visual appeal of physical facilities and cleanliness of employees will indicate that the health care provider offers high quality services. However, a well-structured web page is also of importance in the identification and negotiation phases.

The method for designing KPIs for functional service quality can now be formulated in three steps.

#### *Step 1: Identify Complementary Services*

For each service from the value model, introduce complementary services according to the complementary service types above.

#### *Step 2: Adapt the Standard SDM*

Construct a SDM. This can be done by modifying the standard SDM of Table 1.

#### *Step 3: Apply the SDM*

For each health care service and complementary service from Step 1, apply the SDM from Step 2, i.e. identify for each service item zero or more KPIs. The significance of a service item for a service type will assist in determining whether to introduce a KPI or not.

**Table 1.** The standard SDM

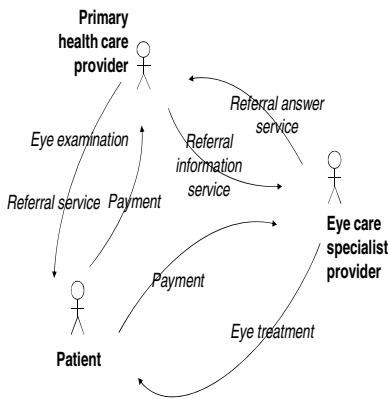
	Actua- lisation	Identi- fication	Nego- tiation	Preac- tualisa- tion	Post- actualisa- tion- result	Post- actualisa- tion- behavior	Post- actualisa- tion complains
<b>RELIABILITY</b>							
consistency of performance	++			+		+	+
correct keeping of records	+				+	+	+
correct billing	N/A	N/A	N/A	N/A	N/A	N/A	+
on-time (promised) delivery	+	N/A	N/A	+	+	+	+
prompt attention to defect services	++			+	+	+	+
<b>RESPONSIVENESS</b>							
prompt service delivery	+	+	+	+	+	+	+
exact information on service delivery time	+	N/A	N/A	+	+	+	+
employee willingness to help	+	+	+	+	+	+	+
employee willingness to help when busy	+	+	+	+	+	+	+
calling the customer back quickly	+	+	+	+	+	+	+
<b>COMPETENCE</b>							
knowledge and skill of the personnel	++			++	++	++	++
education level of the personnel	++			+	+	+	
experiences of the personnel	++			+	+	+	
research capability of the organisation	+	N/A	N/A				N/A
<b>ACCESS</b>							
easy access to services by phone, web and mail		++	++				
short waiting time to receive service	++	++	++	++	++	++	++
convenient hours of operation	+	++	+	+	+	+	+
convenient location of service facility	+						
<b>COURTESY</b>							
attitude of employee in performing the service	+	++	++	+	+	+	++
high level of privacy	++	++	++	N/A	++	N/A	
<b>COMMUNICATION</b>							
understandable explanation of the service			+	++	++	++	++
understandable explanation of cost and trade-offs between service and costs	N/A	N/A	++	N/A	N/A	N/A	++
patient interaction	+		+	+	++	+	+
continuous state updates	+	N/A	N/A	+	+	+	+
<b>UNDERSTANDING/ KNOWING THE CUSTOMER</b>							
learning the customer's specific requirements	++	+	++	++	++	++	+
providing individualised attention	++	+	++	++	++	++	+
<b>TANGIBLES</b>							
up-to-date equipment	+	N/A	N/A	N/A	N/A	N/A	N/A
visual appeal of physical facilities	+				+		
well dressed and neat appearance of employees	+				+		
consistency between the appearance of physical appearance and type of services provided	+				+		
cleanliness of employees	++				+		
cleanliness of physical facilities	++				+		
web page		++	+				

## 4 Method Application

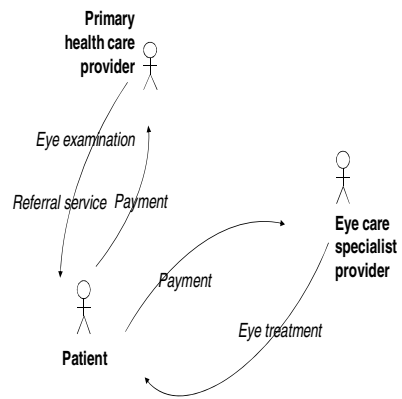
In this section, the proposed method is applied for demonstration purposes. The method is applied on the results from a research project (REMS) in the eye health care domain in Stockholm [16]. The main aim of the project was to develop and evaluate a set of e-services that could be used to create, manage and transfer health care referrals between primary health care and eye specialist providers. In the project, a set of enterprise models were constructed to support the design and evaluation of the e-services. Some of these models are used in this section.

### Task 1: Develop Value Model

The first task in the proposed method is to delimit the domain for which the KPIs are to be designed, i.e. specify which services need to be measured using KPIs. For that purpose, a value model is constructed. In this case, a part of a value model from the REMS project is used, see Fig. 5.



**Fig. 5.** A value model from the REMS case showing actors and resource transfers



**Fig. 6.** The value model from the perspective of the patient

The value model in Fig. 5 shows three actors - patient, primary health care provider and eye care specialist provider - and the transfers of resources between them. The background of the value model is the following: When a patient experiences an eye health problem, she/he will visit the primary health care provider. The basic resource this provider offers is an eye examination service. In order to receive that service, the patient needs to pay a certain fee for the service. The fee transfer is carried out via a payment service. If the patient needs further treatment, either the primary care provider will carry out the treatment (which is not shown in Fig. 5) or the provider refers the patient to an eye care specialist provider who is able to provide advanced treatment services. To do this, the primary health care provider transfers a referral via a referral service to the patient, which allows her to be treated by an eye care specialist provider. The primary care provider will also transfer referral information via a referral information service to the eye specialist provider, which will be used for scheduling a treatment service for the patient, as well as being the base for

resource allocation. The eye care specialist provider offers an eye treatment service to the patient and the patient needs to transfer a fee via a payment service. After the treatment service is carried out, the eye care specialist provider will transfer a referral answer service to the primary health care provider, informing it about the result of the treatment.

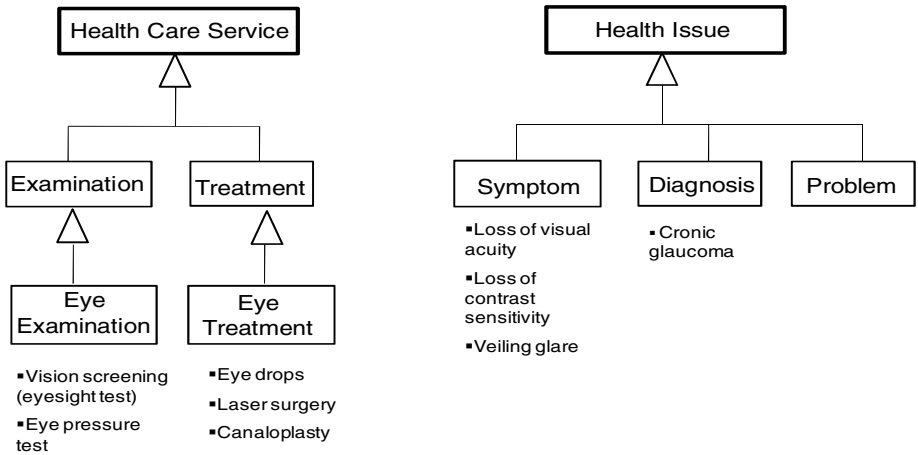
In the method proposed in this paper, only KPIs from the perspective of the patient are to be generated. Therefore, only resourced transferred to and from the patient are included. This will delimit the value model further, see Fig. 6.

**Task 2: Generate Technical KPIs**

In order to generate the KPI for measuring technical service quality, three steps are to be carried out.

*Step 1: Adapt Standard SIM*

In this step, the standard SIM, see Figure 3, is to be adapted to the value model constructed in task 1. In other words, the standard SIM has to include all the health care services from the value model, but not other types of services, such as administrative and resource allocation services. For example, the value model in Fig. 6 includes two health care services, eye examination and eye treatment, which are to be included in the SIM, see Fig. 7.



**Fig. 7.** A SIM with concepts (subclasses) based on the information from the value model. Examples of instances of the concepts are also shown.

*Step 2: Apply Generic KPI Templates*

In this step, the four generic KPI templates presented in section 3.2 are to be instantiated by the concepts (subclasses) of the SIM from step 1, in order to create a set of specific KPI templates. There are seven concepts in the adapted SIM in Fig.7. However, the more generalised concepts examination and treatment will not be used in this step but only their specialised concepts, eye examination and eye treatment. For example, the generic KPI template 4: “Percentage of health care service  $x_1$  that is followed by health issue  $y_1$ ” can be instantiated by the health care service *eye treatment*



and the health care issue *symptom* from the adapted SIM. This will result in a specific KPI template: “Percentage of health care service *eye treatment* that is followed by health issue *symptom*”.

### *Step 3: Apply Specific KPI Templates*

In this step, the resulting specific KPI templates from step 2 need to be instantiated by a set of instances from the SIM. This requires that the SIM in Fig. 7 be instantiated. Examples of instances of the SIM can also be found in Fig. 7. These 9 instances are applied on the specific KPI templates, resulting in a number of KPIs. For example, the template constructed in the previous step, “Percentage of health care service *treatment* that is followed by health issue *symptom*”, can be instantiated by the treatment *eye pressure test* and the symptom *veiling glare*. This will result in the KPI “Percentage of health care service *eye pressure test* that is followed by health issue *veiling glare*”.

### **Task 3: Generate Functional KPIs**

In order to generate KPIs for measuring functional service quality, three steps are to be carried out.

#### *Step 1: Identify Complementary Services*

In this step, the complementary services are applied on each service in the value model, both health care services and other services. For example, the service *eye examination* will have four complementary services: identifying an eye examination service, negotiating an eye examination service, pre-actualising as well as post-actualising an eye examination service.

#### *Step 2: Adapt the Standard SDM*

In this step, the standard SDM, see Table 1, is adapted. This is done by adding, deleting or changing the dimensions and items of the standard service dimension model. In this example case, we do not make any changes to the standard service dimension model but keep it as is.

#### *Step 3: Apply the adapted Service Dimension Model*

For each service identified in Step 1, we consider all the service dimensions and service items in the SDM from Step 2. For each item, we determine whether to include some KPI or not. For example, for the (complementary) service “pre-actualising an eye examination” and the service item “understandable explanation of the service”, we introduce a KPI based on the following question to a patient “Did you get an understandable explanation of how to prepare for your eye examination?”

## **5 Concluding Remarks**

In this paper, we have proposed a method for designing KPIs for a health care organisation with a focus on patient oriented KPIs. The starting point of the method is a value model that delimits the domain under consideration by identifying health care actors and resource transfers, including services. Based on the services identified, the method helps to generate KPIs for both technical and functional service quality. Technical quality KPIs measure the results of health services and the consequences of health issues and are generated by considering a small number of top level goals for

technical quality. Functional quality KPIs refer to the manner in which health care services are delivered to patients and are generated by identifying complementary services and applying dimensions from the predecessor to SERVQUAL. This way of generating KPIs can also be seen as goal driven, as the dimensions of the predecessor to SERVQUAL and its items express subgoals of high level goal customer satisfaction.

The proposed method provides four main advantages:

*Ease of Design.* Designing a set of KPIs for a health care scenario is a complex and time consuming task. The proposed method alleviates this task by offering a designer a large number of potential KPIs from which she can choose the most relevant ones.

*Completeness.* A problem in KPI design is to ensure that all relevant KPIs have been identified, i.e. to ensure the completeness of the KPIs. The proposed method addresses this problem by systematically generating relevant KPIs for technical as well as functional quality.

*Traceability.* KPIs should be possible to justify by relating them to the goals whose fulfillment they measure, i.e. KPIs should be traceable to goals. The proposed method provides traceability by directly generating potential KPIs from goals, which in the case of functional quality are expressed through SERVQUAL dimensions and items.

*Flexibility.* The health care scenario under consideration will influence the choice of KPIs, meaning that a KPI designer needs to be able to describe the specifics of a certain scenario in a convenient way. The proposed method supports this description through the use of value models, SIMs and SDMs, where the KPI designer can describe the most important health care services and health issues as well as relevant service items and their prioritisation.

In this paper, we have only addressed technical and functional service quality. However, a third aspect of service quality is the effect a health service can have on the quality of life of a patient, including the physical, psychological and social functioning of the patient. A direction of future work is to design guidelines for identifying KPIs for this service aspect. This requires measuring how well a patient functions in daily life during an extended period of time, e.g. how often the patient makes or receive phone calls to friends and family.

Another direction of future work is to provide support in actually formulating KPIs for functional quality. The method presented only suggests service items for different services that should be measured by KPIs but does not suggest formulations of these KPIs. We believe that it is possible to find such formulations for a large proportion of the service items.

Finally, only KPIs from the patient's perspective are generated. Future work will also include generation of KPI from the health care providers' perspective, such as revenue and cost efficiency.

**Acknowledgements.** The research described in this paper was partly funded by the Swedish agency for innovation system (VINNOVA).

## References

1. Walburg, J.: The outcome quadrant. In: Walburg, J., Bevan, H., Wilderspin, J., Lemmens, K. (eds.) *Performance Management in Health Care. Improving patient outcomes: an integrated approach*. Routledge, UK (2006)
2. Eckerson, W.: *Performance Dashboards*. In: *Measuring, Monitoring and Managing Your Business*, John Wiley, Chichester (2006)
3. Jun, M., Peterson, R.T., Zsidisin, G.A.: The identification and Measurement of Quality Dimension in Health Care: Focus Group Interview Results. *Health Care Management Review* 23(4), 81–96 (1998)
4. Adair, C.E., Simpson, L., Birdsell, J.M., Omelchuck, K., Casebeer, A.L., Gardiner, H.P., Newman, S., Beckie, A., Clelland, S., Hayden, K.A., Beausejour, P.: *Performance Measurement Systems in Health and Mental Health Services: Models, Practices and Effectiveness. A State of the Science Review*. The Alberta Heritage Foundation for Medical Research (2003)
5. Grönroos, C.: A Service Quality Model and Its Marketing Implication. *European* 18(4), 36–44 (1984)
6. Parasuraman, A., Zeithaml, V.A., Berry, L.L.: SERVQUAL: a multiple-item scale for measuring consumer perception of service quality. *Journal of Retailing* 64(1), 12–40 (1988)
7. Parasuraman, A., Zeithaml, V.A., Berry, L.L.: A conceptual model of service quality and its implication for future research. *Journal of Marketing* 49(4), 41–50 (1985)
8. Gupta, H.D.: Identifying Health Care Quality Constituents: Service Providers' Perspective. *Journal of Management Research* 8(1) (2008)
9. Kilbourne, W.E., Duffy, J.A., Duffy, M., Giarchi, G.: SERVQUAL in cross-national measurements of health care quality. *Journal of Service Marketing* 18(7), 524–533 (2004)
10. Karassavidou, E., Glaveli, N., Papadopoulos, C.T.: Quality in NHS hospitals: no one knows better than patients. *Measuring Business Excellence* 12(1), 34–46 (2009)
11. Gill, L., White, L.: A critical review of patient satisfaction. *Leadership in Health Services*, vol. 22(1), pp. 8–19. Emerald Group Publishing (2009)
12. Kartseva, V., Hulstijn, J., Gordijn, J., Tan, Y.-H.: Control Patterns in a Healthcare Network. In: Brnong, A., and Janssen, M. (eds.) *Proceedings of the 8th Annual International Conference on Digital Government Research (DGO 2007)*, Digital Government Research Center, pp. 66–75 (2007)
13. Johannesson, P., Perjons, E.: Managing Complexity through Value Modeling. *Healthcare IT Management* 4(2) (2009)
14. CONTSYS, EN-13940-1:2007, *Health Informatics – Systems of concepts to support continuity of care. Part 1: Basic Concepts* (2007)
15. UN/CEFACT Modeling Methodology User Guide, <http://www.unece.org/cefact/>
16. Henkel, M., Johannesson, P., Perjons, E., Zdravkovic, J.: Value and Goal Driven Design of E-Services. In: *The IEEE International Conference on e-Business Engineering (ICEBE 2007)*, Hong Kong, China, October 24–26, pp. 295–303 (2007)

**Part II**  
**Biomedical Electronics**  
**and Devices**

# Encapsulated Multi-vesicle Assemblies of Programmable Architecture: Towards Personalized Healthcare

Maik Hadorn<sup>1</sup> and Peter Eggenberger Hotz<sup>2</sup>

<sup>1</sup> Department of Informatics, Artificial Intelligence Laboratory  
University of Zurich, Zurich, Switzerland

<sup>2</sup> The Mærsk Mc-Kinney Møller Institute, University of Southern Denmark  
Odense M, Denmark

hadorn@ifi.uzh.ch, eggen@mimi.sdu.dk

**Abstract.** Although single artificial vesicles are successfully used as delivery vehicles of pharmaceuticals, unilamellarity and restriction to one vessel result in premature content release in physiological environments as well as problems in simultaneous entrapment of a given set of (pharmaceutical) components. Multi-lamellarity and assemblies of distinct populations of vesicles are proposed to solve these problems. In this study, we provide a novel encapsulation protocol to fabricate multilamellar vesicles and we report on the DNA-mediated self-assembly of more than two distinct populations of vesicles. We discuss how these results might be used in personalized healthcare based on custom-tailored encapsulated multicompartment vesicular drug delivery systems.

**Keywords:** Personalized healthcare, Drug delivery, Encapsulation, Compartmentalization, Programmability, Vesicle, Liposome.

## 1 Introduction

Both biological and artificial vesicles feature an aqueous compartment partitioned from an aqueous surrounding by a lipid membrane that is nearly impermeable to hydrophilic substances. This membrane organizes processes by compartmentalizing them. This compartmentalization enables segregation of specific chemical reactions for the purpose of increased controllability, observability, stability, and biochemical efficiency by restricted dissemination and efficient storage of reactants, and/or reaction products. Hence, vesicle-based Living Technology [1] has gained importance in analytics [2-10], synthetics [11-19], and biomedicine [20-33]. Vesicles featuring biocompatibility, biodegradability, low toxicity, and structural variability are successfully utilized as therapeutic agents for the delivery of antibacterial, antiviral, and anticancer drugs, as well as of hormones, enzymes, and nucleotides [34-36].

The use of single unilamellar vesicles prevails in contemporary therapeutic systems. However premature content release in physiological environments limits their reliability [37]. Extending the circulation time of vesicles that results in passive accumulation at tumors or inflammation sites due to the enhanced permeability and retention (EPR) effect [22] is implemented at the molecular level via monomer

design [38] or at the mesoscopic level via encapsulation. The bilayer-within-bilayer structure of encapsulated vesicles not only prevents premature degradation and content release [39] but also offers a division of distinct membrane functions (biocompatibility, cargo release, targeting, and protection) among several membranes of distinct compositions and dimensions.

The applicability of single vesicles is further limited by the need for simultaneous entrapment of a given set of (pharmaceutical) components in one single compartment, which is “not an easy matter” [40, p. 14660]. Multicompartment systems of assembled vesicles can overcome this limitation by conciliating smaller subsets of components entrapped in different compartments.

A combination of encapsulation and (self-)assembled vesicles may provide stable vehicles for multicomponent or multifunctional drug delivery. Zasadzinski *et al.* [30, 39, 41] established a protocol to encapsulate a multicompartment system of tethered vesicles. Both tethering and encapsulation of these vesosomes are based on the molecular recognition process of the biotin-streptavidin complex. In nature, a multivalent “Velcro-like” system allows for selective tethering of a multitude of different cell types. Cell adhesion molecules (CAMs) that provide this selectivity are emulated in the artificial system by several different linking mechanisms [42-51] (for the latest developments in biomimetic supramolecular chemistry see [52]). The multivalent, selective, and sequence-dependent linkage of nucleic acids has considerable potential since it mimics CAMs and offers programmability to the self-assembly process. Hence, single stranded DNA (ssDNA) is used either to induce assembly of hard sphere [53-56] or vesicular [2, 57-59] colloids, to induce programmable fusion of vesicles [2, 59], or to spontaneously and specifically link vesicles to surface supported membranes [2, 57, 60-64]. However, a linkage of more than two populations of vesicles was not implemented. This lack was remedied only recently by the implementation of a DNA-mediated self-assembly of three distinct populations of vesicles reported here and analyzed in depth elsewhere [65].

In this study, we present both a programmable DNA-mediated linkage of three distinct vesicle populations and a novel encapsulation mechanism. Based on the results of this study, we formulate a scenario how encapsulated multicompartment systems might be used to realize custom-tailored vesicular drug delivery systems.


**Table 1.** Membrane composition of vesicles used in experimentation

<i>Encapsulation</i>	
100%	PC(16:0/18:1( $\Delta$ 9-Cis)) = 1-Palmitoyl-2-Oleoyl- <i>sn</i> -Glycero-3-Phosphocholine
<i>Self-Assembly</i>	
99%	PC(16:0/18:1( $\Delta$ 9-Cis)) = 1-Palmitoyl-2-Oleoyl- <i>sn</i> -Glycero-3-Phosphocholine
0.75%	methyl-PEG2000-PE(18:0/18:0) = 1,2-Distearoyl- <i>sn</i> -Glycero-3-Phosphoethanolamine-N- [Methoxy (Polyethylene glycol)-2000]
0.25%	biotin-PEG2000-PE(18:0/18:0) = 1,2-Distearoyl- <i>sn</i> -Glycero-3-Phosphoethanolamine-N- [Biotinyl (Polyethylene Glycol) 2000]


## 2 Material and Methods

Technical modifications of the vesicle formation protocol reported by Pautot *et al.* [66] were: (i) the introduction of 96-well microtiter plates U96 to increase procedural manageability in laboratory experimentation and (ii) a density difference between *inter-* and *intravesicular* solution induced by isomolar solutions of monosaccharids (glucose: *inter*) and disaccharids (sucrose: *intra*). For a description of the modified vesicle protocol see figure 1.A-D. For the membrane composition of the vesicles used in the encapsulation and the self-assembly experiments see table 1. All phospholipids were dissolved in mineral oil.

For details of the encapsulation procedure of untethered vesicles see figure 1.E-H. Encapsulated vesicles exhibited quick random motion within the boundaries of the surrounding vesicle.



**Fig. 1.** Schematic representation of the vesicle formation/encapsulation procedure and micrographs of internally compartmentalized vesicles. (A) A water droplet (solution 1, S1) is added to a phospholipid suspension (S2). (B) A water-in-oil emulsion is produced by mechanical agitation and sonication. (C) The emulsion is placed over an aqueous solution (S3). (D) Induced by centrifugation, the droplets pass the oil/water interface. Due to the density difference of the *inter-* and *intravesicular* fluid and the geometry of the formation chamber, vesicles pelletize in the centre of the well and become easily accessible for pipetting. (E-H) Internally compartmentalized vesicles are prepared by reapplying steps (A-D) using a droplet of the aqueous solution that hosts the vesicles (S4) and an aqueous solution less dense than (S3). (J) Detail of the molecular mechanisms at the water-oil interfaces. Amphiphilic phospholipids, dissolved in mineral oil, stabilize the interfaces by forming monolayers. Two monolayers form a bilayer when a water droplet passes the interface. (K, L) Differential interference contrast micrographs of internally compartmentalized vesicles. Scale bar represents 10µm.



**Fig. 2.** Schematic representation of the self-assembly process. (A) For the vesicle formation procedure see figure 1. Change: Vesicle membranes incorporate biotinylated phospholipids. (B) Vesicle populations (VP) are labelled (grey ring) by incubation with distinct combinations of biotinylated single stranded DNA (b-ssDNA) that differ in sequence and streptavidin (S) that differ in fluorescence labeling (Alexa Fluor 488 / 532 conjugate (AF488 / AF532) or unlabeled). (C) VPs are merged after excess b-ssDNA/S-solutions are removed. VPs become loosely linked in contact. (D) Linker accumulation at the contact sites (solid black) strengthens the linkage. (E) Sequence of complementary biotinylated DNA single strands used in the self-assembly experiments; only bases G/C and A/T pair. (F) Molecular detail of vesicle labeling. B-ssDNA is linked to biotinylated phospholipids through streptavidin as connector. (G) Sequence specific hybridization of ssDNA results in double stranded DNA (dsDNA) and links VPs. Due to their lateral mobility, linkers accumulate at the contact site. The lateral distribution of linkers in the outer leaflet becomes inhomogeneous (sequence dependent).

In the compartmentalization experiments, three distinct vesicle populations were prepared that exposed binary combinations of six ssDNA populations on their surface (1<sup>st</sup> population:  $\alpha$ ,  $\beta$ ; 2<sup>nd</sup>:  $\alpha'$ ,  $\gamma$ ; 3<sup>rd</sup>:  $\beta'$ ,  $\gamma'$ ; for the sequence of biotinylated ssDNA strands see figure 2.E). The DNA strands were biotin-labeled and anchored to biotinylated vesicular membrane via streptavidin as a cross-linking agent. For a detailed protocol of the surface modification and the self-assembly procedure see figure 2.A-D. The sequences of the ssDNA were produced by a genetic algorithm and optimized for minimal DNA-DNA-hybridization among the three pairs.

Light and confocal laser scanning microscopy was performed using an inverted Leica DMR IRE2 SP2 confocal laser scanning microscope.



### 3 Results

The geometry of the microplate wells (U shaped bottom) and the density difference between the *inter*- and *intravesicular* solution induced vesicle pelletization at the centre of the well. The size distribution of vesicles produced in the first round of vesicle formation (Fig. 1.A-D) was shifted to the left when compared to vesicles produced in the second round (Fig. 1.E-H). The two vesicle formation protocols technically differed only in the presence (Fig. 1.B) or absence (Fig. 1.F) of the sonication of the water-in-oil emulsion. To indicate independence of the tethering and encapsulation process, vesicles to be encapsulated were not tethered. Tethered assemblies were encapsulated without any modification of the encapsulation procedure (results not shown). As seen in figures 1.K and 1.L the ratio of vesicles internally compartmentalized to vesicles uncompartimentalized was high. Most of the vesicles produced in the first round were found to be enclosed – encapsulation efficiency was high.

When vesicles that exposed complementary ssDNA came into contact, hybridization of single DNA strands resulted in double stranded DNA. Linkers accumulated in the contact area of the two vesicles formed an adhesion plaque (for a schematic representation see figures 2.D/G; for colour-coded micrographs of real-world results see figures 4.D.1 and 5 of [65], figure 5 of [67], and figure 3 of [68]). Adhesion plaques were found exclusively, when DNA strands were complementary and inorganic ions were present (see [65]). No transfer of linkers between the membranes of different vesicles was observed (data not shown).

### 4 Discussion

Multicomponent or multifunctional custom-tailored vesicular drug delivery systems have to fulfil several requirements: (i) the actual drug containing system should be encapsulated to prevent premature degradation and content release, (ii) the drug containing system should consist of more than two distinct compartments, and (iii) the proper composition of the drug containing system should be controlled.

#### 4.1 Encapsulation

The *in vitro* vesicle formation procedure [12, 66, 69] enables independent tailoring of chemical material properties of the *inter*- and *intravesicular* fluid as well as of the inner and outer membrane leaflet composition. To our knowledge, the entrapment efficiency of this vesicle formation procedure has not, thus far, been analyzed. However one may speculate that its entrapment efficiency is superior to vesicle formation procedures currently used (for an overview of the current vesicle formation procedures see [10]). The potential of an asymmetric leaflet composition was exemplified by the production of phospholipid and polymer hybrids combining biocompatibility and mechanical endurance in single vesicles [66]. We increased procedural manageability of the formation procedure by introducing microtiter plates and vesicle pelletization (due to density differences in the *inter*- and *intravesicular* fluid). By introducing sonication of the water-in-oil emulsion, we could shift the size distribution of the vesicles formed. By refeeding the vesicle containing solution, we established a novel method to produce

multivesicular assemblies. The protocol provides encapsulation of either tethered or untethered vesicular assemblies. The interdependence of tethering and encapsulation, faced in vesosome formation, is therefore resolved.

## 4.2 Compartmentalization

Single stranded DNA provides programmability, specificity, and high degrees of complexity [70]. Streptavidin offers the strongest noncovalent biological interaction known [71], an extensive range of possible vesicle modifications, component modularity, and availability off the shelf. Phospholipid-grafted biotinylated PEG tethers feature lateral mobility [72], high detachment resistance [73], and no intermembrane transfer of linkers. The combination of phospholipid-grafted biotinylated PEG tethers and streptavidin expedites the production of vesicles avoids problems encountered in other approaches using cholesterol-tagged DNA to specifically link different vesicle populations by the hybridization of membrane-anchored DNA [2, 57, 60]: (i) Because the processes of vesicle formation and vesicle modification are not separated (the cholesterol-tagged ssDNA has to be present during vesicle formation), the formation procedure has to be adjusted anew for each change in the vesicle modification. The procedural manageability in laboratory experimentation is therefore reduced. (ii) As previously reported, the cholesterol anchors of the cholesterol-tagged ssDNA spontaneously leave the lipid bilayer and incorporate randomly into (other) lipid bilayers [57]. Thus, the specificity of the linking system is lost over time.


We have presented a DNA-mediated tethering of three distinct vesicle populations, where the linkage of more than two distinct vesicle populations is reported for the first time. Our findings appear to have solved current restrictions, where donor-acceptor mechanisms are binary. The DNA-mediated linkage mechanism, reported here, offers programmability of composition of multicompartment systems. Thus, custom-tailored vesicular drug delivery systems seem feasible.

## 4.3 Composition Control

By loading the vesicular membranes of tethered assemblies by ligand groups not used in the aggregation process, a column chromatographic purification procedure of aggregates may be realized. The ligand groups would be used to purify aggregates from single vesicles (for details see figure 3.A-F). The scenario represents the minimal situation of tethered assemblies of two vesicle populations and two columns in series. If the tethered assemblies consist of three different vesicle populations bearing three different ligand groups not used in the aggregation process, purification of aggregates of proper composition from both single vesicles and incomplete aggregates may be possible.

By a downstream fluorescence activated cell sorting (FACS; for a review of techniques used in cell separation see Pappas and Wang [74]) internally compartmentalized vesicles may be purified from vesicles not compartmentalized properly. By introducing an intermediate separation process, a refeeding (Fig.3.B/E) of single vesicles and incomplete assemblies into the self-assembly process may be realized before they become encapsulated. This may increase encapsulation efficiency and therefore may economize the production of custom-tailored vesicular drug delivery systems.

Encapsulation provides an extended circulation time resulting in accumulation at tumors or inflammation sites due the EPR effect, without the need of specific targeting. On the other hand, multiple compartments offer segregation of multicomponent pharmaceuticals that might be released only when and where they are needed. Permeability control might be realized either by exploitation of stimuli inherent to target site (pH, redox potential, temperature) or externally induced (temperature, magnetic field, ultrasound). For a recent review on stimuli-sensitive pharmaceutical nanocarriers see Torchilin [38].



**Fig. 3.** Scenario how to sort out vesicle assemblies. (A) For vesicle self-assembly procedure see figure 2. Change: Vesicle populations (VPs) are labelled by a pair of complementary biotinylated DNA single strands (b-ssDNA) inducing the self-assembly. Each VP further exposes additional non-complementary b-ssDNA. Complementary strands of these additional strands are exposed on the surface of DNA-coated columns (C1, C2). (A) C1 is loaded by a pool of assembled and unassembled vesicles. (B) Since VP2 and C1 do not complement in sequence of the additional b-ssDNA VP2 is eluted at room temperature (RT). VP1 and vesicle assemblies are retained; (C) but become eluted when temperature is raised by temperature induced DNA denaturation. (D) The new pool of vesicles is fed to C2. (E) At RT VP1 is eluted and VP2 is retained. (F) Only vesicle assemblies are left to be eluted by an increase in temperature. (G) Detail illustrating that both vesicle assemblies and unassembled vesicles of VP1 are retained in C1 at RT. (H) In C2 only vesicle assemblies are retained. (J, K) Molecular details of VP linkage to the columns. Since the DNA strands that establish linkage to columns are shorter than DNA strands that induce self-assembly (cp. Fig.2F,G), they denature at lower temperatures. Linkage of vesicles to the columns hence is less stable than vesicle-vesicle linkage. Release of retained assemblies without breaking the assemblies is feasible.

## 5 Conclusions

Encapsulated multicompartments systems may provide stable vehicles for a multicomponent or multifunctional personalized drug delivery. In this work, we established a novel encapsulation technique and provide evidence for a stable DNA-mediated linkage of more than two vesicle populations. We discussed how these techniques may personalize the individual healthcare by providing custom-tailored vesicular drug delivery systems.

**Acknowledgements.** Maik Hadorn was supported by the Swiss National Foundation Project 200020-118127 Embryogenic Evolution: From Simulations to Robotic Applications. The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 249032. All experiments were performed at the Institute of Plant Biology, University of Zurich. We thank Enrico Martinoia for providing laboratory equipment and Bo Burla (both Institute of Plant Biology, University of Zurich) for his qualified assistance in various experiments. Eva Bönzli (Faculty of Veterinary Medicine, University of Zurich) and Shaun Peters (Institute of Plant Biology, University of Zurich) provided thoughtful discussion and comments on the manuscript. Moreover, we thank the reviewers for their helpful comments.

## References

1. Bedau, M.A., McCaskill, J.S., Packard, N.H., Rasmussen, S.: Living Technology: Exploiting Life's Principles in Technology. *Artif. Life* 16, 89–97 (2009)
2. Chan, Y.H.M., van Lengerich, B., Boxer, S.G.: Effects of linker sequences on vesicle fusion mediated by lipid-anchored DNA oligonucleotides. *Proc. Natl. Acad. Sci. USA* 106, 979–984 (2009)
3. Hase, M., Yoshikawa, K.: Structural transition of actin filament in a cell-sized water droplet with a phospholipid membrane. *J. Chem. Phys.* 124, 104903 (2006)
4. Hotani, H., Nomura, F., Suzuki, Y.: Giant liposomes: from membrane dynamics to cell morphogenesis. *Curr. Opin. Colloid Interface Sci.* 4, 358–368 (1999)
5. Limozin, L., Roth, A., Sackmann, E.: Microviscoelastic moduli of biomimetic cell envelopes. *Phys. Rev. Lett.* 95, 178101 (2005)
6. Luisi, P., Walde, P.: Giant vesicles. John Wiley & Sons, Ltd, Chichester (2000)
7. Gomez-Hens, A., Fernandez-Romero, J.M.: The role of liposomes in analytical processes. *Trac-Trends Anal. Chem.* 24, 9–19 (2005)
8. Owen, R.L., Strasters, J.K., Breyer, E.D.: Lipid vesicles in capillary electrophoretic techniques: Characterization of structural properties and associated membrane-molecule interactions. *Electrophoresis* 26, 735–751 (2005)
9. Wiedmer, S.K., Jussila, M.S., Riekkola, M.L.: Phospholipids and liposomes in liquid chromatographic and capillary electromigration techniques. *Trac-Trends Anal. Chem.* 23, 562–582 (2004)
10. Jesorka, A., Orwar, O.: Liposomes: Technologies and Analytical Applications. *Annu. Rev. Anal. Chem.* 1, 801–832 (2008)
11. Michel, M., Winterhalter, M., Darbois, L., Hemmerle, J., Voegel, J.C., Schaaf, P., Ball, V.: Giant liposome microreactors for controlled production of calcium phosphate crystals. *Langmuir* 20, 6127–6133 (2004)

12. Noireaux, V., Libchaber, A.: A vesicle bioreactor as a step toward an artificial cell assembly. *Proc. Natl. Acad. Sci. USA* 101, 17669–17674 (2004)
13. Nomura, S., Tsumoto, K., Hamada, T., Akiyoshi, K., Nakatani, Y., Yoshikawa, K.: Gene expression within cell-sized lipid vesicles. *Chembiochem.* 4, 1172–1175 (2003)
14. Bolinger, P.Y., Stamou, D., Vogel, H.: Integrated nanoreactor systems: Triggering the release and mixing of compounds inside single vesicles. *J. Am. Chem. Soc.* 126, 8594–8595 (2004)
15. Bolinger, P.Y., Stamou, D., Vogel, H.: An integrated self-assembled nanofluidic system for controlled biological chemistries. *Angew. Chem.-Int. Edit.* 47, 5544–5549 (2008)
16. Chiu, D.T., Wilson, C.F., Ryttsen, F., Stromberg, A., Farre, C., Karlsson, A., Nordholm, S., Gaggari, A., Modi, B.P., Moscho, A., Garza-Lopez, R.A., Orwar, O., Zare, R.N.: Chemical transformations in individual ultrasmall biomimetic containers. *Science* 283, 1892–1895 (1999)
17. Kuruma, Y., Stano, P., Ueda, T., Luisi, P.L.: A synthetic biology approach to the construction of membrane proteins in semi-synthetic minimal cells. *Biochim. Biophys. Acta-Biomembr.* 1788, 567–574 (2009)
18. Kita, H., Matsuura, T., Sunami, T., Hosoda, K., Ichihashi, N., Tsukada, K., Urabe, I., Yomo, T.: Replication of Genetic Information with Self-Encoded Replicase in Liposomes. *Chembiochem.* 9, 2403–2410 (2008)
19. Chiarabelli, C., Stano, P., Luisi, P.L.: Chemical approaches to synthetic biology. *Curr. Opin. Biotechnol.* 20, 492–497 (2009)
20. Abraham, S.A., Waterhouse, D.N., Mayer, L.D., Cullis, P.R., Madden, T.D., Bally, M.B.: The liposomal formulation of doxorubicin. In: *Liposomes*. Pt E. Elsevier Academic Press Inc., San Diego (2005)
21. Allen, T.M.: Liposomal drug delivery. *Curr. Opin. Colloid Interface Sci.* 1, 645–651 (1996)
22. Allen, T.M., Cullis, P.R.: Drug delivery systems: Entering the mainstream. *Science* 303, 1818–1822 (2004)
23. Allen, T.M., Hansen, C., Martin, F., Redemann, C., Yauyoung, A.: Liposomes containing synthetic lipid derivatives of poly(ethylene glycol) show prolonged circulation half-lives in vivo. *Biochimica Et Biophysica Acta* 1066, 29–36 (1991)
24. Allen, T.M., Hansen, C.B., Demenezes, D.E.L.: Pharmacokinetics of long-circulating liposomes. *Adv. Drug Deliv. Rev.* 16, 267–284 (1995)
25. Boutorin, A.S., Guskova, L.V., Ivanova, E.M., Kobetz, N.D., Zarytova, V.F., Rytte, A.S., Yurchenko, L.V., Vlassov, V.V.: Synthesis of alkylating oligonucleotide derivatives containing cholesterol or phenazinium residues at their 3'-terminus and their interaction with DNA within mammalian-cells. *FEBS Lett.* 254, 129–132 (1989)
26. Marjan, J.M.J., Allen, T.M.: Long circulating liposomes: Past, present and future. *Biotechnology Advances* 14, 151–175 (1996)
27. Tardi, P.G., Boman, N.L., Cullis, P.R.: Liposomal doxorubicin. *J. Drug Target.* 4, 129–140 (1996)
28. Sengupta, S., Eavarone, D., Capila, I., Zhao, G.L., Watson, N., Kiziltepe, T., Sasisekharan, R.: Temporal targeting of tumour cells and neovasculature with a nanoscale delivery system. *Nature* 436, 568–572 (2005)
29. Hadorn, M., Eggenberger Hotz, P.: Towards Personalized Drug Delivery: Preparation of an Encapsulated Multicompartment System. In: *3rd International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC)*, Valencia, Spain (in press, 2010)
30. Kisak, E., Coldren, B., Evans, C., Boyer, C., Zasadzinski, J.: The vesosome - A multi-compartment drug delivery vehicle. *Current medicinal chemistry* 11, 199–220 (2004)
31. Torchilin, V.P.: Recent advances with liposomes as pharmaceutical carriers. *Nat. Rev. Drug Discov.* 4, 145–160 (2005)

32. Bonacucina, G., Cespi, M., Misici-Falzi, M., Palmieri, G.F.: Colloidal Soft Matter as Drug Delivery System. *J. Pharm. Sci.* 98, 1–42 (2009)
33. Theodossiou, T.A., Galanou, M.C., Paleos, C.M.: Novel amiodarone-doxorubicin cocktail liposomes enhance doxorubicin retention and cytotoxicity in DU145 human prostate carcinoma cells. *J. Med. Chem.* 51, 6067–6074 (2008)
34. Lasic, D., Vallner, J., Working, P.: Sterically stabilized liposomes in cancer therapy and gene delivery. *Current opinion in molecular therapeutics* 1, 177–185 (1999)
35. Eckstein, F.: The versatility of oligonucleotides as potential therapeutics. *Expert Opin. Biol. Ther.* 7, 1021–1034 (2007)
36. Weissig, V., Boddapati, S., Cheng, S., D'souza, G.: Liposomes and liposome-like vesicles for drug and DNA delivery to mitochondria. *Journal of Liposome Research* 16, 249–264 (2006)
37. Bakker-Woudenberg, I., Schifflers, R.M., Storm, G., Becker, M.J., Guo, L.: Long-circulating sterically stabilized liposomes in the treatment of infections. In: *Liposomes*. Pt E. Elsevier Academic Press Inc., San Diego (2005)
38. Torchilin, V.: Multifunctional and stimuli-sensitive pharmaceutical nanocarriers. *Eur. J. Pharm. Biopharm.* 71, 431–444 (2009)
39. Boyer, C., Zasadzinski, J.A.: Multiple lipid compartments slow vesicle contents release in lipases and serum. *ACS Nano* 1, 176–182 (2007)
40. Luisi, P.L., de Souza, T.P., Stano, P.: Vesicle Behavior: In Search of Explanations. *J. Phys. Chem. B* 112, 14655–14664 (2008)
41. Walker, S.A., Kennedy, M.T., Zasadzinski, J.A.: Encapsulation of bilayer vesicles by self-assembly. *Nature* 387, 61–64 (1997)
42. Vermette, P., Taylor, S., Dunstan, D., Meagher, L.: Control over PEGylated-liposome aggregation by NeutrAvidin-biotin interactions investigated by photon correlation spectroscopy. *Langmuir* 18, 505–511 (2002)
43. Menger, F.M., Seredyuk, V.A., Yaroslavov, A.A.: Adhesive and anti-adhesive agents in giant vesicles. *Angew. Chem.-Int. Edit.* 41, 1350–1352 (2002)
44. Berti, D., Baglioni, P., Bonaccio, S., Barsacchi-Bo, G., Luisi, P.L.: Base complementarity and nucleoside recognition in phosphatidyl nucleoside vesicles. *J. Phys. Chem. B* 102, 303–308 (1998)
45. Sideratou, Z., Foundis, J., Tsiourvas, D., Nezis, I.P., Papadimas, G., Paleos, C.M.: A novel dendrimeric “glue” for adhesion of phosphatidyl choline-based liposomes. *Langmuir* 18, 5036–5039 (2002)
46. Marchi-Artzner, V., Gulik-Krzywicki, T., Guedeau-Boudeville, M.A., Gosse, C., Sander-son, J.M., Dedieu, J.C., Lehn, J.M.: Selective adhesion, lipid exchange and membrane-fusion processes between vesicles of various sizes bearing complementary molecular recognition groups. *Chem. Phys. Chem.* 2, 367–376 (2001)
47. Paleos, C.M., Sideratou, Z., Tsiourvas, D.: Mixed vesicles of didodecyldimethylammonium bromide with recognizable moieties at the interface. *J. Phys. Chem.* 100, 13898–13900 (1996)
48. Constable, E.C., Meier, W., Nardin, C., Mundwiler, S.: Reversible metal-directed assembly of clusters of vesicles. *Chem. Commun.*, 1483–1484 (1999)
49. Chiruvolu, S., Walker, S., Israelachvili, J., Schmitt, F.J., Leckband, D., Zasadzinski, J.A.: Higher-order self-assembly of vesicles by site-specific binding. *Science* 264, 1753–1756 (1994)
50. Noppleson, D.A., Needham, D.: Avidin-biotin interactions at vesicle surfaces: Adsorption and binding, cross-bridge formation, and lateral interactions. *Biophys. J.* 70, 1391–1401 (1996)
51. Weikl, T.R., Groves, J.T., Lipowsky, R.: Pattern formation during adhesion of multicomponent membranes. *Europhys. Lett.* 59, 916–922 (2002)
52. Voskuhl, J., Ravoo, B.J.: Molecular recognition of bilayer vesicles. *Chem. Soc. Rev.* 38, 495–505 (2009)

53. Biancaniello, P.L., Crocker, J.C., Hammer, D.A., Milam, V.T.: DNA-mediated phase behavior of microsphere suspensions. *Langmuir* 23, 2688–2693 (2007)
54. Mirkin, C.A., Letsinger, R.L., Mucic, R.C., Storhoff, J.J.: A DNA-based method for rationally assembling nanoparticles into macroscopic materials. *Nature* 382, 607–609 (1996)
55. Valignat, M.P., Theodoly, O., Crocker, J.C., Russel, W.B., Chaikin, P.M.: Reversible self-assembly and directed assembly of DNA-linked micrometer-sized colloids. *Proc. Natl. Acad. Sci. USA* 102, 4225–4229 (2005)
56. Biancaniello, P., Kim, A., Crocker, J.: Colloidal interactions and self-assembly using DNA hybridization. *Phys. Rev. Lett.* 94, 058302 (2005)
57. Beales, P.A., Vanderlick, T.K.: Specific binding of different vesicle populations by the hybridization of membrane-anchored DNA. *J. Phys. Chem. A* 111, 12372–12380 (2007)
58. Beales, P.A., Vanderlick, T.K.: DNA as Membrane-Bound Ligand-Receptor Pairs: Duplex Stability Is Tuned by Intermembrane Forces. *Biophys. J.* 96, 1554–1565 (2009)
59. Stengel, G., Zahn, R., Hook, F.: DNA-induced programmable fusion of phospholipid vesicles. *J. Am. Chem. Soc.* 129, 9584–9585 (2007)
60. Benkoski, J.J., Hook, F.: Lateral mobility of tethered vesicle - DNA assemblies. *J. Phys. Chem. B* 109, 9773–9779 (2005)
61. Yoshina-Ishii, C., Boxer, S.G.: Arrays of mobile tethered vesicles on supported lipid bilayers. *J. Am. Chem. Soc.* 125, 3696–3697 (2003)
62. Li, F., Pincet, F., Perez, E., Eng, W.S., Melia, T.B.J., Rothman, J.E., Tareste, D.: Energetics and dynamics of SNAREpin folding across lipid bilayers. *Nat. Struct. Mol. Biol.* 14, 890–896 (2007)
63. Svedhem, S., Pfeiffer, I., Larsson, C., Wingren, C., Borrebaeck, C., Hook, F.: Patterns of DNA-labeled and scFv-antibody-carrying lipid vesicles directed by material-specific immobilization of DNA and supported lipid bilayer formation on an Au/SiO<sub>2</sub> template. *ChemBiochem.* 4, 339–343 (2003)
64. Stadler, B., Falconnet, D., Pfeiffer, I., Hook, F., Voros, J.: Micropatterning of DNA-tagged vesicles. *Langmuir* 20, 11348–11354 (2004)
65. Hadorn, M., Eggenberger Hotz, P.: DNA-Mediated Self-Assembly of Artificial Vesicles. *PLoS One* 5 5, e9886 (2010)
66. Pautot, S., Frisken, B.J., Weitz, D.A.: Engineering asymmetric vesicles. *Proc. Natl. Acad. Sci. USA* 100, 10718–10721 (2003)
67. Hadorn, M., Burla, B., Eggenberger Hotz, P.: Towards tailored communication networks in assemblies of artificial cells. In: Korb, K., Randall, M., Hendtlass, T. (eds.) *ACAL 2009. LNCS*, vol. 5865, pp. 126–135. Springer, Heidelberg (2009)
68. Hadorn, M., Eggenberger Hotz, P.: Multivesicular assemblies as real-world testbeds for embryogenic evolutionary systems. In: Korb, K., Randall, M., Hendtlass, T. (eds.) *ACAL 2009. LNCS*, vol. 5865, pp. 169–178. Springer, Heidelberg (2009)
69. Träuble, H., Grell, E.: Carriers and specificity in membranes. IV. Model vesicles and membranes. The formation of asymmetrical spherical lecithin vesicles. *Neurosciences Research Program bulletin* 9, 373–380 (1971)
70. Licata, N.A., Tkachenko, A.V.: Errorproof programmable self-assembly of DNA-nanoparticle clusters. *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)* 74, 41406 (2006)
71. Green, N.M.: Avidin and streptavidin. *Method Enzymol.* 184, 51–67 (1990)
72. Singer, S.J., Nicolson, G.L.: Fluid mosaic model of structure of cell-membranes. *Science* 175, 720–731 (1972)
73. Burridge, K.A., Figa, M.A., Wong, J.Y.: Patterning adjacent supported lipid bilayers of desired composition to investigate receptor-ligand binding under shear flow. *Langmuir* 20, 10252–10259 (2004)
74. Pappas, D., Wang, K.: Cellular separations: A review of new challenges in analytical chemistry. *Anal. Chim. Acta* 601, 26–35 (2007)

# Adaptive RBF-HMM Bi-Stage Classifier Applied to Brain Computer Interface

José Luis Martínez Pérez and Antonio Barrientos Cruz

Grupo de Robótica y Cibernética, Universidad Politécnica de Madrid  
C/José Gutierrez Abascal 2, Madrid, Spain  
jlmartinez@etsii.upm.es, antonio.barrientos@upm.es  
<http://www.robcib.etsii.upm.es>

**Abstract.** Brain Computer Interface is a new technology aimed to communicate the user's intentions without using nerves or muscles. To obtain this objective, BCI devices make use of classifiers which translate inputs from the user's brain signals into commands for external devices. This paper describes an adaptive bi-stage classifier based on RBF neural networks and Hidden Markov Models. The first stage analyses the user's electroencephalographic input signal and provides sequences of pre-assignments to the second stage. The segment of EEG signal is assigned to the HMM with the highest probability of generating the pre-assignment sequence.

The algorithm is tested with real samples of electroencephalographic signal, from five healthy volunteers using the cross-validation method. The results allow to conclude that it is possible to implement this algorithm in an on-line BCI device, but a huge dependency in the percentage of the correct classification from the user and the setup parameters has been detected.

**Keywords:** Electroencephalography, Brain computer interface, Linear discriminant analysis, Spectral analysis, Biomedical signal detection, Pattern recognition.


## 1 Introduction

Since the discovery of the electrical nature of the brain by Dr. Hans Berger, it has been considered the possibility to communicate persons with external devices only through the use of the brain waves [1].

Brain Computer Interface technology [2] is aimed at communicating with persons using external computerised devices via the electroencephalographic signal as the primary command source [3]; in the first international meeting for BCI technology it was established that BCI “*must not depend on the brain's normal output pathways of peripheral nerves and muscles*” [4]. The primary uses of this technology are to benefit persons with blocking diseases, such as: Amiotrophic Lateral Sclerosis (ALS), brain-stem stroke, or cerebral palsy; or persons whom have suffered some kind of traumatic accident like for example paraplegic [5].

One way to control an external device using thoughts is associating selected brain activity to device commands. Therefore, an algorithm that detects, acquires, filters, and classifies the EEG signal is required [6] [7].





**Fig. 1.** Block diagram of a BCI device

Actually different types of classifications can be established for BCI technology, from the physiologic point of view BCI devices can be classified in exogenous and endogenous. The devices in the first group provide some kind of stimuli to the user and they analyse the user's responds to them, examples of this class are devices based on visual evoked potential or P300 [5]. On the contrary, the endogenous devices do not depend on the user's respond to external stimuli, they base their operation in the detection and recognition of brain-wave patterns controlled autonomously by the user, examples of this class are devices based on the desynchronization and synchronization of  $\mu$  and  $\beta$  rhythms [4], [8], [9].

In this paper is presented an endogenous classifier composed of two adaptive stages.

In the first stage a Radial Basis Function neural network [10] performs a pre-classification of the segment of EEG input signal and provides a pre-assignment sequence of data to the second stage.

In this second stage is computed the generation probability of the input sequence by three different Hidden Markov Models, each one had been previously trained with data sequences from the different cognitive activities between classifying.

The content of this paper is as follows:

- Section 2 describes the experimental procedure, the cognitive activities or mental tasks, and the equipment to develop the experiment.
- Section 3 presents the proposed two-stage classifier and describes the algorithms employed in it for training and operation.
- Section 4 presents the results obtained from each volunteer.
- Section 5 analyses the previous results.
- And finally section 6 is devoted to making conclusions.

## 2 Experimental Procedure

Two sessions of the tests described below were carried out on five healthy male subjects, one of whom had been trained before but the other four of whom were novices in the use of the system.

In order to facilitate the mental concentration necessary for the proposed activities, the experiments were carried out in a room with a low noise level and under controlled environmental conditions. The experiments were carry out between 10 a.m. and 14 p.m.; the period of time between sessions was one day.

The subjects were directed to sit down 50 cm from the screen of the acquisition system monitor, and with their hands in a visible position. The supervisor of the experiment ensured the correct enactment of this process.

### 2.1 Flow of Activities in the Experimental Process

The experimental process is shown in Figure 2.


*System devices test.* The correct battery level and correct state of the electrodes were checked.

*System assembly.* Device connections: superficial electrodes (Grass Au-Cu), battery, bio-amplifier (g.BSamp by g.tec), acquisition signal card (PCI-MIO-16/E-4 by National Instrument), computer.

*System test.* The correct operation of the whole system was verified. To minimise noise from the electrical network the Notch filter (50Hz) of the bio-amplifier was switched on.

*Subject preparation for the experiment.* Electrodes were applied to the subject's head. Electrode impedance  $\leq 4\text{KOhms}$ .

*System initialisation and setup.* The data register was verified. The signal evolution was checked; it was imperative that a very low component of 50 Hz appeared within the spectrogram.



**Fig. 2.** Diagram of the experiment carried out

*Experiment setup.* The experiment supervisor set up the number of replications  $N_{rep} = 10$ , and the quantity of different mental activities. The duration of each trial was  $t = 7s$ , and the acquisition frequency was  $f_s = 384Hz$ . The system suggested that the subject think about the proposed mental activity. A short relaxation period was allowed at the end of each trial; between replications the relax time is  $t = 7s$ .

## 2.2 Electrode Position

Electrodes were placed in the central zone of the skull, next to C3 and C4 [11], and two pairs of electrodes were placed in front of and behind the Rolandic sulcus. This zone has the highest discriminatory power and receives signals from motor and sensory areas of the brain [3] [12]. A reference electrode was placed on the right mastoid and two more electrodes were placed near the corners of the eyes to register blinking.

## 2.3 Description of Cognitive Activities

The experiment supervisor asked the subject to figure out the following mental activities; tasks that were used to differentiate between cerebral patterns [12]:



Fig. 3. Electrode placement

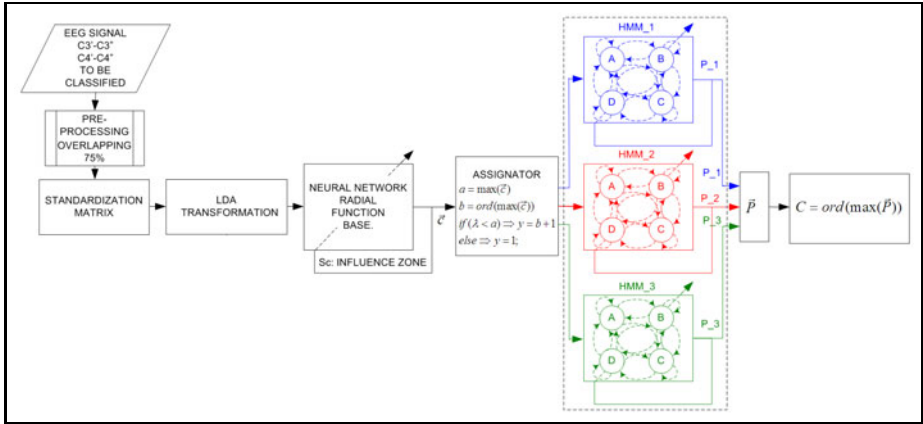


Fig. 4. Block diagram of the classifier

Activity A. *Mathematical task.* Recursive subtraction of a prime number, i.e. 7, from a large quantity, i.e. 3.000.000.

Activity B. *Movement task.* The subject imagines moving their limbs or hands, but without actually doing so. It is movement imagery.

Activity C. *Relax.* The subject relaxes.

### 3 Description of the Classifier

#### 3.1 Introduction

In Fig. 4 is shown the block diagram of the algorithm for the proposed classifier.

In it can be appreciated how the classification of the considered segment of the EEG signal is obtained after the evaluation of the probability generation of the pre-assignment sequence provided by three Hidden Markov Models.

There are as many Hidden Markov Models as cognitive activities to be considered for the classification, each model is trained with pre-assignment sequences of data of the cognitive activity associated to it.

The pre-assignment sequence of data are provided by a neural network, which inputs are the vectors of features obtained after the preprocessing of the segment of EEG signal, as it is described in the following subsections.




Fig. 5. Block diagram for the preprocessing phase

### 3.2 Preprocessing

In Fig. 5 is depicted the operations associated to the preprocessing phase.

In a first step the EEG signal coming from both electroencephalographic channels C3'-C3'' and C4'-C4'' are sampled and quantified at  $f_s = 384Hz$ .

In the next step the samples are bundled in package of  $N_w = 128$  samples, it is equivalent to  $T_w = 1/3s$ . Each group of samples is normalised, in order to obtain homogeneous groups of transformed samples with zero mean value and unity as standard deviation.

$$x'_i = \frac{x_i - \mu}{\sigma} \quad (1)$$

This transformation does not affect the frequency properties of the signal, but allows the comparison of groups of samples in the same session or between sessions, it avoids that changes in the impedance of the electrodes or skin conductivity affect the next procedures.

After this the samples are processed by a pass band Butterworth filter of order  $n = 6$ , with lower and higher frequencies of  $f_1 = 4Hz$  and  $f_2 = 40Hz$  [13]; frequencies outside this band are not common in sane conscious users.

In the next step the filtered samples are convoluted with a Tukey's window of length  $N = 128$ , this attenuates the leakage effect associated to the package procedure, see Fig. 6. Previous studies allow to conclude that the convolution of the signal with this kind of window increase the discrimination capability that the one obtained with other kind of windows as for example: rectangular, triangular, Blackman's, Hamming's, Hanning's or Kaiser's [14].

After the convolution, the Fast Fourier Transform, eq 2 is applied in order to obtain the spectral power estimation, eq 3; the obtained frequency resolution is in eq 4

$$X(k) = \sum_{j=1}^N x(j)w_N^{(j-1)(k-1)}; \quad w_N = e^{-\frac{2\pi i}{N}} \quad (2)$$

$$P(k) = \frac{X(k)^2}{N} \quad (3)$$

$$\Delta f = \frac{f_s}{N_w} = \frac{384}{128} = 3Hz \quad (4)$$




Fig. 6. Frequency leakage effect

After the estimation of each frequency band, it is computed the vector of features considering the power average of the involved bands as it is shown on Table 1.

Table 1. Feature vector

Index	Denomination.	Frequency (Hz).
1	$\theta.$	6 - 8
2	$\alpha_1.$	9 - 11
3	$\alpha_2.$	12 - 14
4	$\beta_1.$	15 - 20
5	$\beta_2.$	21 - 29
6	$\beta_3.$	30 - 38

In case of presence of artifacts the algorithm detects them and during the learning phase it substitutes its value by the average value of the samples in that package, if the artifacts are detected in the on-line phase, it instructs the classifier to discard that group of samples.

A group of samples is considered with artifacts if one sample differs more than three standard deviations from the previous one.

### 3.3 Training of the Neural Network

The considered neural network is the type of Radial Function Basis. This type of neural network is characterised by the learning of the position of the samples in the training set and by the interpolation capability between them [10].

In Fig. 7 is represented the architecture of this type of neural network.




Fig. 7. Architecture of the RBF neural network

From previous studies it has been concluded that this type of neural network behaves better than other types of neural networks, as for example Multi-Layer Perceptrons or Probabilistic Neural Networks [15].

The activation function is:

$$\text{radbas}(x) = e^{-x^2}; \quad x = (w - p) * S_c \quad (5)$$

In where  $w$  and  $S_c$  are respectively the weights and influence zone constant of each neuron, and  $p$  is the position of the considered sample.

During the learning phase the neurons of the hidden layer learn the position of the samples of the learning set,  $w$ ; during the test phase when a new sample  $p$  is presented, it is computed the distance between the sample and the learned positions, the nearest neurons to the sample will proportionate higher activation values than the rest of the neurons.

For the learning process are considered vectors of features from the EEG signal, acquired when the user was performing one of the different cognitive activities considered for the classification. The learning set is composed by the 75% of all the sample set, and the other 25% is considered for validation. After the determination of the learning and validation sets, the input vectors to the neural network are normalised, and with LDA technique is reduced their dimensionality projecting the original input vectors in the direction of the highest discrimination capability [16].

In order to minimise the over-learning effect, the RBF learning process allows a dynamic growth of the number of neurons in the hidden layer. In the output layer are considered as many linear neurons as cognitive activities between discriminate. Finally in the assignment block on Figures 4 it is weighted the output vector of the neural network and it assigns the input vector to the activity with highest output value provided it is higher than a threshold  $\lambda$ , on the contrary if the value is lower than  $\lambda$ , the input vector is labelled as unclassified.

On operation, once the neuronal network has been trained, when a new vector is presented the cognitive activity, with samples nearer to it, will provide a higher activation level, and the corresponding output will have a higher value than the others cognitive activities.




Fig. 8. Training of the RBF neural network

### 3.4 Description of Hidden Markov Models

A Hidden Markov Model is a double stochastic statistical model, it consists of a Markov process with unknown and non-observable parameters, and a observed model which parameters depend stochastically from the hidden states. A stochastic process is called a Markovian process if the future does not depend from the past, only from the known present; considering the stochastic variable  $q(t - 1)$  the transition probability in the instant  $t$  is defined as  $P(q_t = \sigma_t | q_{t-1} = \sigma_{t-1})$ . A Markov chain is formally defined with the pair  $(Q, A)$ , where  $Q = \{1, 2, \dots, N\}$  are the possible estates of the chain and  $A = [a_{ij}]$  is the transition matrix of the model, with the constrains:

$$0 \leq a_{ij} \leq 1; \quad 1 \leq i, j \leq N \tag{6}$$

$$\sum_{j=1}^N a_{ij} = 1; \quad 1 \leq i \leq N \tag{7}$$

The transition and emission probabilities depends from the actual estate and no from the former estates.

$$\begin{aligned} P(q_t = j | q_{t-1} = i, q_{t-2} = k, \dots) &= \\ &= P(q_t = j | q_{t-1} = i) = a_{ij}(t) \end{aligned} \tag{8}$$

Formally a discrete HMM of first grade is defined by the 5-tuple:

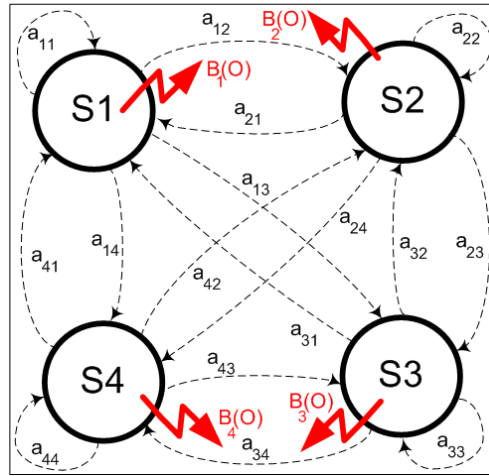
$\lambda = \{Z, Q, A, B, \pi\}$ , in where:

- $Z = \{V_1, V_2, \dots, V_M\}$  is the alphabet or discrete set of  $M$  symbols.
- $Q = \{1, 2, \dots, N\}$  is the set of  $N$  finite estates.
- $A = [a_{ij}]$  is the transition matrix of the model.
- $B = (b_j(Q_i))_{N \times M}$  is the matrix of emission symbols, also known as observation matrix.
- $\pi = (\pi_1, \pi_2, \dots, \pi_N)$  is the prior probability vector of the initial estate.

The parameters of a HMM are  $\lambda = \{A, B, \pi\}$ , see Fig. 9

There are three types of canonic problems associated to HMM [17] [18]:





**Fig. 9.** Example of Hidden Markov Model

1. Given the parameters of the model, obtain the probability of a particular output sequence. This problem is solved through a forward-backwards algorithm.
2. Given the parameters of the model, find the most probable sequence of hidden states, that could generate the given output sequence. This problem is solved through the use of Viterbi algorithm.
3. Given an output sequence, find the parameters of the HMM. This problem is solved through the use of Baum-Welch algorithm.

The HMM have been applied specially in speech recognition an generally in temporal sequences, hand written, gestures recognition, and bioinformatics [18].

### 3.5 Training of the Hidden Markov Models

The HMM's are trained with sequences of pre-assignments coming from the EEG samples, as it is shown in the Figure 10.

For each cognitive activity a particular HMM, with the following characteristics, is trained:

- Number of hidden states: 4.
- Number of different observable objects: 4

In the training phase, chains of nine pre-assignments were used. In a previous experiment with synthetic samples, it was concluded that for the proposed architecture of Hidden Markov Models the highest percentage of correct classifications were obtained with chains of nine elements.

After the training or solution of the third canonic problem, the probability matrices of state transitions and observation matrices are determined. The Viterbi algorithm is used in order to determine the probability that a model generates the proposed sequence.




Fig. 10. Training of the HMM

### 4 Results

In order to test the behaviour of the proposed algorithm, the influence of the threshold assignation parameter ( $\lambda$ ), and the influence zone of the neuron ( $S_c$ ), the EEG samples of the session tests from the volunteers were used as follows:

#### 4.1 Evaluation of the Learning Capability

With a subset of 75% of the all EEG samples the algorithm was trained with different  $\lambda$  and  $S_c$  values:

$\lambda =$	0.55	0.65	0.8
$S_c =$	0.5	0.95	

These values have been fixed after a seek in wide with the samples of the first volunteer.

After the learning, the same samples were processed with the trained algorithm, and a comparison between the results obtained with the algorithm and the ones employed for the learning was done, in all cases a 100% of correct classification has been obtained.

#### 4.2 Evaluation of the Generalisation Capability

After the good results obtained from the learning phase, a cross-validation methodology is used to estimate the generalisation capability. From the whole ten sessions, nine are used for learning and one is used for validation, the process is repeated ten times changing each time the session used for validation.

In the following tables are shown the results obtained for each volunteer, considering the  $\lambda$  and  $S_c$  parameters.

For each combination, the process is replicated three times. In the upper row it is shown the number of correct classifications. In the lower row it is shown the percentage of improvement against a naive classifier. The results of the second session follow the results of the first one.

**Table 2.** Volunteer: AI01 & AI02

$S_c = 0.5$ $\lambda = 0.65$			$S_c = 0.95$ $\lambda = 0.55$			$S_c = 0.5$ $\lambda = 0.55$			$S_c = 0.95$ $\lambda = 0.80$		
94	103	103	94	81	87	93	92	87	86	97	81
+4	+14	+14	+4	-10	-3	+3	+2	-3	-4	+8	-10
98	107	97	103	98	97	96	90	99	112	117	107
+9	+19	+8	+14	+9	+8	+7	0	+10	+24	+30	+19

**Table 3.** Volunteer: Ro01 & Ro02

$S_c = 0.5$ $\lambda = 0.65$			$S_c = 0.95$ $\lambda = 0.55$			$S_c = 0.5$ $\lambda = 0.55$			$S_c = 0.95$ $\lambda = 0.80$		
103	97	92	118	109	118	97	87	86	117	106	110
+14	+8	+2	+31	+21	+31	+8	-3	-4	+30	+18	+22
115	106	110	103	98	102	102	83	81	115	101	117
+28	+18	+22	+14	+9	+12	+12	-8	-10	+27	+12	+30

**Table 4.** Volunteer: Ja01 & Ja02

$S_c = 0.5$ $\lambda = 0.65$			$S_c = 0.95$ $\lambda = 0.55$			$S_c = 0.5$ $\lambda = 0.55$			$S_c = 0.95$ $\lambda = 0.80$		
106	97	110	87	90	107	99	106	107	98	108	99
+18	+8	+22	-3	0	+19	+10	+18	+19	+9	+20	+10
86	90	103	106	109	104	103	97	105	106	112	109
-4	0	+14	+18	+21	+15	+14	+8	+16	+18	+24	+21

**Table 5.** Volunteer: Da01 & Da02

$S_c = 0.5$ $\lambda = 0.65$			$S_c = 0.95$ $\lambda = 0.55$			$S_c = 0.5$ $\lambda = 0.55$			$S_c = 0.95$ $\lambda = 0.80$		
109	102	104	83	92	92	106	91	110	86	87	92
+21	+13	+15	-8	+2	+2	+18	+1	+22	-4	-3	+2
106	94	110	84	85	92	109	104	106	116	108	110
+18	+4	+22	-7	-5	+1	+21	+15	+18	+28	+20	+22

**Table 6.** Volunteer: Ra01 & Ra02

$S_c = 0.5$ $\lambda = 0.65$			$S_c = 0.95$ $\lambda = 0.55$			$S_c = 0.5$ $\lambda = 0.55$			$S_c = 0.95$ $\lambda = 0.80$		
106	97	110	87	90	107	99	106	107	91	76	99
+18	+8	+22	-3	0	+19	+10	+18	+19	+1	-15	+10
102	102	98	102	107	114	103	105	96	116	99	98
+13	+13	+9	+13	+19	+26	+14	+16	+6	+29	+10	+9

## 5 Discussion

From the results of the proposed classification algorithm, it is observed that:

- The learning capability is better than the one achieved only with the RBF neural network [15].
- From the analysis of the results of the replicas it has been detected that the variability in the percentage of correct classifications is caused by the HMM's, both in the learning and validation phases. The sequences of pre-assignments provided by the neural network were stable, but the generation probabilities of the HMM's changed in each replica. In the learning phase the HMM's probabilities allowed a perfect classification, but they were not maintained in the cross validation phase; for this stage a lower percentage of correct classification was obtained, as it is summarised in the tables 2 to 6. But until in this case, almost in all replicas, the cross-validation test results were better than the ones hoped from a naive classifier.
- The values of correct classifications depend highly from the user. There has not been identified a pair of  $\lambda$  and  $S_c$  values which proportionate the highest percentage of correct classification for all users. The discrepancy between the results of the first and the second session is explained by the user's learning process, the same effect was discussed in [19].

## 6 Conclusions

The information inside the pre-assignment sequences improves the classification capability, therefore the Hidden Markov Model technique is useful for the extraction and use of this information in an On-line BCI device.

The scattering of the maximum values, of the correct classifications obtained from the cross-validation tests, shows that the combination of  $\lambda$  and  $S_c$  parameters are highly dependent on the user for all of the cases, for this reason a BCI device based in this kind of algorithm should have a setup stage, that allows to initialise correctly these parameters.

On the other hand, the algorithm behaves better than a naive algorithm, but it is not as good as it should be taking into account the good results obtained during the learning phase. The size of the learning data set is critical in the results obtained during the validation phase. With a bigger learning data set the validation results will improve, because of the minimisation of the overlearning.

In future applications the algorithm presented in this paper will be used as kernel for an on-line classifier embedded in a BCI device. The on-line use of this device will allow to assess how the different kinds of user's feedbacks modify the classification capability.

## References

1. Vidal, J.J.: Toward direct brain-computer communication (1973)
2. Wolpaw, J.R., et al.: Brain-Computer Interface Technology: A Review of the First International Meeting. IEEE Trans. on Reh. Eng. 8(2), 164–171 (2000)

3. Birbaumer, N., et al.: The thought translation device (TTD) for completely paralyzed patients. *IEEE Trans. on Reh. Eng.* 8(2), 190–193 (2000)
4. Wolpaw, J.R., et al.: Brain-Computer interface for communication and control. *Clinical Neurophysiology* 113, 767–791 (2002)
5. Donchin, E., Spencer, K.M., Wijesinghe, R.: The mental prosthesis: assessing the speed of a P300-based brain-computer interface. *IEEE Transactions on Rehabilitation Engineering* 8(2), 174–179 (2000)
6. Kostov, A., Polak, M.: Parallel man-machine training in development of EEG-based cursor control. *IEEE Trans. on Reh. Eng.* 8(2), 203–205 (2000)
7. Pfurtscheller, et al.: Current trends in Graz brain-computer interface (BCI) research. *IEEE Trans. on Reh. Eng.* 8(2), 216–219 (2000)
8. Pfurtscheller, et al.: Brain oscillations control hand orthosis in a tetraplegic. *Neuroscience Letters* 1(292), 211–214 (2000)
9. Pineda, J.A., et al.: Learning to Control Brain Rhythms: Making a Brain-Computer Interface Possible. *IEEE Trans. on Reh. Eng.* 11(2), 181–184 (2003)
10. Ripley, B.: *Pattern Recognition and Neural Networks*, 2nd edn. Cambridge University Press, London (2000)
11. Penny, W.D., et al.: EEG-based communication: A pattern recognition approach. *IEEE Trans. on Reh. Eng.* 8(2), 214–215 (2000)
12. Neuper, C., et al.: Motor Imagery and Direct Brain-Computer Communication. *Proceedings of the IEEE* 89(7), 1123–1134 (2001)
13. Proakis, J.G., Manolakis, D.G.: *Tratamiento digital de señales: [principios, algoritmos y aplicaciones]*. Prentice-Hall, Madrid (1997)
14. Martínez, J.L., Barrientos, A.: The windowing Effect in Cerebral Pattern Classification. In: *An Application to BCI Technology. IASTED Biomedical Engineering BioMED 2006*, pp. 1186–1191 (February 2006)
15. Martínez, J.L., Barrientos, A.: Brain Computer Interface. Comparison of Neural Networks Classifiers. In: *Proceedings of the BIODEVICES International Conference on Biomedical Electronics and Devices*, vol. 1(1), pp. 3–10 (January 2008)
16. Martínez, J.L., Barrientos, A.: Linear Discriminant Analysis on Brain Computer Interface. In: *IEEE International Symposium on Intelligent Signal Processing, Conference Proceedings Book*, pp. 859–864 (October 2007)
17. Rabiner, L.R.: *A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition* (February 1989)
18. Rabiner, L.R., Juang, B.H.: *An Introduction to Hidden Markov Models* (January 1986)
19. Martínez, J.L., Barrientos, A.: Brain Computer Interface. In: *Application of an adaptive bi-stage classifier based on RBF-HMM, Proceedings of the BIODEVICES International Conference on Biomedical Electronics and Devices*, vol. 1(1), pp. 13–20 (January 2010)

# Thermal Characteristics of Microcantilever Biosensors

Mohd. Zahid Ansari and Chongdu Cho

Department of Mechanical Engineering, Inha University, 253 Yonghyon-dong  
Nam-Ku, Incheon, 402751, Republic of Korea  
cdcho@inha.ac.kr

**Abstract.** This study investigates the thermal deflection characteristics of microcantilever biosensors due to bimetallic effects. Thermal deflections are a major source of noise in microcantilever biosensors. Therefore, by characterising the thermal behaviour, the accuracy of the biosensors can be improved significantly. A commercial finite element analysis software ANSYS Multiphysics is used to analyse the deflection characteristics of gold-coated silicon microcantilevers. The cantilever is first subject to temperature increase and then to combined temperature increase and surface stress. The simulation results are compared against analytical and available experimental results. Results show that both thermal stress and surface stress have linear effect on the cantilever deflection and can be added algebraically to determine the absolute deflection produced entirely by the surface stress. Further, we show that by using double-coated silicon microcantilevers thermal deflections in the biosensors can be eliminated.

**Keywords:** Coefficient of thermal expansion (CTE), Thermal strain, Bimetallic effects, Microcantilever biosensor.

## 1 Introduction

Microcantilever based sensors are getting increasingly popular in a variety of physical, chemical, and biological studies. They have been operated in both liquid and gaseous environments. For instance, microcantilever sensors have been successfully used in calorimetric [1], rheometric [2], acoustic [3], infrared [4], pH meter [5], NO<sub>2</sub> gas sensors [6], atrazine pesticide detector [7], DNA hybridization [8], biomarking of myoglobin and kinase proteins [9], detection of biomarker transcripts in human RNA [10], assaying amyloid growth and protein aggregation [11], and DNA hybridization using hydration induced tension in nucleic acid films [12]. Microcantilever sensors are found to be highly sensitive and accurate device. These sensors normally use optical deflection readout techniques to measure the deflections generated upon the change in the surface stress on the functionalized surface of the sensor. By measuring the deflection, the relevant phenomena is monitored and measured.

The overall accuracy of a microcantilever sensor depends on its design and measurement sensitivities. An efficient cantilever design should convert the surface stress induced stimulus into large deflection of the cantilever, whereas, an efficient measurement technique should ensure that the deflections measured are induced

entirely because of the change in the surface stress. An efficient measurement system should be able to isolate and eliminate the noise present in the deflection signals. The design sensitivity of microcantilevers can be improved by changing the cantilever shape [13], profile [14], or both [15]. In addition, by use of a polymer cantilever such as SU8 can also increase the deflection [16-18]. Most of the noise in the deflection is due to thermal [19] and flow [20] induced excitations. Thermal noise arise mostly because of the bimetallic effects generated due to ambient temperature change, whereas, the flow effects are due to the oscillating drag force acting on the cantilever profile. Present work investigates the sources of thermal noise in microcantilever biosensors and proposes measures to reduce or eliminate the noise. The biosensors use the adsorbate-induced surface stress change in measuring and assaying the unknown species present in a media. The adsorption alters the surface stress distribution on the adsorbing surface, which results in cantilever motion. The biosensor system normally comprises gold film coated thin silicon or polymer substrate cantilever structure and a laser-based optical readout scheme. The gold film helps formation of monolayer of receptor molecules on the cantilever surface during functionalization, and also acts as reflecting medium during deflection measurements. The gold film coating nevertheless introduces bimetallic effects in the biosensor. Depending on the thermo-elastic properties of the film and substrate materials and the temperature variation, thermal deflection can exceed the surface stress-induced deflection, leading to large noise in the deflection signal.

In this work we show that knowing the thermo-elastic properties of the film and cantilever material, the analytical and simulation relation can be used to predict the bimetallic effects in a microcantilever biosensor. A commercial finite element analysis (FEA) code ANSYS Multiphysics is used in the simulations. The thermal deflections induced due to the variation in the ambient temperature are calculated using Stoney and beam models, and simulation results. These results are then compared against available experimental result. In the next step, the cantilever is subjected to combined effect of thermal and adsorbate-induced surface stress changes. Finally, the bimetallic effects in SU8 polymer biosensor are discussed.

## 2 Theory and Numerical Analysis

In microcantilever biosensors, the change in the surface stress distribution at the functionalized surface of the cantilevers is used to determine the unknown molecules adsorbed on the surface. Since the induced surface stress strongly depends on the molecular species and its concentration, by measuring the cantilever deflection the attaching species as well as its concentration can be determined. Thus, the basic principle of microcantilever-based biosensor is to detect and measure the surface-stress induced deflections in the cantilever (Fig. 1).


Surface stresses are generated either by the redistribution of the electronic charge at the surface due to the change in the equilibrium positions of the atoms near the surface, or by the adsorption of foreign atoms onto the surface to saturate the dangling bonds [21]. Dangling bonds are of particular interest in biosensors, because during the functionalization of the biosensor, when the complementary bioreceptor molecules are deposited onto the sensing surface of the biosensor, we basically introduce

dangling bonds on the biosensor surface to provide sites for the analyte molecules to bond at. It should be, however, noted that surface stress is different from the normal stress encountered in solids. First, compared to the normal stresses which are distributed along and are supported by the entire cross-sectional area of the body, the surface stresses are usually confined to within few atomic planes of the exposed surface. Second, the unit of surface stress measurement is N/m, whereas, that of normal stress is N/m<sup>2</sup>. For convenience, surface stress in solids can be assumed analogous to surface tension in liquids.

Assuming the film thickness ( $t_f$ ) is infinitesimal compared to the substrate thickness ( $t_s$ ), Stoney expression [22] relating the transverse deflection ( $\Delta z_{ss}$ ) in a microcantilever to adsorbate-induced surface stress change ( $\Delta\sigma_{ss}$ ) can be given as

$$\Delta z_{ss} = \frac{4(1-\nu_s)\Delta\sigma_{ss}}{E_s} \left(\frac{L}{t_s}\right)^2. \quad (1)$$

where  $E_s$  and  $\nu_s$  are elastic modulus and Poisson's ratio of the cantilever material, and  $L$  is cantilever length.



**Fig. 1.** Schematic showing (a) functionalized biosensors cantilever deposited with known receptor molecules and (b) deflection in the cantilever due to adsorption of unknown analyte molecules onto its functionalized surface

As mentioned above, thermal properties of the cantilever structure can itself be a major source of noise in microcantilever biosensors. Bimetallic effects are common phenomena that arise in multilayered structures when subjected to temperature change. The mismatch in their coefficients of thermal expansions (CTEs) results in thermal strain, causing deflections in the structure. The amount and direction of deflection depends on the CTEs of the layers. For instance, if the CTE of the substrate is lower than that of the film and the change in temperature is positive, the thermal deflection will be downwards.

Fig. 2 shows the schematic designs of single- and double-coated microcantilever. In single-coated cantilever, only top surface of the cantilever is coated with a thin film of gold (Fig. 2 (a)). In double-coated, both top and bottom surface of the cantilever is coated with gold (Fig. 2 (b)). In both the cases, the thicknesses of the gold film are same and the film covers the entire surface of the cantilever. It can be deduced from Fig. 2 (b) that thermal deflections in a microcantilever biosensor can be reduced or



eliminated by using a double-coated cantilever. The thermal deflection produced by the expansion of top film layer is annulled by the expansion of the bottom film layer. In other words, the downward deflection created by the bimetallic effect of the top layer is totally balanced by the upward deflection produced by the bottom layer. Both the layers expand simultaneously and produce axial rather transverse movement in the cantilever.


Based on classical beam theory, the transverse deflection in a bilayer beam (Fig. 2(a)) due to bimetallic effect can be given as [23]

$$\Delta z_{th} = \frac{3(\alpha_s - \alpha_f)(t_s + t_f)}{K_1} \left(\frac{L}{t_s}\right)^2 \Delta T \quad (2)$$

where,

$$K_1 = 4 + 6\left(\frac{t_f}{t_s}\right) + 4\left(\frac{t_f}{t_s}\right)^2 + \left(\frac{E_f}{E_s}\right)\left(\frac{t_f}{t_s}\right)^3 + \left(\frac{E_s}{E_f}\right)\left(\frac{t_s}{t_f}\right)$$

and  $\Delta T$  is temperature change,  $\alpha_s$  and  $\alpha_f$  are the CTEs of the substrate and film, respectively. Eq. 2 is commonly used in micro-electro-mechanical systems (MEMS) applications to determine the residual thermal stresses in thin film structures.



**Fig. 2.** Schematic showing bimetallic effects in (a) single-coated and (b) double-coated microcantilever

In this work, we show that by replacing the adsorbate-induced surface stress ( $\Delta\sigma_{ss}$ ) by the thermal-induced film stress ( $\Delta\sigma_{th}$ ), Stoney equation can be used to find the thermal deflection ( $\Delta z_{th}$ ) in the microcantilever biosensor. The expression relating film stress to temperature change can be given as [24]

$$\Delta\sigma_{th} = t_f E_f (\alpha_s - \alpha_f) \Delta T. \quad (3)$$

Thus, the Stoney relation between thermal deflection and thermal-induced film stress can be given as


$$\Delta z_{th} = \frac{4(1-\nu_s) \Delta\sigma_{th} \left(\frac{L}{t_s}\right)^2}{E_s}. \quad (4)$$

## 2.1 Numerical Analysis

A commercial finite element analysis (FEA) software package ANSYS Multiphysics was used for numerical analysis. By taking advantage of the cantilever geometry, 2-D FE model was used. The model was meshed by 8-node coupled-field solid PLANE223 elements. Solution convergence and mesh-size effects were analysed before the final simulations. First a validation study on thermal deflection was performed by comparing the analytical and FEA results against the experimental results of Ramos et al [25]. After validation, thermal characterisation and a coupled thermal-structural analysis involving both thermal- and adsorbtion-induced stresses was performed on both single- and double-coated cantilevers (see Fig. 3). In the first case, the microcantilever was subjected to temperature change alone, whereas, in the other, it was subjected to a constant surface stress as well as temperature variation. The cantilever was subject to an adsorbtion-induced surface stress of 0.05 N/m on its top surface. The surface stress was modelled as a tensile force applied to the top edge of the cantilever [15]. The applied tensile force was  $F = \Delta\sigma_{ss} \times b = 0.05 \times 100 \times 10^{-6} = 5 \times 10^{-6}$  N. Thermal stress was also applied in the same manner by replacing  $\Delta\sigma_{ss}$  by  $\Delta\sigma_{th}$ . The simulation parameters used in the analysis are listed in Table 1.

**Table 1.** Geometric and thermo-elastic properties of the microcantilever

Length of cantilever, $l$ ( $\mu\text{m}$ )	500
Width of cantilever, $b$ ( $\mu\text{m}$ )	100
Thickness of substrate, $t_s$ ( $\mu\text{m}$ )	1
Thickness of film, $t_f$ ( $\mu\text{m}$ )	0.02
Elastic modulus of substrate, $E_s$ (GPa)	130
Elastic modulus of film, $E_f$ (GPa)	78
Poisson's ratio of substrate, $\nu_s$	0.28
Poisson's ratio of film, $\nu_f$	0.44
Coef. of thermal expansion of substrate, $\alpha_s$ ( $1/^\circ\text{C}$ )	$2.6 \times 10^{-6}$
Coef. of thermal expansion of film, $\alpha_f$ ( $1/^\circ\text{C}$ )	$14.2 \times 10^{-6}$
Thermal conductivity of substrate, $k_s$ (W/m $^\circ\text{C}$ )	149
Thermal conductivity of film, $k_f$ (W/m $^\circ\text{C}$ )	315



**Fig. 3.** Schematic showing total deflection produced in (a) single-coated and (b) double-coated microcantilever biosensor due to analyte-receptor binding and temperature variation


### 3 Results and Discussion

Table 2 shows the validation study results comparing the analytical and simulation results against the experimental result in [25]. In the experiment, a silicon cantilever of size  $400 \times 100 \times 1 \mu\text{m}$  was coated with a 20-nm-thick film of gold. And the cantilever was subjected to a temperature variation of  $-12 \text{ }^\circ\text{C}$ . The thermo-elastic material and geometric properties for this validation study are adopted from [25]. The rest of the analysis used the gold-coated microcantilever detailed in Table 1. The analytical results were calculated using beam model (i.e., Eq. 2) and Stoney model (i.e., Eq. 4) and FEA result was from ANSYS. It is obvious in the table that the experimental and the analytical and simulation results have good accord in predicting the thermal deflection, indicating the conformity of analytical and simulation analysis.

**Table 2.** Comparison between experimental, analytical and FEA results

$\Delta T$ ( $^\circ\text{C}$ )	Thermal deflection, $\Delta z_{\text{th}}$ ( $\mu\text{m}$ )			
	Exp. [25]	Stoney [Eq. 4]	Beam [Eq. 2]	FEA
-12	0.56	0.60	0.61	0.62

Fig. 4 shows the effect of temperature variation on thermal characteristics of the gold-coated silicon cantilever shown in Fig. 2 (a). The simulation results for thermal deflections are compared against beam model (Eq. 2) and Stoney model (Eq. 4). As can be seen in the figure, all the three curves show good accord in predicting the thermal deflections. Furthermore, the simulation results corroborate the linear relation between deflection and temperature, which is also suggested by the analytical models.




**Fig. 4.** Thermal deflections in a gold-coated silicon microcantilever

Fig. 5 shows the combined effect of temperature variation and adsorbate-induced surface stress on the deflection characteristics of single- and double-coated cantilevers. The cantilever was subject to a constant surface stress of 0.05 N/m and the ambient temperature was increased. The total deflection ( $\Delta z_{\text{total}}$ ) is sum of thermal deflection and surface stress-induced deflection. The analytical study used two different approaches. In the first case, Stoney + beam models (i.e., Eq. 1 and Eq. 2) were used to calculate the total deflection. In the other case, Stoney models (i.e., Eq. 1 and Eq. 4) were used to calculate the total deflection.

Fig. 5 (a) shows the simulation results for the single-coated microcantilever subject to combined action of film and surface stresses. As can be seen in Fig. 5 (a), at  $\Delta T = 0$  the cantilever shows a deflection of about 0.28  $\mu\text{m}$ , which is induced entirely by the surface stress of 0.05 N/m. In the figure, for  $\Delta T > 3^\circ\text{C}$ , thermal deflections exceed the surface stress-induced deflections. The higher the temperature change, the higher the thermal noise will be. Since both the thermal and surface-stress deflections show linear and additive characteristics, thermal noise can be isolated easily by deducting the thermal deflection from the total deflection. Thus, by knowing the thermo-elastic and geometric properties of the cantilever, thermal deflection can be accurately predicted from analytical relations. It can also be observed that as the temperature increases, the deflection curves show deviation, which suggests the analytical and simulation results are accurate for  $\Delta T < 10^\circ\text{C}$ .

Fig. 5 (b) shows the simulation results for the double-coated microcantilever subject to combined action of film and surface stresses. It is evident in the figure that the proposed cantilever shows negligible thermal deflection (i.e.,  $\Delta z_{\text{th}} \approx 0$ ) when exposed

to temperature increase. And the total deflection is only produced by the surface stress. Since the top and bottom layers expand simultaneously, only axial movement occurs and the cantilever is pulled along the length. Thus, we may conclude that the thermal noise in microcantilever biosensors can be reduced by using a double-coated microcantilever.



**Fig. 5.** Total deflection due to combined action of thermal and surface stresses in (a) single-coated and (b) double-coated silicon microcantilever

Thermal noise problem is more pronounced in liquid medium than in gaseous. In microcantilever biosensor system, the cantilevers array is immersed in a fluidic cell chamber and solution stream containing the analyte molecules is flowed across. The temperature distribution around a microcantilever placed in a fluidic cell can change because of extrinsic and intrinsic sources. The extrinsic sources of temperature change are mostly associated with physical process. As extrinsic source, during purging the chamber and/or injecting the solution samples, the fluid filling the chamber can alter the ambient temperature around the cantilever array. It can induce thermal strain in the cantilever, and eventually produce thermal deflections. Intrinsic sources

arise due chemical processes that occur on the cantilever surface during the reaction between analyte and receptor molecules. If the reaction is exothermic, i.e., heat is released during the reaction, the temperature distribution around the cantilever is severely perturbed, inducing bimetallic effects in the cantilever.

In practice, the surface stress-induced deflections normally range few tens to few hundreds of a nanometre. Fritz et al. [19] reported a 50 nm deflection of the microcantilever due to the bimetallic effect. This was five times the microcantilever deflection due to analyte-receptor binding. And it can also be readily concluded from Fig. 4 that a temperature variation of even 1°C can produce a thermal noise of about 100 nm. Thus, elimination of thermal noise is vital for accurate measurements using microcantilever biosensors and double-coated cantilevers can be efficiently used in this regard.

Bimetallic effects present a far more serious problem in SU8 polymer cantilever biosensors than in silicon. Due to the large CTEs mismatch between the SU8 substrate and the gold film, large thermal deflections are produced. Since the material properties of SU8 are strongly affected by the fabrication process [26], we adopted standard material properties for SU8 from the manufacturer's data sheet [27]. The CTE of SU8 is  $52 \times 10^{-6}$  /°C, whereas, that of gold is  $14.2 \times 10^{-6}$  /°C. Since the mismatch between the CTEs of the layers is the main cause of bimetallic effect in cantilevers, such a large CTE difference will generate considerably high thermal deflection. Further, since the CTE of SU8 substrate is higher than that of film, the cantilever will bend upwards when the temperature is increased.

In addition to the CTE mismatch, the elastic moduli difference between SU8 and gold is also very high. For example, compared to the elastic modulus of about 78 GPa for gold, the modulus of SU8 is about 2GPa. The large difference between the moduli will induce proportionately large strains in the layers. The strain mismatch will generate large shear strain at the gold-SU8 interface, which may eventually lead to delaminating of the film from the substrate if the temperature is raised significantly. The delaminating effect is also present in gold-coated silicon cantilevers to a lesser extent.

If the film thickness is significant but still much less than substrate thickness, Eq. 2 and Eq. 4 can be used for predicting deflections by replacing the substrate modulus ( $E_s$ ) with the effective substrate modulus ( $E_{\text{eff}}$ ) given as [28]

$$E_{\text{eff}} = \frac{E_f^2 a^4 + E_s^2 b^4 + 2E_f E_s ab(2a^2 + 2b^2 + 3ab)}{E_f a + E_s b} \quad (5)$$

where  $a = t_f/t$  and  $b = t_s/t$ , and  $t$  is total thickness of cantilever. To elucidate Eq. 5, consider the geometric and material properties of the gold-coated cantilever analyzed in this study. Though  $E_s$  is 130 GPa, the  $E_{\text{eff}}$  is about 127 GPa.

## 4 Conclusions

Microcantilever biosensors provide a universal, rapid, and highly sensitive mean to study many applications. The measurement accuracy of the biosensor depends on its ability to isolate and eliminate noise from the signal. Since bimetallic effects are a major source of noise in the signal, this study investigated the effect of temperature on the deflection characteristics of the cantilever. The results indicated that thermal

deflections can be determined accurately using analytical and simulation models. By studying the combined effect of thermal- and adsorbate-induced surface stresses on the cantilever deflection, we found that the two stresses act linearly and additively. We found that even for small temperature variation, the thermal deflections can easily exceed the adsorbate-induced deflections, and hence contribute significant noise in the deflection signal. However, by deducting the thermal deflection from the total deflection amount, the exact adsorbate-induced deflection or stress can be determined. We also showed that double-coated silicon microcantilever can be used to completely eliminate the thermal deflection in the biosensor. Finally, the large mismatch in the thermo-elastic properties between gold and SU8 makes SU8 microcantilever biosensors more susceptible to bimetallic effects.

**Acknowledgements.** This study was supported by Inha University.

## References

1. Barnes, J.R., Stephenson, R.J., Woodburn, C.N.: A femtojoule calorimeter using micromechanical sensors. *Rev. Sci. Instrum.* 65, 3793–3798 (1994)
2. Hennemeyer, M., Burghardt, S., Stark, R.W.: Cantilever micro-rheometer for the characterization of sugar solutions. *Sensors* 8, 10–22 (2008)
3. Passian, A., Evans, P.G., Varma, V.K.: Piezoresistive detection of acoustic waves. *Rev. Sci. Instrum.* 74, 1031–1035 (2003)
4. Wachter, A.E.A., Thundat, T., Oden, P.I.: Remote optical detection using microcantilevers. *Rev. Sci. Instrum.* 67, 3434–3439 (1996)
5. Nordstrom, M., Keller, S., Lillemose, M., Johansson, A., Dohn, S., Haefliger, D., Blagoi, G., Havsteen-Jakobsen, M., Boisen, A.: SU-8 cantilevers for bio/chemical sensing; fabrication, characterisation and development of novel read-out methods. *Sensors* 8, 1595–1612 (2008)
6. Qazi, M., Koley, G.: NO<sub>2</sub> detection using microcantilever based potentiometry. *Sensors* 8, 7144–7156 (2008)
7. Suri, A.R., Kaur, J., Gandhi, S., Shekhawat, G.S.: Label-free ultra-sensitive detection of atrazine based on nanomechanics. *Nanotechnol.* 19, 235502–235600 (2008)
8. McKendry, R., Zhang, J., Arntz, Y., Strunz, T., Hegner, M., Lang, H.P., Baller, M.K., Certa, U., Meyer, E., Guntherodt, H.J., Gerber, C.: Multiple label-free biodetection and quantitative DNA-binding assays on a nanomechanical cantilever array. *Proc. Natl. Acad. Sci.* 99, 9783–9788 (2002)
9. Arntz, Y., Seelig, J.D., Lang, H.P., Zhang, J., Hunziker, P., Ramseyer, J.P., Meyer, E., Hegner, M., Gerber, C.: Label-free protein assay based on a nanomechanical cantilever array. *Nanotechnol.* 14, 86–90 (2003)
10. Zhang, J., Lang, H.P., Huber, F., Bietsch, A., Grange, W., Certa, U., McKendry, R., Guntherodt, H.J., Hegner, M., Gerber, C.H.: Rapid and label-free nanomechanical detection of biomarker transcripts in human RNA. *Nature Nanotechnol.* 1, 214–220 (2006)
11. Knowles, T.P.J., Shu, W., Huber, F., Lang, H.P., Gerber, C., Dobson, C.M., Welland, M.E.: Label-free detection of amyloid growth with microcantilever sensors. *Nanotechnol.* 19, 384007 (2008)
12. Mertens, J., Rogero, C., Calleja, M., Ramos, D., Martin-Gago, J.A., Briones, C., Tamayo, J.: Label-free detection of DNA hybridization based on hydration induced tension in nucleic acid films. *Nature Nanotechnol.* 3, 301–307 (2008)

13. Ansari, M.Z., Cho, C.: A study on increasing sensitivity of rectangular microcantilevers used in biosensors. *Sensors* 8, 7530–7544 (2008)
14. Fernando, S., Austin, M., Chaffey, J.: Improved cantilever profiles for elements. *J. Phys. D: Appl. Phys.* 40, 7652–7655 (2007)
15. Ansari, M.Z., Cho, C., Kim, J., Bang, B.: Comparison between deflection and vibration characteristics of rectangular and trapezoidal profile microcantilevers. *Sensors* 9, 2706–2718 (2009)
16. Calleja, M., Nordstrom, M., Alvarez, M., Tamayo, J., Lechuga, L.M., Boisen, A.: Highly sensitive polymer-based cantilever-sensors for DNA detection. *Ultramicroscopy* 105, 215–222 (2005)
17. Ransley, J.H.T., Watari, M., Sukumaran, D., McKendry, R.A., Seshia, A.A.: SU8 biochemical microarrays. *Micro. Engg.* 83, 1621–1625 (2006)
18. Johansson, A., Blagoi, G., Boisen, A.: Polymeric cantilever-based biosensor with integrated readout. *Appl. Phys. Lett.* 89, 173505 (2006)
19. Fritz, J., Baller, M.K., Lang, H.P., Rothuizen, H., Vettiger, P., Meyer, E., Guntherodt, H.J., Gerber, C., Gimzewski, J.K.: Translating biomolecular recognition into nanomechanics. *Science* 288, 316–318 (2000)
20. Khaled, A.-R.A., Vafai, K.: Analysis of oscillatory flow disturbances and thermal characteristics inside fluidic cells due to fluid leakage and wall slip conditions. *J. Biomech.* 37, 721–729 (2004)
21. Ibach, H.: *Physics of Surfaces and Interfaces*. Springer, Heidelberg (2006)
22. Stoney, G.G.: The tension of metallic films deposited by electrolysis. *Proc. Roy. Soc. Lond. A* 82, 172–175 (1909)
23. Finot, B., Passian, A., Thundat, T.: Measurement of mechanical properties of cantilever shaped materials. *Sensors* 8, 3497–3541 (2008)
24. Hsueh, A.H.: Modeling of elastic deformation of multilayers due to residual stresses and external bending. *J. Appl. Phys.* 91, 12 (2002)
25. Ramos, C., Mertens, J., Calleja, M., Tamayo, J.: Study of the origin of bending induced by bimetallic effect on microcantilever. *Sensors* 7, 1757–1765 (2007)
26. Feng, R., Farris, R.J.: Influence of processing conditions on the thermal and mechanical properties of SU8 negative photoresist coatings. *J. Micromech. Microeng.* 13, 80–88 (2003)
27. <http://www.microchem.com/products/pdf/SU-8-table-of-properties.pdf>
28. Yi, J.W., Shih, W.Y., Shih, W.H.: Effect of length, width, and mode on the mass detection sensitivity of piezoelectric unimorph cantilevers. *J. Appl. Phys.* 91, 1680–1686 (2002)



# Prosodic Speech Restoration Device: Glottal Excitation Restoration Using a Multi-resolution Approach

Olaf Schleusing<sup>1</sup>, Rolf Vetter<sup>1</sup>, Philippe Renevey<sup>1</sup>,  
Jean-Marc Vesin<sup>2</sup>, and Valérie Schweizer<sup>3</sup>

<sup>1</sup> Swiss Center for Electronics and Microtechnology, CSEM, Neuchâtel, Switzerland

<sup>2</sup> Swiss Federal Institute of Technology Lausanne, EPFL, Lausanne, Switzerland

<sup>3</sup> University Hospital Lausanne, CHUV, Lausanne, Switzerland

{osg, rve, prv}@csem.ch,

jean-marc.vesin@epfl.ch, valerie.schweizer@chuv.ch

**Abstract.** We describe a novel device for the restoration of authentic characteristics in pathological speech uttered by subjects with laryngeal disorders. The original speech signal is acquired and analyzed by the device and a speech signal with improved, healthy-like features is reconstructed. A concatenation of randomly chosen healthy reference patterns replaces the pathological excitation. In order to restore authentic features, intervals between subsequent reference patterns are obtained through a multi-resolution approach. Short-term pitch variability is reproduced through a statistical variation model. For middle-term pitch variability the correlation between pitch and signal envelope at the middle-term time scale is exploited. Long-term variability is obtained through adaptive wavetable oscillators; a novel, reliable and computationally efficient pitch estimation method. Two authentic features were used to assess performance, namely breathiness and prosody. Preliminary results indicate that breathiness of the restored signal is clearly reduced and prosody related features are improved.

**Keywords:** Speech processing, Speech restoration, Adaptive filters, Multi-resolution, Laryngectomy.

## 1 Introduction

The degree of degradation in pathological voices often engenders a decrease in a patient's speech intelligibility and thereby a severe limitation in his social life and oral interaction [1]. For example, subjects who have undergone laryngectomy suffer from degradation of their natural vocal excitation [2,3,4]. Laryngectomy is the common treatment after diagnosis of larynx cancer in an advanced stage and constitutes the partial or total removal of the larynx. This significantly reduces the patient's ability to produce voiced sounds due to the reduced or missing vocal cord functionality [5,6]. During speech rehabilitation, the patient may learn to use alternative voicing methods, but the result usually is a noisy and intermittently obstructed voice. It lacks power and pitch-variability and typically has an unnaturally low pitch. Alaryngeal voices are usually perceived as not gender-discriminative, breathy and speakers find it difficult

to express prosody. In accordance with the widely accepted source-filter-model in healthy speech processing [7], the vocal cords are essential since they provide an excitation signal with distinctive, periodical energy concentrations. This signal undergoes further spectral shaping due to resonances of oral and nasal cavities as well as the lip radiation function. In contrast, the alaryngeal voice excitation consists of a flawed, distorted excitation signal where the glottal peaks are much less concentrated in the time domain. This results in an unpleasant and unnatural voice with a fluctuating and often intermitted periodicity. In addition, the speaker loses most of its control over pitch variability.

Several advanced voice restoration systems and methods have been presented in the past aiming at the improvement of the quality and intelligibility of the alaryngeal speech. In [8] methods based on linear prediction (LPC) for analysis and synthesis were used to enhance the perceived, subjective voice quality. In [9] modified voice conversion methods combined with formant enhancement were utilized to reduce the pathological speech signal's spectral distortions. In [10] a voice restoration system is described that synthesizes speech from electroglottograph (EGG) pitch information and a jitter reduction model. In [13] a system is presented that restores a pathological speech signal by replacing its pathological excitation with a concatenation of glottal reference patterns randomly chosen from a database extracted from healthy speakers. The intervals between successive healthy glottal patterns are determined by the fundamental frequency ( $f_0$ ) extracted from the original, pathological speech signal. Promising performances have been obtained in terms of reduction of breathiness and increase of the average pitch, but the resulting speech lacks authenticity due to the significantly reduced  $f_0$ -variability in pathological speech.

To improve these deficiencies, we propose a speech restoration device based on a multi-resolution method to increase the variability of the restored pitch. Natural prosody is restored by obtaining the intervals between subsequent glottal waves through a multi-resolution approach. The long-term variability is deduced from the  $f_0$ -trend in the original speech signal. Middle-term variability is restored using the correlation between  $f_0$  and the signal variance in natural speech. Short-term variability is restored using a statistical variation model and the signal envelope. The speech signal is reconstructed subsequently with the enhanced excitation and can be deployed in manifold applications such as voice enhancement systems or interactive support systems for voice rehabilitation and tutoring.

In the next Sections we outline common forms of speech rehabilitation and illustrate the resulting pathological speech characteristics, which motivated the development of our speech restoration system. In Sect. 4 we describe the multi-resolution pitch restoration method used by the proposed device. The results obtained by subjective listening tests are then presented and discussed in Sect. 5.

## 2 Speech Rehabilitation for Alaryngeal Speakers

During laryngectomy, the larynx including the vocal cords and the laryngeal muscles is partially or totally removed [5]. Generally, postlaryngectomy patients may regain means of verbal communication in two ways.


On one hand there exist electro-mechanical devices called electrolarynx that use a membrane to generate an external, synthetic speech excitation when held against the neck. This sets the air volume in the vocal tract into vibration and the patient can articulate in a natural manner. Unfortunately, the voice quality achievable with electrolarynx devices is low since there is no intuitive control over the fundamental frequency and voice quality parameters such as breathiness. The resulting speech sounds very monotonous and robot-like. Advantages of this method are its simplicity and short learning phase. The patient does not need to undergo additional surgery and can start communicating verbally almost immediately.

On the other hand, postlaryngectomy speakers may learn to use other tissues to substitute the functionality of the vocal cords. In tracheo-esophageal speech, the speaker utilizes pulmonary air to produce voicing with the substitutional tissues. The speaker may retain intuitive control over aspects of its voice, but only to a very limited extent. The expression of prosody such as variations of the fundamental frequency or modulation of the voice quality is greatly reduced compared with healthy speakers. In addition, the aptitude of the remaining tissue to produce a rich, harmonic sound is very limited and its physical properties vary greatly among speakers and differ significantly from those of the vocal cords.

### 3 Characteristics of Pathological Speech

In healthy speech production, subglottal air pressure leads to a sudden, non-symmetric opening of the vocal cords and a release of this pressure. Subsequently, the musculature surrounding the glottis and the accelerated air flowing through the glottis force a closure of the vocal cords and the process starts all over. Varieties in glottis shape and size amongst humans lead to speaker-specific patterns for the opening and closing process as well as to the introduction of jitter in the period between subsequent glottal cycles. These effects amongst others lead to speaker specific voice characteristics.

In comparison to healthy voices, these aerodynamic-myoplastic processes are not very well studied in voiced, alaryngeal speech production [5]. Alaryngeal voice characteristics have been found to differ remarkably from that of healthy voices. Among subjects the position and shape of the neoglottis vary significantly [8]. Often incomplete glottal closure can be observed. Furthermore, the flexibility and controllability of the neoglottis lacks greatly when compared with a healthy glottis, especially due to the absence of the laryngeal musculature. The high mass of the neoglottis and low resistance to mucus aggregation influence the absolute value and stability of the fundamental frequency in a disadvantageous manner. The alaryngeal oscillator tends to break down intermittently [14]. Eventually, the resulting voice has an unnaturally low and unstable pitch and often is found to have a hoarse, croaky and breathy voice quality [15]. Figure 1 depicts segments of LPC-residues of a laryngeal and an alaryngeal voiced excitation signal. This figure highlights the alteration of the produced harmonic excitation due to the changed physiologic conditions. The glottal wave patterns in the excitation of the healthy speaker are highly concentrated in the time domain, whereas the excitation of the alaryngeal speaker appears merely as a modulated noise signal.



**Fig. 1.** Segments of a laryngeal (a) and an alaryngeal (b) voice excitation signal obtained by LPC inverse filtering and their respective envelopes

## 4 Multi-resolution Voice Restoration

### 4.1 Method

A block scheme of the method implemented in the device is depicted in Figure 2. The articulation information and the voice excitation are separated in a primer LPC-based signal analysis. The obtained excitation signal is then divided into voiced and unvoiced segments, since we are only interested in restoring the voiced excitation signal. The voiced excitation segments then are replaced by concatenating healthy glottal reference patterns. These reference patterns were previously extracted from laryngeal speakers and are randomly chosen from a lookup-table. The intervals between successive reference patterns determine the fundamental frequency of the reconstructed speech signal. The fundamental frequency in pathological speech is degraded in terms of variability and stability and thus insufficient for a successful restoration of an authentic speech signal. To increase authenticity, the intervals between subsequent glottal waves are obtained through a multi-resolution approach on three different time scales:

- Long-term pitch variability  $f_{0,LT}$  is restored with instantaneous frequency estimations obtained from the alaryngeal voice excitation with adaptive wavetable oscillators (AWO).
- Middle-term pitch variability  $f_{0,MT}$  is strongly related to prosody and is reconstructed by exploiting the correlation between pitch variations and instantaneous signal energy [16].
- Short-term pitch variability  $f_{0,ST}$  is introduced through a statistical model to mimic the presence of pitch jitter as in healthy speech.

Finally, the improved excitation signal is recombined with the unmodified unvoiced speech segments and the estimated articulation information.




Fig. 2. Block diagram of the multi-resolution pitch restoration method

## 4.2 Long-Term Pitch Estimation

The objective of the long-term pitch estimation is to grasp what remains of the pitch information in the alaryngeal speech. The selection of method for the extraction of the fundamental frequency of a specific signal depends on different characteristics of the signal itself:

- The signal’s nature in terms of time-frequency distribution
- The amount and characteristics of additional harmful background noise
- The affordable computational complexity.

In general, pitch estimation methods can be classified into event detection methods and short-term averaging methods. Event detection methods such as for example zero-crossing [17] or threshold-guided maxima localization [17] are computationally inexpensive and yield high performance for well shaped signals in low-noise environments. Signals with higher harmonic complexity or increasing noise level require more advanced methods such as the matched filter method [18] or auto-correlation method [19]. They are based on short-term averaging and computationally more expensive. More advanced methods with yet increased computational complexity, decompose the signal into its Eigenspace components [20]. Conjoint approaches [21] combine three different methods, namely in time, frequency and cepstrum domain.

For the device presented in this paper, the focus is clearly on the efficient utilization of the given computational resources. We propose to use a new pitch estimation method taking into account the demand for low computational load and for the pertinence and simplicity of fixed-point real-time implementation. The method is based on adaptive wavetable oscillators, a method recently published in [22]. An evaluation of the method comparing it with other state-of-the-art methods for fundamental frequency estimation was presented in [23].

AWOs constitute a time-frequency method combining wavetables and adaptive oscillators. Wavetables generate periodic output signals by cyclic indexing of a lookup

table that stores a single period of the waveform. Adaptive oscillators synchronize their output to both frequency and phase of the input signal. The indexing parameters of the AWO are determined by optimizing a well defined cost function such that the error between the wavetable output and an incoming, periodic signal is minimized.

The first step in the design of an AWO requires the selection of an appropriate pattern. This pattern should represent a high similarity with the signal pattern to be extracted and is stored in a wavetable as numerical, digital information. With respect to the above consideration, we use the energy distribution of the glottic excitation envelope as input (see Figure 1(b)) and a Gaussian function as wavetable pattern. As one can observe, the envelope of energy during glottal patterns of the excitation signal has a high similarity with a Gaussian shape. A Gaussian function is easily controllable with only a few parameters such as a time index  $n$ , a phase offset in samples  $\beta$  and a temporal width  $\sigma$ :

$$w(n) = e^{-\frac{1}{2} \frac{(n-\beta)^2}{\sigma^2}} \quad (1)$$

Cyclic sampling is used to generate a periodical reference signal  $v(n)$ :

$$v(n) = w(k(n) \bmod N) \quad (2)$$

where  $k(n)$  is the cyclic sampling index

$$k(n) = (k(n-1) + \alpha) \bmod N \quad (3)$$

$k(0)$  is initialized to 0,  $x \bmod N$  is the remainder operator, and  $\alpha$  is the sampling step determining the sub-sampling rate of the wavetable pattern. The control parameters of the periodic output of Equation 1 are adaptively updated by using well understood gradient techniques [24]. The output of the wavetable oscillator thus is locked to the input signal and the parameter  $\alpha$  is related to the fundamental frequency of the alaryngeal excitation signal by  $\alpha/(NT_s)$ , with  $T_s$  being the sampling period. The phase depends on the offset  $\beta$  of the sampling index. The adaption of the indexing parameters is achieved by minimizing a well defined cost function that gauges the error between the wavetable output and an incoming, periodic signal:

$$J(n) = s(n)v(n) \quad (4)$$

where  $s(n)$  is the envelope of the extracted speech excitation signal.


Assuming that the phase and frequency of the input signal vary slowly over time one can follow these changes by moving the argument of the cost function slowly into the direction of the derivative:

$$\alpha(n+1) = \alpha(n) + \mu_\alpha \left. \frac{\partial J}{\partial \alpha} \right|_{\alpha=\alpha(n)} \quad (5)$$

and

$$\beta(n+1) = \beta(n) + \mu_\beta \left. \frac{\partial J}{\partial \beta} \right|_{\beta=\beta(n)} \quad (6)$$

It can easily be seen that the gradients  $\frac{\partial J}{\partial \alpha}$  and  $\frac{\partial J}{\partial \beta}$  are similar up to a constant.



**Fig. 3.** Performance of the pitch estimation methods for the healthy speech signal with different levels of AWGN

$$\frac{\partial J}{\partial x} = \frac{\partial J}{\partial w} \frac{\partial w}{\partial x} \quad (7)$$

Indeed, they include the partial derivative of  $w$ , which is typically stored in a wavetable of  $N$  samples to minimize the computational load. The learning gains  $\mu_\alpha$  and  $\mu_\beta$  should be chosen such that the oscillator can change rapidly enough to follow changes in the fundamental frequency and minimize noise influences. Since the adaptation of the frequency is much more sensible than that of the phase,  $\mu_\alpha$  should be much smaller than  $\mu_\beta$ .

In [23] the AWO method was compared to several other methods such as the correlation method and the matched filter method. A quantitative validation was performed on healthy, phonetically equilibrated French speech sentences with additive white Gaussian noise at SNRs ranging from  $-10dB$  to  $20dB$ . To obtain an objective performance assessment we evaluated first the most likely fundamental frequency during voiced segments as the median value of the estimations from the three different methods at each sample. Quantitative performances were then assessed as the Mean Relative Absolute Error (MRAE) between an estimation result from a specific estimator and the most likely value. The results displayed in Fig. 3 confirm the results on healthy speech signals. The AWO and XCorr perform better than the MF method, particularly at low levels of SNR.

For pathological voices, very likely instantaneous fundamental frequencies to serve as a reference are not available. Thus, a subjective validation by listeners was carried out, based on the Mean Opinion Score (MOS) [11] of a signal generated by the presented method.

The results of the subjective evaluation displayed in Table 1 highlight that XCorr outperforms AWO and MF in the mean with respect to listener specific subjective evaluation. However, an analysis of the variance of the XCorr and AWO methods yields a p-value of 0.61, which suggests that from a statistical point of view this result is not significant [12]. The performance of the matched filter method drops due to insufficient reliable support points for its restoration, which made it impossible to follow the changes in the fundamental frequency.

The above results have shown that the proposed AWO method performs similar to the correlation method under a variety of experimental conditions. In fact, the correlation method is outperformed by the AWO method when applied to healthy voices and when significant amounts of additive background noise is present. When applied to

**Table 1.** Mean Opinion Score (mean  $\pm$  standard deviation) of seven listeners assessing the performance of specific pitch estimation methods as a preprocessing unit to a voice restoration method for pathological voices. Applied MOS-scale: bad-1, poor-2, fair-3, good-4 and excellent-5.

Method	XCorr	AWO	MF
MOS	3.7 $\pm$ 0.49	3.6 $\pm$ 0.53	1.7 $\pm$ 0.49

pathological voices, the performances of AWO and the correlation method are nearly the same. Both methods reconstruct the pitch in the speech analysis and restoration system to a quality, where it is rated by listeners between fair and good. However, as the computational load of the AWO method is much lower than that of the correlation method, AWO appears to be a very promising pitch estimation method for real-time fixed point implementation on embedded platforms.


### 4.3 Middle-Term Pitch Restoration

Middle-term pitch variability  $f_{0,MT}$  is restored by exploiting the correlation between pitch and the signal envelope at this time scale. It has been shown that prosody is not only strongly related to variations in pitch, but also to variations in the envelope of the speech signal [16]. Figure 4 shows a segment of a healthy speech signal and the  $f_0$  variations, which are directly related to variations of the signal envelope. In healthy speakers, the  $f_0$ -variations may span several octaves, but pathological speakers have lost much of the ability to vary the fundamental frequency. Nevertheless, modulations of the intensity of the speech signal may allow a reconstruction of the  $f_0$ -variability due to their tight correlation. The key point of the presented method is to infer variations in pitch from variations in the signal envelope of the alaryngeal speech signal. Therefore, the segmental signal envelope is bandpass-filtered (2 – 8 Hz) and then used to restore the middle-term pitch variability. This will give the pathological speaker a means to intuitively manipulate the pitch of the restored speech signal by manipulating the intensity of his speech.

### 4.4 Short-Term Pitch Restoration

An important characteristic of natural voices are small imperfections, such as involuntary variations in pitch. The human perception expects this short-term variability in speech signals and the lack of them results in an unnatural, unpleasant and buzz-like sound. In the presented device, short-term pitch variability  $f_{0,ST}$  is induced through high-pass-filtered ( $f_c = 8$  Hz) and weighted additive white Gaussian noise (AWGN). The weighting of this AWGN is determined through a signal-envelope-dependent, non-linear weighting inspired by recent findings in healthy subjects [25]. In healthy voices, jitter in  $f_0$  was found to be constantly low during voicing with sound pressure levels of 70 – 75 dB and above. At lower intensity levels though, jitter was found to steadily and sharply increase with falling sound pressure levels. The produced pitch is increasingly





**Fig. 4.** Correlation between the speech signal envelope (middle panel) and pitch (bottom panel) in a healthy speech signal (top panel)

unsteady with decreasing intensity of the produced voice. We model this non-linear behavior with a piecewise linear function.

$$f_{PL}(s_{env}(n)) = \begin{cases} 0.1 & \text{if } s_{env}(n) \geq \gamma \\ 0.1 + 0.9 \frac{\gamma - s_{env}(n)}{\gamma - \kappa} & \text{if } \gamma > s_{env}(n) > \kappa \\ 1 & \text{if } \kappa \geq s_{env}(n) \end{cases} \quad (8)$$

where  $s_{env}(n)$  is the logarithm of the averaged instantaneous signal envelope normalized with respect to the given acoustical configuration,  $\gamma$  and  $\kappa$  have been adjusted with respect to subjective listening tests.

## 5 Results


An evaluation was performed to assess the successful restoration of authentic characteristics from pathological speech to a higher quality. A sustained sound of a vowel  $a$  of a pathological, male speaker with varying pitch was recorded at a sampling rate of  $8kHz$  and 16 bits quantization. The speech signal was restored using the method as described in Sect. 4.1, implemented in the Matlab programming language [26]. Twelve amateur listeners quantified the performance in terms of prosody and breathiness using a mean opinion score (MOS) by listening to the restored speech signals using ordinary headphones. The relative contributions of short-term, middle-term and long-term pitch variabilities to the improved speech quality were assessed using three different system configurations. The pitch was restored from:

- LT: long-term pitch variability
- LT-MT: long-term and middle-term variability
- MR: multi-resolution approach


The results displayed in Table 2 show that the proposed restoration approach improves performance with respect to both criteria. The contribution of the  $f_0$ -variability restored at the middle-term scale appears to be most significant (1.1 points compared

**Table 2.** Mean Opinion Score (mean  $\pm$  standard deviation) of 12 listeners that assessed the quality of 3 different restored voices. Applied MOS-scale: highly improved-1, no improvement-3, highly degraded-5.

Method	Improved Feature	
	Prosody	Breathiness
LT	2.9 $\pm$ 0.7	1.7 $\pm$ 0.8
LT+MT	1.9 $\pm$ 0.7	1.6 $\pm$ 0.5
MR	2.0 $\pm$ 1.2	2.0 $\pm$ 1.2




(a)



(b)

**Fig. 5.** Lower frequencies of a spectrogram of a sustained vowel *a* of an alaryngeal speaker (top panel) and restored with the LT+MT system configuration (bottom panel). The reduced pitch fluctuations and increased pitch variability in the restored speech yielded a perception of improved prosody.

to 0.1 points long-term variability alone). This seems to emphasize our assumption that additional pitch variability at the middle-term scale (Fig. 5) can contribute to the restoration of prosody. The contribution of the MR approach yields no significant improvement to the perceived prosody. The high amount of standard deviation and the relatively small amount of listeners prohibits to draw general conclusions. Nevertheless, a positive trend can be recognized. Regarding the breathiness of the restored voice, a clear improvement (1.0 to 1.4 points) in all voices can be observed. (Figure 6) illustrates the effect of the low SNR in alaryngeal voices, where high frequency harmonics are covered by the noise and their lack contributes to the perception of a less harmonic, breathy voice. For voices restored with the MR approach, the additionally induced short-term variability seems to imply a degradation in terms of increased breathiness compared to the result of the LT+MT system configuration. This could be due to the fact that short-term variability



**Fig. 6.** Spectra of a sustained vowel *a* performed by an alaryngeal speaker (elevated by 20 dB for illustrative purposes) and restored with the proposed MR system. The structure of the higher frequency harmonics is restored in reconstructed speech signal due to the restoration of the laryngeal speech excitation leading to a reduction of the perceived breathiness.

is related to the jitter of speech. Indeed, jitter may be perceived as a desired, authentic feature at a very low intensity level but becomes certainly harmful over a given threshold. This threshold depends on the subject's idiosyncrasies and may be adjusted to the alaryngeal speakers's desire. We suggest that a more carefully designed non-linear model for the short-time variability contribution or a spectrally shaped noise instead of the AWGN may reduce this undesired effect of the short-time pitch variability.

## 6 Conclusions

We presented a method for the restoration of authentic features in pathological voices. The implementation of this method on an embedded device can be regarded as an attractive alternative to currently used electro-larynx devices due to its hands-free mode of operation and superior acoustic quality. The Adaptive Wavetable Oscillator has evolved as a sound method for the estimation of the long-term pitch variations in pathological speech and stands out due to its low computational complexity. The Multi-Resolution approach for the restoration of the  $f_0$ -variability at different time scales improved the prosody and breathiness of pathological voices. Furthermore, the possibility to influence prosodic characteristics of the voice such as the  $f_0$  in an intuitive manner can be regarded as a clear advantage and motivate more alaryngeal speakers to learn or improve tracheo-esophageal speech.

## References

1. Weinberg, B.: Acoustical properties of esophageal and tracheoesophageal speech. In: Keith, R.L., Darley, F.L. (eds.) *Laryngectomy Rehabilitation*, pp. 113–127. College-Hill Press, San Diego (1986)
2. Williams, S.E., Barber Watson, J.: Speaking proficiency variations according to method of alaryngeal voicing. *Laryngoscope* 97(6), 737–739 (1987)
3. Most, T., Tobin, Y., Mimran, R.C.: Acoustic and Perceptual Characteristics of Esophageal and Tracheoesophageal Speech Production. *J. of Communication Disorders* 33(2), 165–180 (2000)

4. Moerman, M., Pieters, G., Martens, J.P., van der Borgt, M.J., Dejonckere, P.: Objective Evaluation of Quality of Substitution Voices. *Eur. Arch. Otorhinolaryngol.* 261, 541–547 (2004)
5. van As, C.J.: Tracheoesophageal Speech: A multidimensional assessment of voice quality. PhD Thesis, University of Amsterdam (2001)
6. Pindzola, R.H., Cain, B.H.: Acceptability Ratings of Tracheoesophageal Speech. *Laryngoscope* 98(4), 394–397 (1988)
7. Fant, G.: The Source Filter Concept in Voice Production. *Speech Research Summary Report STL-QPSR 22*, pp. 21–37 (1981)
8. Qi, Y., Weinberg, B., Bi, N.: Enhancement of Female Esophageal and Tracheoesophageal Speech. *J. Acoust. Soc. of America* 98(5 Pt 1), 2461–2465 (1995)
9. Bi, N., Qi, Y.: Application of Speech Conversion to Alaryngeal Speech Enhancement. *IEEE Transactions on Speech and Audio Processing* 5(2), 97–105 (1997)
10. del Pozo, A., Young, S.: Continuous Tracheoesophageal Speech Repair. In: *Proc. of EUSIPCO* (2006)
11. Virag, N.: Speech Enhancement based on Masking Properties of the Human Auditory System. PhD thesis, Ecole Polytechnique Federale de Lausanne (1996)
12. Papoulis, A.: *Probability and Statistics*. Prentice-Hall, Englewood Cliffs (1989)
13. Vetter, R., Cornuz, J., Vuadens, P., Sola, I.C.J., Renevey, P.: Method and System for Converting Voice. *European Patent (EP1710788)* (2006)
14. Kasuya, H., Ogawa, S., Kikuchi, Y., Ebihara, S.: An Acoustic Analysis of Pathological Voice and its Application to the Evaluation of Laryngeal Pathology. *Speech Communication* 5(2), 171–181 (1986)
15. Verma, A., Kumar, A.: Introducing Roughness in Individuality Transformation through Jitter Modelling and Modification. In: *Proc. of ICASSP*, vol. (1), pp. 5–8 (2005)
16. Rosenberg, A., Hirschberg, J.: On the Correlation Between Energy and Pitch Accent in Read English Speech. In: *Proc. of Interspeech 1294-Mon2A3O.2* (2006)
17. Gerhard, D.: Pitch Extraction and Fundamental Frequency: History and Current Techniques. *Tech. Report*, University of Regina, CA (2003)
18. Turin, G.L.: An Introduction to Matched Filters. *IRE Transactions on Information Theory* 6(3), 311–329 (1960)
19. Un, C., Yang, S.-C.: A Pitch Extraction Algorithm Based on LPC Inverse Filtering and AMDF. *IEEE Trans. ASSP* 25(6), 565–572 (1977)
20. Murakami, T., Ishida, Y.: Fundamental Frequency Estimation of Speech Signals using MUSIC Algorithm. *Acoust. Sci. Technol.* 22(4), 293–297 (2001)
21. Mitev, P., Hadjitodorov, S.: Fundamental Frequency Estimation of Voice of Patients with Laryngeal Disorders. *Information Sciences* 156(1-2), 3–29 (2003)
22. Arora, R., Sethares, W.A.: Adaptive Wavetable Oscillators. *IEEE Trans. on Signal Processing* 55(9), 4382–4392 (2007)
23. Schleusing, O., Vetter, R., Renevey, P., Krauss, J., Reale, F.N., Schweizer, V., Vesin, J.-M.: Restoration of Authentic Features in Tracheoesophageal Speech by a Multi-Resolution Approach. In: *Proc. of SPPRA*, pp. 643-042, Innsbruck (2009)
24. Haykin, S.: *Adaptive Filter Theory*. Prentice-Hall, Englewood Cliffs (2001)
25. Brockmann, M., Storck, C., Carding, P.N., Drinnan, M.J.: Voice Loudness and Gender Effects on Jitter and Shimmer in Healthy Adults. *J. of Speech, Language and Hearing Research* volume (51), 1152–1160 (2008)
26. The Mathworks: *Matlab 2006b* (2006)

# An Investigation on ALD Thin Film Evanescent Waveguide Sensor for Biomedical Application

Agung Purniawan<sup>1</sup>, Paddy French<sup>1</sup>, Gregory Pandraud<sup>2</sup>, and Pasqualina M. Sarro<sup>2</sup>

<sup>1</sup> Electronic Instrumentation Laboratory (EI)-DIMES, TU Delft  
Mekelweg 4, Delft, The Netherlands

<sup>2</sup> Electronic Components, Technology and Materials (ECTM)-DIMES, TU Delft  
Feldmannweg 17, Delft, The Netherlands

**Abstract.** In this paper, we investigate the use of  $\text{Al}_2\text{O}_3$  and  $\text{TiO}_2$  deposited by Atomic Layer Deposition (ALD) as evanescent waveguide sensors. These sensors will be employed to detect bacteria concentration in drain fluid in post anastomosis surgery. Surface roughness, conformality, and homogeneity of the sensor material are very important factors to obtain high sensitive sensor. ALD fulfill these requirements. Surface roughness before and after fabrication are investigated using AFM. As we aim at freestanding structure the buckling of freestanding ALD films is studied. Finally we build an optical characterization set up and measured the propagation loss of  $\text{Al}_2\text{O}_3$  and  $\text{TiO}_2$  waveguides at 1.3  $\mu\text{m}$ . The results show that ALD thin films can be used as waveguide material to obtain very high sensitive sensors.

**Keywords:** Evanescent waveguide, Atomic layer deposition, Propagation loss.

## 1 Introduction

In biomedical engineering, modern health care involves interdisciplinary fields such as sensors, electrical engineering, material science, signal processing, information technology etc. to get more informations from patients. In this research, we focus on developing biomedical sensor especially evanescent waveguide sensor for anastomosis leakage detection. If the leakage occurs in the colon after surgery, bacteria could enter the abdominal cavity and cause severe infections [1]. The anastomosis leakage causes multi organ failure, second complication and even mortality. As comparison, the mortality cases after anastomosis are 2.6% and 18.6% without and with leakage, respectively. In addition, the multi organ failure cases are 1.1% and 15.9% without and with leakage, respectively [2]. The detection of leakage must be done early so that further medical action can be taken soon. Unfortunately conventional methods take up to seven days to detect the leakage.

Atomic layer deposition is a technique to deposit thin atomic films layer by layer. Therefore the resulted thin film has a roughness below 1 nm, good electrical properties, good bio compatibility, good thickness control [3], and conformal [4]. By taking the advantages of the technique, thin film ALD is successfully used in Microelectronic Mechanical System (MEMs) and sensors [4-6], organic transistor [3, 7], etc.

In the following we present the design, the fabrication and the testing of high sensitive evanescent waveguide sensors (EWS) for biomedical applications especially for early detection of bacteria's in the case of anastomosis. We investigated specifically  $\text{Al}_2\text{O}_3$  and  $\text{TiO}_2$  deposited by ALD.

## 2 Fabrication

The expected structure of the waveguide sensor we want to use, is a freestanding structure. These structures are proven to give the best achievable sensitivity as the measuring is on both side of the waveguide. However making flat and long structures is not always possible. Therefore studying the buckling of thin film that composes the waveguide has to be done to avoid that too much light scatters during propagation. The buckling is related to the properties, the dimension and the stress of the material.

### 2.1 Buckling of Thin Film

The ALD thin film stress was measured by laser interferometer stress measurement. The equipment uses a laser interferometer to measure the curvature of the wafer which is then used to calculate the stress in the thin film. The measured stress in the  $\text{Al}_2\text{O}_3$  and  $\text{TiO}_2$  was found to be 347 MPa and 282 MPa, respectively. The critical stress is defined as the maximum stress that can be held by the structure without buckling. The stress is influenced by the material properties and the dimensions of the structure as follow [8]:


$$\sigma_{cr} = \frac{k \cdot \pi^2 \cdot E}{12(1-\nu^2)} \frac{h^2}{b^2} \quad (1)$$

$$h < \sqrt{\frac{12(1-\nu^2) \cdot b^2 \cdot \sigma}{k \cdot \pi^2 \cdot E}} \quad (2)$$

The coefficient  $k$  is a constant that depends on the dimensions (based on ref.[3], if ratio between length/width ( $a/b$ ) = 20, then  $k = 4$ ). By taking a poisson ratio ( $\nu$ ) = 0.3, a waveguide thickness ( $h$ ) = 100 nm and width of waveguide ( $b$ ) = 5  $\mu\text{m}$  the critical stress is 53.5 MPa. Since the stress measured is more than critical stress ( $\sigma > \sigma_{cr}$ ), the buckling of ALD thin film will be present in the structure. By substituting the measured stress of thin film to the critical stress in Eq. (1) we can calculate the maximum buckling of the thin film (Eq. (2)). For an ALD film, the buckling of the thin film is found to be less than 80.5 nm. As the length of waveguide was considered to be 100  $\mu\text{m}$  and assuming that the buckling happens in the middle of the structure, it can be estimated as an angle of  $2.8 \cdot 10^{-5}$  degree. Since this is very small, we can conclude that the buckling does not affect the propagation of the light through the structure.


## 2.2 Front Side Fabrication

In a waveguide configuration,  $\text{SiO}_2$  and  $\text{SiN}$  are used as interlayer between thin film ALD and Si-wafer due to high refractive index of silicon[9]. To fabricate waveguides we use a lithographic process and plasma etching to pattern a rib (Fig. 1). However plasma etching affects the surface roughness and high power will create high surface roughness.  $\text{Al}_2\text{O}_3$  and  $\text{TiO}_2$  are often seen as hard materials to etch. However in this work we used plasma etching with low power (20 W) to reduce to have surface roughness. The cross section of waveguide structure is rib waveguide as shown in Fig.1.



**Fig. 1.** Cross section of waveguide



By using low power (20 watt), the roughness is slightly increased after etching from  $0.4 \text{ nm}$  to  $4 \text{ nm}$ . For the front side patterning we used RIE to achieve a rib depth of  $50 \text{ nm}$  ( $\text{CF}_4 = 50 \text{ sccm}$ ,  $\text{CH}_3 = 25 \text{ sccm}$ ,  $\text{H}_2 = 40 \text{ sccm}$ , power = 20 watt), etching time 10 minute. Fig. 2 shows top view of the fabricated waveguide.



**Fig. 2.** Top view of waveguide (a)  $\text{Al}_2\text{O}_3$  and (b)  $\text{TiO}_2$

## 3 Propagation Losses

To obtain the optical properties of the waveguide which resulted from the fabrication process, we tested them optically. The waveguides tested are  $\text{TiO}_2$  and  $\text{Al}_2\text{O}_3$  rib waveguides. Propagation and coupling losses were investigated using optical measurement set up, schematically depicted in Fig. 3.



**Fig. 3.** Setup and schematic of optical measurement system

To measure the losses we used a LED as light source and a photodiode to convert intensity into voltage. Losses can be obtained by recording the output signal and using [10]:

$$dB = 10 \log \frac{P_{out}}{P_{in}} \tag{3}$$

In these measurements, various lengths of specimens were used to obtain the transmission of every length. Then the readout of transmission is converted into losses and then plotted. The method is called the cut back methods. By using this technique the propagation loss can be estimated. Fig. 4 shows the transmission of specimens of various lengths.

The propagation loss is defined as the slope of the linear regression ( $dy/dx$ ) and the coupling loss are the loss value when the length is zero. From Fig. 3 the propagation



losses of  $\text{TiO}_2$  and  $\text{Al}_2\text{O}_3$  are 6.8dB/cm and 3.7dB/cm, respectively. The coupling losses of  $\text{TiO}_2$  and  $\text{Al}_2\text{O}_3$  are 19 dB and 17.4 dB, respectively.




Fig. 4. Measured losses of waveguide

## 4 Free Standing Structure

In this research we want to build freestanding waveguide to obtain highly sensitive sensors.  $\text{SiO}_2$  and  $\text{SiN}$  nitride were investigated as etch stop layer during the releasing by etching in TMAOH and KO. The structures dimensions and the properties of the waveguides are simulated by EIM method.

### 4.1 Free Standing Processing

To obtain free standing waveguides Si is etched from backside until a stopping layer. Selective etching is very important to fabricate nanolayered free standing structures. The aimed freestanding structure is as shown in Fig. 5.






Fig. 5. Schematic of evanescent waveguide sensor of  $\text{TiO}_2$  waveguide

For backside etching two materials are used to protect the backside of wafer:  $\text{SiO}_2$  for TMAOH etching and  $\text{SiN}$  for KOH etching. Firstly  $2\ \mu\text{m}$   $\text{SiO}_2$  was employed as stopping layer when using TMAOH wet etching due to low etching rate of  $\text{SiO}_2$  in TMAOH compare with Si.  $\text{SiO}_2$  is obtained from thermal oxidation at  $800\text{-}1100^\circ\text{C}$ . However  $\text{SiO}_2$  thermal oxidation has a high stress and consequently the  $\text{Al}_2\text{O}_3$  or  $\text{TiO}_2$  deposited on  $\text{SiO}_2$  also have high stress. Experiment results show that the membranes break when etching until a  $\text{SiO}_2$  stop layer. Fig. 6 shows the broken membrane and the high stress in the membrane.



**Fig. 6.**  $\text{SiO}_2$  membranes with  $\text{Al}_2\text{O}_3$  waveguide resulted by TMAOH etching

Based on these results we concluded that thermal oxide can not be used with ALD thin film on top. We then used  $\text{SiN}$  as stop layer using this time KOH. No stress appears on membranes and conformal etching can be seen in Fig. 7.




**Fig. 7.**  $\text{SiN}$  membranes with  $\text{Al}_2\text{O}_3$  thin film resulted by KOH etching

The results presented in Fig. 7 (KOH etching with  $\text{SiN}$  as stop layer) prove that this method can be used to make free standing ALD waveguides.

## 4.2 New Structure Definition

We have shown that we need to use  $\text{SiN}$  as a etch stop layer. Unfortunately only  $\text{TiO}_2$  ( $n > 2.4$ ) here satisfy waveguiding conditions. Further to prevent leakage to the

substrate. The film has to be quite thick around 1 micron. In contrary in the free standing region the thickness has to be reduced to obtain evanescent effect on both sides (front and back side). Fig. 8 shows the simulation of waveguide SiN/TiO<sub>2</sub> in free standing region using effective index method [11].



**Fig. 8.** Vertical TM field Profile of Freestanding structure of 250 nm TiO<sub>2</sub> on SiN with the thickness of SiN being (a) 1  $\mu\text{m}$  (b) 500 nm (c) 250 nm and (d) 100 nm

SiN thickness influences on the evanescent tail are obviously shown in Fig. 8. The figure shows TM field profiles of waveguide with difference SiN thickness. By reducing SiN thickness in freestanding region, it can be found an evanescence on both surfaces. On the other hand the fabrication process will become difficult when reducing SiN thickness to the optimum thickness. Based on these simulations a structure with 250nm TiO<sub>2</sub> on 250 nm SiN has evanescent tails both sides (Fig. 8(c)). The stack will be in our future investigations used as freestanding waveguide. In case of a waveguide consisting of 250 nm SiN and 250 nm TiO<sub>2</sub> the sensitivity is found to be 0.3.

## 5 Conclusions

We can conclude that buckling of ALD film when freestanding is very low ( $2.8 \cdot 10^{-5}$  degree) which it is not significant to affect the propagation of light through the waveguide. TiO<sub>2</sub> and Al<sub>2</sub>O<sub>3</sub> films have acceptable losses and can be considered as

waveguide materials. We showed that thermal oxide  $\text{SiO}_2$  was difficult to use as stopping layer for the releasing of ALD films due to its high stress. Instead we proposed the use of  $\text{SiN}$  and  $\text{KOH}$  etching experiments showed good result.  $\text{TiO}_2$  can only be used with  $\text{SiN}$  but we simulated that 250 nm  $\text{TiO}_2$  on 250 nm  $\text{SiN}$  has evanescent tale in both sides and can exhibit a sensitivity of 0.3

**Acknowledgements.** The authors would like to acknowledge the Dutch Technical Foundation, STW and the support provided by Electronic Instrumentation Laboratory (EI)–DIMES and Electronic Components, Technology and Materials (ECTM)–DIMES, TU Delft, The Netherlands and wish to thank IC process group of DIMES for technical assistance. The authors also thank to Eduardo Margalo for his assistance of optical measurement setup and loss propagation measurement.

## References

1. Chaeron, S.R.A.: An optical measurement system for the analysis of bacteria in drain fluid, in Dept. of Microelectronics. vol. Master Delft. TU Delft, The Netherlands (2007)
2. Kube, R., Mroczkowski, P., Granowski, D., Benedix, F., Sahm, M., Schmidt, U., Gastinger, I., Lippert, H.: Anastomotic leakage after colon cancer surgery: A predictor of significant morbidity and hospital mortality, and diminished tumour-free survival. *European Journal of Surgical Oncology (EJSO)* 36, 120–124 (2010)
3. Fumagalli, L., Natali, D., Sampietro, M., Peron, E., Perissinotti, F., Tallarida, G., Ferrari, S.:  $\text{Al}_2\text{O}_3$  as gate dielectric for organic transistors: Charge transport phenomena in poly-(3-hexylthiophene) based devices. *Organic Electronics* 9, 198–208 (2008)
4. Hoivik, N.D., Elam, J.W., Linderman, R.J., Bright, V.M., George, S.M., Lee, Y.C.: Atomic layer deposited protective coatings for micro-electromechanical systems. *Sensors and Actuators A: Physical* 103, 100–108 (2003)
5. Herrmann, C.F., DelRio, F.W., Miller, D.C., George, S.M., Bright, V.M., Ebel, J.L., Strawser, R.E., Cortez, R., Leedy, K.D.: Alternative dielectric films for rf MEMS capacitive switches deposited using atomic layer deposited  $\text{Al}_2\text{O}_3/\text{ZnO}$  alloys. *Sensors and Actuators A: Physical* 135, 262–272 (2007)
6. Mayer, T.M., Elam, J.W., George, S.M., Kotula, P.G., Goeke, R.S.: Atomic-layer deposition of wear-resistant coatings for microelectromechanical devices. *Applied Physics Letters* 82, 2883–2885 (2003)
7. Ferrari, S., Perissinotti, F., Peron, E., Fumagalli, L., Natali, D., Sampietro, M.: Atomic layer deposited  $\text{Al}_2\text{O}_3$  as a capping layer for polymer based transistors. *Organic Electronics* 8, 407–414 (2007)
8. Timoshenko, S.P., Gere, J.M.: *Theory of Elastic Stability*, 2nd edn. McGraw-Hill, New York (1985)
9. Jakschik, S., Schroeder, U., Hecht, T., Dollinger, G., Bergmaier, A., Bartha, J.W.: Physical properties of ALD- $\text{Al}_2\text{O}_3$  in a DRAM-capacitor equivalent structure comparing interfaces and oxygen precursors. *Materials Science and Engineering B* 107, 251–254 (2004)
10. Pollock, C., Lipson, M.: *Integrated Photonics*. Kluwer Academic Publication, Massachusetts (2003)
11. Hammer, M.: 2-D multilayer waveguide mode solver effective index approximation (2010), <http://wwwhome.math.utwente.nl/~hammer/eims.html>

**Part III**  
**Bio-inspired Systems and**  
**Signal Processing**

# Probabilistic Patient Monitoring with Multivariate, Multimodal Extreme Value Theory

Samuel Hugueny, David A. Clifton, and Lionel Tarassenko

Institute of Biomedical Engineering, Department of Engineering Science  
University of Oxford, Oxford, U.K.

samuel.hugueny@eng.ox.ac.uk, {davidc, lionel}@robots.ox.ac.uk

**Abstract.** Conventional patient monitoring is performed by generating alarms when vital signs exceed pre-determined thresholds, but the false-alarm rate of such monitors in hospitals is so high that alarms are typically ignored. We propose a principled, probabilistic method for combining vital signs into a multivariate model of patient state, using extreme value theory (EVT) to generate robust alarms if a patient's vital signs are deemed to have become sufficiently "extreme". Our proposed formulation operates many orders of magnitude faster than existing methods, allowing on-line learning of models, leading ultimately to patient-specific monitoring.

**Keywords:** Patient monitoring, Extreme value theory, Extreme value distributions, Density estimation.

## 1 Introduction

Many patients die in hospital every year because deterioration in physiological condition is not identified. It has been estimated that 23,000 cardiac arrests and 20,000 unforeseen admissions to ICU could be avoided each year in the UK alone, if deterioration were identified and acted upon sufficiently early [1,2]. Thus, there is a great need for patient monitoring systems that perform this automatic identification of patient deterioration.

### 1.1 Existing Patient Monitors

Conventional hospital patient monitors take frequent measurements of vital signs, such as heart rate, respiration rate, peripheral blood oxygen saturation ( $\text{SpO}_2$ ), temperature, and blood pressure, and then generate an alarm if any of these parameters exceed a fixed upper or lower threshold defined for that parameter. For example, many patient monitors will generate an alarm if the patient heart rate exceeds 160 BPM, or decreases below 40 BPM [3]. However, these single-channel alarming methods suffer from such high false-alarm rates that they are typically ignored in clinical practice; the study reported in [4] concluded that 86% of alarms generated by conventional monitors were false-positive.

## 1.2 Intelligent Patient Monitoring

The investigation described by this paper models the distribution of vital signs under “normal” patient conditions, and then detects when patient vital signs begin to deteriorate with respect to that model. This is the so-called “novelty detection” approach, in which patient deterioration corresponds to novelty with respect to a model of normality. We have previously applied this technique to the monitoring of other critical systems, such as jet engines [5] and manufacturing processes [6].

[7] and [3] used a Parzen window density estimator [8] to form a probabilistic model  $p(\mathbf{x})$  of the distribution of patient vital signs  $\mathbf{x}$  from a training set of vital signs observed from a population of stable, high-risk patients. However, alarms were generated by comparison of test data to a heuristic threshold set on  $p(\mathbf{x})$ . This threshold is termed the *novelty threshold*, because data exceeding it (in a probabilistic sense; i.e. data less probable than the threshold value) are classified “abnormal”.

Previous work presented in [9] and [10] has shown that such heuristic novelty thresholds do not allow on-line learning of patient models, because thresholds are not portable between models - primarily because they have no direct probabilistic interpretation. In that work, we described the use of Extreme Value Theory (EVT) as a principled method for determining if test data are “abnormal”, or “extreme”, with respect to some model of normality (such as a Gaussian Mixture Model, or GMM), which is summarised in Section 1.4. This process is automatic, and requires only the selection of a probabilistic novelty threshold (e.g.,  $P(\mathbf{x}) \leq 0.99$ ) in order to achieve accurate identification of patient deterioration.

## 1.3 Contributions in This Paper

Our previously-proposed work has a number of limitations:

1. The system described in [9] uses EVT for determining when multivariate test data are “extreme” with respect to a model of normality. In this case, a fully multimodal model is allowed, such as a GMM comprised of many Gaussian kernels. However, it is a numerical algorithm that requires large quantities of sampling, making it unsuitable for on-line learning of models that are frequently updated.
2. The system described in [10] provides a closed-form solution to the problems posed in (1) such that sampling is avoided, but is valid only for *unimodal* multivariate models consisting of a single Gaussian kernel. In practice, such single-kernel models are too simple to describe the distribution of training data accurately.

Thus, there is a need for an EVT algorithm that allows multimodal, multivariate models of normality to be constructed, overcoming the unimodal limitation of (2), while being computationally light-weight, overcoming the heavy sampling-based limitation of (1). This paper proposes such a method, described in Section 2, illustrates its use with synthetic data in Section 3, and presents results from a large patient monitoring investigation in Section 4.

## 1.4 Classical Extreme Value Theory

If we have a univariate probability distribution describing some univariate data,  $F(x)$ , classical EVT [11] provides a distribution describing where the most “extreme” of  $m$  points drawn from that distribution will lie. For example, if we draw  $m$  samples from a univariate Gaussian distribution, EVT provides a distribution that describes where the largest of those  $m$  samples will lie. It also provides a distribution that describes where the smallest of those  $m$  samples will lie. These distributions determined by EVT are termed the *Extreme Value Distributions* (EVDs). The EVDs tell us where the most “extreme” data generated from our original distribution will lie under “normal” conditions after observing  $m$  data. Thus, if we observe data which are more extreme than where we would expect (as determined by the EVDs), we can classify these data “abnormal”, and generate an alarm. This process lies at the heart of using EVT for patient monitoring, where we can classify observed vital signs as “extreme” if EVT determines that they lie further than one would expect under “normal” conditions (given by the EVDs).

Though classical EVT is defined only for univariate data, we present a generalisation of EVT to multivariate, multimodal models as described later in this paper.

To state this introduction more formally, consider  $\{x_m\}$ , a set of  $m$  independent and identically distributed random variables (iid rvs), which are univariate, and where each  $x_i \in \mathbb{R}$  is drawn from some underlying distribution  $F(x)$ . We define the maximum of this set of  $m$  samples to be  $M_m = \max(x_1, x_2, \dots, x_m)$ . EVT tells us the distribution of where to expect this maximum,  $M_m$ , and, noting that  $\min(x_1, x_2, \dots, x_m) = -\max(-x_1, x_2, \dots, x_m)$ , where to expect the minimum in our dataset. The fundamental theorem of EVT, the Fisher-Tippett theorem [12], shows that the distribution of the maximum,  $M_m$ , depends on the form of the distribution  $F(x)$ , and that this distribution of  $M_m$  can only take one of three well-known asymptotic forms in the limit  $m \rightarrow \infty$ : the Gumbel, Fréchet, or Weibull distributions.

The Fisher-Tippett theorem also holds for the distribution of minima and the EVDs of minima are the same as the EVDs of maxima, with a reverse  $x$ -axis. The Gumbel, Fréchet, and Weibull distributions are all special cases of the Generalised Extreme Value (GEV) distribution,

$$H_{GEV}^+(x; \gamma) = \exp\left(-[1 + \gamma x]^{-1/\gamma}\right), \quad (1)$$

where  $\gamma$  is known as the *shape parameter*. The cases  $\gamma \rightarrow 0$ ,  $\gamma > 0$  and  $\gamma < 0$  give the Gumbel, Fréchet and Weibull distributions, respectively. In the above, the superscript ‘+’ indicates that this is the EVD describing the maximum of the  $m$  samples generated from  $F$ .

## 1.5 Redefining “Extrema”

Classical univariate EVT (uEVT), as described above, cannot be directly applied to the estimation of multivariate EVDs. In the case of patient monitoring, for example, our data will be multivariate, where each dimension of the data corresponds to a different channel of measurement (heart rate, respiration rate, SpO<sub>2</sub>, etc.) In this multivariate



case, we no longer wish to answer the question “how is the sample of *greatest magnitude* distributed?”, but rather “how is the *most improbable* sample distributed?” This will allow us, as will be shown in Section 2, to generalise uEVT to a multivariate EVT (mEVT). As proposed in [9], we consider the following definition of extrema:

**Definition 1.** Let  $m \in \mathbb{N}^*$  and  $\{x_m\}$  be a sequence of (possibly multivariate) iid rvs, drawn from a distribution  $F$  with probability density function  $f$ . We define the extremum to be the random variable

$$E_m = \operatorname{argmin} \{f(X_1), \dots, f(X_m)\} .$$

## 1.6 Density Estimation

If a large number of actual observed extrema are available, or if it is possible to draw extrema from a generative model, then it is tempting to try and fit an EVD to those extrema, via Maximum Likelihood Estimation (MLE), for instance. If the form of the EVD for our dataset is known (i.e., whether it is Gumbel, Fréchet, or Weibull), one could attempt to fit a Gumbel, Fréchet or Weibull distribution directly to the extrema. Even if the form of the EVD is not known, the distribution of extrema is theoretically guaranteed to converge to one of the three instances of the GEV distribution, as stated by the Fisher-Tippett theorem.

This approach was taken in [9], in which a method was proposed to estimate the EVD in the case where the generative model is known to be a mixture of multivariate Gaussian distributions (a GMM). The GMM  $f(\mathbf{x})$  was constructed using a training set of observed multivariate data  $\{\mathbf{x}\}$ . The method is based on our capacity to generate (via sampling) a large number of extrema from the GMM. Each extremum is defined as being the sample of minimum probability density  $f(\mathbf{x})$  out of a set of  $m$  samples. Thus, if we require a large number of extrema (say,  $N = 10^6$ ), then we must generate  $N$  sets of  $m$  samples (where each set gives a single extremum).

In [13], this method was used for the purpose of patient monitoring. A GMM was trained using multivariate patient data, and the EVD for that model was estimated using the sampling method described above. A sliding window of length  $m$  was applied to the time-series of test patient data, where  $m$  was determined empirically. A window of test data was classified “abnormal” if its most extreme datum lay outside the estimated EVD.

This approach has a number of disadvantages. Estimating the EVD by generating extrema from the GMM is time-consuming. However, testing a range of values for  $m$  in order to find the optimal value is even more time-consuming: it requires us to generate a large number (e.g.,  $N = 10^6$ ) of extrema for *each value of  $m$  that we test*. If we wish to perform on-line learning, in which models are constructed in real-time from newly-acquired patient data, then these disadvantages must be overcome.

In Section 2, we propose a method to estimate numerically the EVD for a multivariate, multimodal model (such as a GMM) which does *not* require sampling of extrema, and so overcomes the disadvantages described above.

## 2 Method

### 2.1 Introduction

Though the Fisher-Tippett theorem (described in Section 1) is valid only for univariate data, we can use it to determine the EVD of an  $n$ -dimensional multivariate model  $F_n(\mathbf{x})$  using an approach from [9]. Rather than consider the EVD in the  $n$ -dimensional data space of  $\mathbf{x} \in \mathbb{R}^n$ , we can consider the EVD in the model's corresponding probability space  $F_n(\mathbb{R}) \subset \mathbb{R}^+$ . That is, we find the probability distribution over the model's probability density values. This new distribution (over probability density values) is univariate, and the Fisher-Tippett theorem applies.

We have previously shown in [10] that this can be used for multivariate, *unimodal* data; this paper proposes an extension to the method to allow us to cope with multivariate, *multimodal* data, as required when using a GMM to model the distribution of vital signs in patient monitoring.

### 2.2 Detail of Method

Define  $F_n$  to be a mixture of  $n$ -dimensional Gaussian kernels (i.e., a GMM), trained using example training data, for multivariate data  $\mathbf{x} \in \mathbb{R}^n$ . Now, consider the GMM's corresponding probability space: let  $\mathcal{P}$  be  $F_n(\mathbb{R}^n)$ , the image of  $\mathbb{R}^n$  under  $F_n$ . That is,  $\mathcal{P}$  is the set of all probability densities taken by the GMM, which will cover the range  $]0, p_{\max}]$ , where  $p_{\max}$  is the largest probability density taken by the GMM.

We can find the model's distribution over probability densities, which we define to be  $G_n$ :

$$\forall y \in \mathcal{P}, \quad G_n(y) = \int_{f_n^{-1}(]0, y])} f_n(\mathbf{x}) d\mathbf{x}, \quad (2)$$


where  $f_n^{-1}(]0, y])$  is the preimage of  $]0, y]$  under  $f_n$  (the set of all values of  $\mathbf{x}$  that give probability densities in the range  $]0, y]$ ). Thus,  $G_n(y)$  is the probability that data  $\mathbf{x}$  generated from the GMM will have probability density  $y$  or lower. The lower end of this distribution will be  $G_n(0) = 0$  because the probability of data having probability density  $p(\mathbf{x}) \leq 0$  is 0, and the upper end of this distribution will be  $G_n(p_{\max}) = 1$  because the probability of data having probability density  $p(\mathbf{x}) \leq p_{\max}$  is 1 (recalling that  $p_{\max}$  is the maximum probability density taken by the GMM).

Figure 1 shows  $G_n$  and its corresponding probability density function (pdf)  $g_n$  for an example 4-dimensional, bimodal GMM (in light grey). Note that the probability mass for models with dimensionality  $n > 2$  tends towards lower probability density values, as shown in [9]: a sample drawn from the GMM is more likely to have a low probability density  $y$  than a high value of  $y$ .

If  $F_n$  is composed of a single Gaussian kernel, an analytical form of  $G_n$  is derived in [10] and its pdf shown to be:

$$k_n(y, \beta) = \Omega_n \beta \left[ -2 \ln \left( (2\pi)^{n/2} \beta y \right) \right]^{n-2/2}, \quad (3)$$

where  $\Omega = \frac{2\pi^{n/2}}{\Gamma(\frac{n}{2})}$  (the total solid angle subtended by the unit  $n$ -sphere) and  $\beta = |\Sigma|^{1/2}$ , for covariance matrix  $\Sigma$ .



**Fig. 1.** Distributions in probability space  $y \in \mathcal{P}$  for an example bimodal GMM of dimensionality  $n = 4$ . In the upper plot, the pdf  $g_n(y)$  over probability density values  $y$  shows that the maximum probability density for this GMM is  $p_{\max} \approx 0.015$ . The estimating distribution  $k_n$  shows that the proposed method closely approximates the actual  $g_n$ . In the lower plot, the corresponding cdfs  $G_n$  and  $K_n$ . The histograms are obtained using  $10^7$  samples and 512 bins. Only the 64 bins on the left are used to estimate the tail of  $g_n(y)$ .

We can see from Equation (3) that  $k_n$  is independent of the mean of  $F_n$ , which is unsurprising: the probability density values taken by a Gaussian kernel are invariant under translations in the data space (as occurs when the mean is changed), but change if the kernel covariance is changed.

If  $F_n$  is composed of more than one Gaussian kernel, there is no analytical form for  $G_n$  or its pdf  $g_n$ . However, we can make the assumption that sufficiently far away from the modes of the distribution, a mixture of Gaussian kernels behaves approximately like a single Gaussian kernel. This assumption is typically valid because the EVD lies in the tails of  $F_n$ , not near its modes. This corresponds to the tail of  $g_n$ , where  $F_n(\mathbf{x})$  is close to zero, for which we wish to find the EVD.

Thus, for  $F_n(\mathbf{x})$  sufficiently close to zero,  $g_n$  can be approximated by  $k_n$  for some (positive) value of  $\beta$ . The family of parametric functions  $k_n$  can therefore be used to estimate  $g_n$ . A convenient feature of this method is that the family of  $k_n$  functions have a single scalar parameter,  $\beta$ . To estimate the value of  $\beta$  that best approximates the tail

of our  $g_n$ , we can estimate  $g_n$  using a histogram, and then find the value of  $\beta$  that minimises the least-square error in the tail.

Figure 1 shows that  $k_n$  and  $K_n$  accurately estimate  $g_n$  and  $G_n$  in the left-hand tail (where  $\mathcal{P}$  is close to zero), which is the area of interest for determining the EVD. So, if we can determine the EVD for  $k_n$  (and thus  $K_n$ ), we will have an accurate estimate of the EVD of our desired distribution  $G_n$ , and hence for our GMM,  $F_n$ .

From [10],  $k_n$  is known to be in the domain of attraction of the minimal Weibull EVD:

$$H^-(y; d_m, c_m, \alpha_m) = 1 - \exp \left[ - \left( \frac{y - d_m}{c_m} \right)^{\alpha_m} \right], \tag{4}$$

where its scale, location, and shape parameters  $c_m$ ,  $d_m$ , and  $\alpha_m$ , respectively, are given by:

$$c_m = K_n^{\leftarrow} \left( \frac{1}{m} \right), \tag{5}$$

$$d_m = 0, \tag{6}$$

$$\alpha_m = m c_m k_n[c_m], \tag{7}$$

where  $K_n$  is the integral of  $k_n$ , which is given in [10], and where  $K_n^{\leftarrow} \left( \frac{1}{m} \right)$  is the  $1/m$  quantile of  $K_n$ .

After estimation of  $\beta$ , we can use Equations (5), (6), and (7) to define entirely the EVD of our  $G_n$ .

### 2.3 Novelty Score Assignment


Having estimated  $c_m$ ,  $d_m$ , and  $\alpha_m$ , let  $\mathbf{x}_m = \{x_1, \dots, x_m\}$  be a set of  $m$  samples drawn from  $F_n$ . Let  $y$  be  $\min [f(x_1), \dots, f(x_m)]$ . Then, the quantity  $1 - H^-(y; d_m, c_m, \alpha_m)$  is the probability of drawing an extremum out of  $m$  samples with a greater density value; i.e. a *more likely* extremum. This is interpreted as the probability for the extremum to be novel with respect to the model. As it is desirable for novelty scores to take low values for normal data and higher values for increasingly abnormal data, we define the novelty score function:

$$q(\mathbf{x}_m) = \phi \left[ 1 - H^-(y; d_m, c_m, \alpha_m) \right], \tag{8}$$

where  $y$  is defined above, and  $\phi$  is a monotonically increasing function with domain  $]0, 1]$ .  $\phi$  must be the identity function to keep the probabilistic interpretation of novelty scores, but can be taken to be some other function for visualisation purposes. Figure 2 shows an example of novelty score assignment for an example bivariate GMM.

## 3 Validation on Simulated Data

To validate our approach, we compare EVDs obtained using Equations (5), (6), and (7) with the EVDs obtained using Maximum Likelihood Estimation (MLE) of the Weibull parameters, using simulated data. An application using real patient vital-sign data is shown in Section 4.



**Fig. 2.** From top left to bottom right: pdf of example bivariate 4-kernel GMM, associated novelty scores for  $m = 5, 15, 50, 150$  and  $500$ .  $\phi$  is the cubic root function. Black and white indicate a probability zero and one, respectively, of drawing an extrema of greater density value. As  $m$  increases, extrema move further away from the kernel centres and ultimately further away from the distribution centre.


For dimensionality  $n = 1$  to  $6$ , we define  $F_n$  to be the  $n$ -dimensional mixture of Gaussians comprised of two multivariate standard Gaussian distributions with equal priors and a Euclidean distance between their centres equal to two.

In order to estimate the EVD using MLE, for each dimensionality  $n = 1 \dots 6$ , and for increasing values of  $m$ , a large number of extrema (e.g.,  $N = 10^6$ ) must be sampled. Figure 3 shows estimates obtained using MLE for both the scale  $c_m$  and shape  $\alpha_m$  parameters of the EVD. The figure also shows parameters estimated using the method proposed in Section 2.

The scale parameter appears to be accurately estimated even for small values of  $m$ . However, the proposed method's use of Equation (7) to estimate the shape parameter only matches the MLE estimate for values of  $m$  greater than 15. This was expected, as the Fisher-Tippett theorem tells us that the Weibull distribution is the EVD for asymptotically increasing  $m$ , and that actual EVDs are not expected to match the Weibull distribution closely for small values of  $m$ .

Figure 4 presents a comparison between the cdfs of the corresponding distributions estimated using MLE and with the proposed method, for  $n = 4$  and a range of values of  $m$ . Taking into account the logarithmic scale in  $y$ , we conclude that solutions obtained using the new method are a good match to the maximum likelihood estimates.

The main advantage of our approach is that it does not require sampling of extrema, which is a particularly intensive process. Assuming a model  $F_n$ , we only need to obtain  $N$  samples from that model to build a histogram approximating  $G_n$ , then we solve a simple least-squares estimation problem (as described in Section 2), and finally apply




**Fig. 3.** Comparison of results of MLE estimates of the scale parameter  $c_m$  (top) and the shape parameter  $\alpha_m$  (bottom) parameter (shown as points in the plots), and values obtained using Equations (5) and (7) for increasing  $m$  for  $n = 1$  to 6 and increasing values of  $m$  (shown as continuous lines). For each dimensionality  $n$ , the GMM  $F_n$  is composed of two standard Gaussian kernels with equal priors, and a Euclidean distance between their centres equals to two. Error bars are too small to be visible at this scale.

the closed-form Equations (5), (6), and (7) to obtain an estimate of the Weibull parameters for *any* value of  $m$ . On the other hand, the MLE (which in itself is more intensive than the least-square estimation problem) requires  $m \times N$  samples to be drawn to obtain  $N$  extrema, and this is for a *single* value of  $m$ . To test all values of  $m$  between 1 and 100 for instance, our algorithm requires up to 5,000 times less sampling, and none of the 100 iterations of the MLE algorithm.

## 4 Application to Vital-sign Data

In this section, we describe an application of our methodology to a patient monitoring problem, using a large dataset of patient vital-sign data obtained from a clinical trial [3]. The dataset is comprised of 18,000 hours of vital-sign data collected from 332 high-risk adult patients. Measurements of heart rate (HR), breathing rate (BR) and oxygen saturation (SpO<sub>2</sub>) are available at 1 Hz. The data were reviewed by clinical experts and “crisis events” were labelled, corresponding to those events that should have resulted in a call to a Medical Emergency Team [14,15] being made on the patient’s behalf.



**Fig. 4.** Logarithmic plot of cumulative distributions obtained using our proposed method (black) and Maximum Likelihood Estimation (grey). Dimensionality  $n = 4$ , histograms and MLE use  $10^5$  samples. From right to left, the values of  $m$  are 2, 5, 10, 30, 50, 100, 200 and 500.


We split the available data into three subsets: (i) a training set and (ii) a control set, each consisting of data from 144 “normal” patients (and each containing approximately 8000 hours of data); (iii) a test set consisting of data from the 44 patients who went on to have crisis events (approximately 2000 hours) which includes “abnormal” data labelled by clinical experts (approximately 43 hours).

The training set is used to construct a model of normality  $F$  (with pdf  $f$ ), consisting of a trivariate GMM (noting that  $n = 3$ , corresponding to the number of physiological parameters available in the dataset). The number of kernels in the GMM was estimated via cross-validation, which showed that 9 kernels provided the lowest overall cross-validation error.

Given a value of  $m$ , the values of  $c_m$ ,  $d_m$ , and  $\alpha_m$  can be computed as described in Section 2. Novelty scores are then assigned to all patient data using eq 8, with  $\phi$  the identity function and,  $y = \min [f(\mathbf{x}_{t-m+1}), f(\mathbf{x}_{t-m+2}), \dots, f(\mathbf{x}_t)]$ . That is,  $y$  is the datum with minimum probability density within a window containing the last  $m$  vital-sign measurements. This definition of  $y$  ensures that the extremum of  $m$  samples is considered at each time step. The value of  $m$  conditions the width of the sliding time-window used to assign novelty scores.

Setting a threshold on the novelty score function  $q$  allows us to separate “normal” from “abnormal” data, and therefore compute a true positive rate (TPR) and a false positive rate (FPR) for each of the three data subsets described above, with respect to the known labels provided by clinical experts. Varying this threshold yields the ROC curves shown in Figure 5.

We note that the setting of a novelty threshold on the EVD is different to the conventional method of setting a novelty threshold on the pdf  $f_n$  given by the GMM. In EVT-based approaches, the threshold corresponds to a direct probabilistic interpretation (e.g., “these data are abnormal with a probability of 0.99”), whereas conventional thresholding of the GMM  $f_n$  is heuristic (as described in Section 1.2), being based on probability density values, and so such thresholds are not portable between different models.



**Fig. 5.** True positive rate vs. false alarm rate for the “control” and “test” group, plotted for different values of  $m$ .

The absence of data points above a true positive rate of 92% is due to the heterogeneity of the data within the crisis windows, a portion of which cannot be considered abnormal with respect to the model.

As the dynamic range of a change in patient status is not known, it is in our best interest to be able to explore a range of values for  $m$ . Depending on what is considered an acceptable true positive rate for the crisis data, one can choose the value of  $m$  that minimises the false alarms rate for the control group. A small value of  $m$  seems to be preferable if the desired TPR is between 0.65 and 0.8. If we wish to maximise the TPR, however, our results suggest that we should take a large value of  $m$ .

## 5 Discussion

### 5.1 Conclusions

This paper has proposed a new method for estimating the extreme value distributions of multivariate, multimodal mixture models, as is required for the analysis of complex datasets such as those encountered in patient vital-signs monitoring. The method overcomes the limitations of previous methods, by (i) providing a light-weight formulation that is shown to be significantly faster than previous maximum-likelihood methods, which require large amounts of sampling, and (ii) providing solutions for *multimodal* multivariate models, as are required for the analysis of complex datasets, whereas previous closed-form approaches were limited to *unimodal* multivariate models.

We have validated our methodology using synthetic data and patient vital-sign data from a large clinical trial, and have shown that EVDs estimated using the method are a good match to those obtained using maximum-likelihood methods, particularly when the value of EVT parameter  $m$  (the window length) is greater than 15. For most real datasets, in which the sampling rate is relatively fast, larger values of  $m$  will be necessary in order to model system dynamics. For example, in the case of patient vital-signs



monitoring presented in this paper, in which vital-signs data were obtained at 1 Hz, a value of  $m = 15$  corresponds to a window length of 15s.

As shown in Section 3 because the EVD is known in closed form and is parameterised by  $m$ , the value of  $m$  can be optimised in real-time. The light-weight formulation allows on-line learning of models, ultimately allowing patient-specific monitoring to take place, in which models are constructed in real-time using data observed from a new monitored patient.

## 5.2 Future Work

The solutions proposed in this paper, while validated only for mixtures of Gaussian kernels are sufficiently general that they should apply to any kernel mixture model. For example, the proposed method could also be used to find the extreme value distributions corresponding to Parzen windows estimators (themselves also mixtures of Gaussian distributions); mixtures of Gamma distributions, as used by [16]; mixtures of Student's  $t$  distributions, as proposed by [17], and mixtures of Weibull distributions, as proposed by [18].

These solutions are based on closed form formulae, and so the light-weight approach could facilitate the use of Bayesian parameter estimation.

In application to patient monitoring, as well as demonstrating benefit on existing datasets (as shown in this paper), we hope to have provided the facility to perform on-line learning of patient-specific models, which forms an important part of our future work.

## 5.3 Downloadable Content

Example MATLAB code on how to use our multivariate EVT is available for download at <http://www.robots.ox.ac.uk/~davidc>.

**Acknowledgements.** SH was supported by the EPSRC LSI Doctoral Training Centre, Oxford, and DAC was supported by the NIHR Biomedical Research Centre, Oxford and the Centre of Excellence in Personalised Healthcare funded by the Wellcome Trust and EPSRC under grant number WT 088877/Z/09/Z. DAC wishes to thank the Abbey-Santander fund for the award that made publication of this paper possible. The authors wish to thank Iain G.D. Strachan of Oxford BioSignals Ltd. and Lei A. Clifton of the University of Oxford for useful discussions.

## References

1. Hodgetts, T., Kenward, G., Vlackonikolis, I., Payne, S., Castle, N., Crouch, R., Ineson, N., Shaikh, L.: Incidence, location, and reasons for avoidable in-hospital cardiac arrest in a district general hospital. *Resuscitation* 54, 115–123 (2002)
2. McQuillan, P., Pilkington, S., Allan, A., Taylor, B., Short, A., Morgan, G., Nielsen, M., Barrett, D., Smith, G.: Confidential inquiry into quality of care before admission to intensive care. *British Medical Journal* 316, 1853–1858 (1998)

3. Hann, A.: Multi-parameter monitoring for early warning of patient deterioration. PhD thesis, University of Oxford (2008)
4. Tsien, C., Fackler, J.: Poor prognosis for existing monitors in the intensive care unit. *Critical Care Medicine* 25(4), 614–619 (1997)
5. Clifton, D., McGrogan, N., Tarassenko, L., King, S., Anuzis, P., King, D.: Bayesian extreme value statistics for novelty detection in gas-turbine engines. In: *Proceedings of IEEE Aerospace*, Montana, USA, pp. 1–11 (2008)
6. Clifton, D., Tarassenko, L., Sage, C., Sundaram, S.: Condition monitoring of manufacturing processes. In: *Proceedings of Condition Monitoring 2008*, Edinburgh, UK, pp. 273–279 (2008)
7. Tarassenko, L., Hann, A., Young, D.: Integrated monitoring and analysis for early warning of patient deterioration. *British Journal of Anaesthesia* 98(1), 149–152 (2006)
8. Parzen, E.: On estimation of a probability density function and mode. *Annals of Mathematical Statistics* 33, 1065–1076 (1962)
9. Clifton, D., Hugueny, S., Tarassenko, L.: Novelty detection with multivariate extreme value theory, part I: Numerical approach for multimodal estimation. In: *Proceedings of IEEE Workshop on Machine Learning in Signal Processing* (2009)
10. Hugueny, S., Clifton, D., Tarassenko, L.: Novelty detection with multivariate extreme value theory, part II: Analytical approach for unimodal estimation. In: *Proceedings of IEEE Workshop on Machine Learning in Signal Processing* (2009)
11. Embrechts, P., Klüppelberg, C., Mikosch, T.: *Modelling Extremal Events for Insurance and Finance*. Springer, Heidelberg (1997)
12. Fisher, R., Tippett, L.: Limiting forms of the frequency distributions of the largest or smallest members of a sample. *Proc. Camb. Philos. Soc.* 24, 180–190 (1928)
13. Clifton, D., Hugueny, S., Tarassenko, L.: A comparison of approaches to multivariate extreme value theory for novelty detection. In: *Proceedings of IEEE Workshop on Statistical Signal Processing* (2009)
14. Hourihan, F., Bishop, G., Hillman, K., Daffurn, K.: The medical emergency team: a new strategy to identify and intervene in high-risk patients. *Clinical Intensive Care* 6(6), 269–272 (1995)
15. Hillman, K., Chen, J., Cretikos, M., Bellomo, R., Brown, D., Doig, G., Finfer, S., Flabouris, A.: Introduction of the medical emergency team (MET) system: a cluster-randomised controlled trial. *The Lancet* 365(9477), 2091–2097 (2005)
16. Mayrose, I., Friedman, N., Pupko, T.: A gamma mixture model better accounts for among site rate heterogeneity. *Bioinformatics* 21(2), 151–158 (2005)
17. Svensen, M., Bishop, C.: Robust Bayesian mixture modelling. *Neurocomputing* 64, 235–252 (2005)
18. Ebden, M., Stranjak, A., Dutta, P., Rogers, A.: A multi-agent simulation system for prediction and scheduling of aero engine overhaul. In: *Proceedings of the 7th Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Estoril, Portugal (2008)

# Weighted SVMs and Feature Relevance Assessment in Supervised Heart Beat Classification

Gael de Lannoy<sup>1,2</sup>, Damien François<sup>1</sup>, Jean Delbeke<sup>2</sup>, and Michel Verleysen<sup>1,3</sup>

<sup>1</sup> Université catholique de Louvain, Machine Learning Group  
Place du levant 3, 1348 Louvain-La-Neuve, Belgium

<sup>2</sup> Université catholique de Louvain, Departement of physiology and pharmacology  
av. Hippocrate 54, 1200 Bruxelles, Belgium

<sup>3</sup> Université Paris I Panthéon-Sarbonne, SAMM  
Rue de Tolbiac 90, F-75634 Paris Cedex 13, France  
{gael.delannoy,damien.francois,jean.delbeke,  
michel.verleysen}@uclouvain.be

**Abstract.** The diagnosis of cardiac dysfunctions requires the analysis of long-term ECG signal recordings, often containing hundreds to thousands of heart beats. In this work, automatic inter-patient classification of heart beats following AAMI guidelines is investigated. The prior of the normal class is by far larger than the other classes, and the classifier obtained by a standard SVM training is likely to act as the trivial acceptor. To avoid this inconvenience, a SVM classifier optimizing a convex approximation of the balanced classification rate rather than the standard accuracy is used. First, the assessment of feature sets previously proposed in the literature is investigated. Second, the performances of this SVM model is compared to those of previously reported inter-patient classification models. The results show that the choice of the features is of major importance, and that some previously reported feature sets do not serve the classification performances. Also, the weighted SVM model with the best feature set selection achieves results better than previously reported inter-patient models with features extracted only from the R spike annotations.

## 1 Introduction

The analysis of electrocardiogram (ECG) signals provides with critical information on the cardiac function of patients. Long-term recordings of ECG's are required, for example, for the clinical diagnosis of some disease conditions, and for the evaluation of new drugs during phase-one studies. Such long-term recordings are usually obtained using the popular Holter recorders.

Holter ambulatory systems record from 24 to 48 hours of heart activity, providing with data containing thousands of heart beats. The analysis is usually performed off-line by cardiologists, whose diagnosis may rely on just a few transient patterns. Because of the high number of beats to scrutinize, this task is very expensive and reliable visual inspection is difficult. Computer-aided classification of pathological beats is therefore of great importance to help physicians perform correct diagnoses.

Automating this task, however, is difficult in real situations for several reasons. First, several sources of noise pollute the ECG signal. Among these, power line interferences,

muscular artifacts, poor electrode contacts and baseline wandering due to respiration can sometimes be identified. Second, the vast majority of the heart beats are normal healthy beats while just only a small number of beats are pathological, and of course those are of uttermost importance. Third, artificial intelligence methods used to automate classification require the extraction of discriminative features from the ECG signals. Unfortunately, very few information is available to decide how to extract and build those features.

Several features characterizing the heart beats and several classification models have been investigated in the literature [6]. However, very few reported experiments follow the standards defined by the American Association for Medical Instrumentation (AAMI), which makes it very difficult to assess the relative merits of the methods and of the proposed extracted features [1]. Furthermore, critical issues, such as the class imbalance are often not taken specifically into account.

Another important aspect of heart beat classification is the way the training and test sets are composed. Most published methods for classifying the heart beats of a patient require access to previous data from that particular patient. In other words, data for each patient must be present both in training and test sets. We refer to this as “intra-patient” classification. Of course, one cannot always ensure that previous data is known and available for each patient when needed, especially new ones. Therefore, “inter-patient” classification – classifying the beats of a new tested patient according to a reference database built from data from other patients – is much more useful, but at the same time, much harder and demanding in terms of generalization capabilities of the models.

In this paper, inter-patient classification of heart beats following the AAMI guidelines is investigated. The class imbalance is taken into account by using a *weighted* support vector machine classifier. A large number of features extracted from the ECG signals described in the literature are combined and evaluated, and their relevance is assessed.

The remaining of this paper is organized as follows. Section 2 briefly reviews the state of the art in heart beat classification. Section 3 provides a short overview of the theoretical background for the mathematical methods used in this work. Section 4 describes the ECG database and Section 5 describes the features extracted from the heart beats. Finally, Section 6 shows the experiments and the results.

## 2 State of the Art

This section provides a short review of the state of the art in supervised heart beat classification. Two kinds of heart beat classification paradigms can be distinguished, corresponding to either *intra-patient* classification or *inter-patient* classification.

Inter-patient classification consists in classifying the beats of a new tested patient according to a reference database and a model built from data from other patients. This process thus implies generalization from one patient to another. As far as intra-patient classification is concerned, the reference database must contain previously labeled beats from the tested patient. The results that can be achieved are naturally better than when inter-patient classification is performed, but the patient labeled beats are usually not available in real situations. Furthermore, because pathological beats can be very rare,

there is no guarantee that the few training beats that would be labeled for this patient would contain representatives for each class; and the classifier could possibly fail in predicting something it has not learned.

Despite these major drawbacks, the majority of previously reported work is about intra-patient classification. Different models have been proposed for this task, including neural networks [13],  $k$ -nearest neighbors [5], hidden Markov models [4] and support vector machines (SVMs) [12]. A comprehensive review of intra-patient classification methods and their results can be found in [6].

As far as inter-patient classification is concerned, the first study to our knowledge to establish a reliable inter-patient classification methodology following AAMI standards is the work of [3]. A linear discriminant classifier model is trained and the results are evaluated on an independent test set. The unbalanced classes issue is addressed by introducing weights in the linear discriminant functions. Another linear discriminant model and a Mahalanobis distance classifier are used in [17] on the same dataset with different features and worst results. In [15], hierarchical SVMs are used to reduce the effect of the class unbalance on the same dataset.

The classification performances heavily rely on the extraction of relevant features from the heart beat time series. A variety of features have been proposed to characterize the heart beats. The representation of the heart beat signal by the coefficients of Hermite basis functions expansions is introduced for a clustering application in [11], and later used for classification by [14] and [15]. Another type of features that has been proposed is the representation of the heart beats by higher order statistics, and in particular the cumulants of order 2, 3 and 4 [13].

Other widely used groups of features are segmentation interval features and morphological features [53]. Morphological features correspond to the raw amplitude values of the ECG signal. Segmentation features require the annotation of the ECG characteristic points and then summarize the shape of the P, QRS and T waves by their duration, area, height, Q-T intervals, S-T intervals, etc. In most of previously reported works, the successive time differences between the R spikes of heart beats (later referred to as R-R interval features) are always combined to the other features. However, the intrinsic relevance of each type of features remains unknown. In this paper, this relevance is investigated using feature selection techniques [79].

### 3 Theoretical Background

This section provides a brief theoretical background on mathematical methods that are used in this work.

#### 3.1 Weighted SVM Classifier

A support vector machine (SVM) is a supervised learning method introduced by Vapnik [18]. The two-class case is described here, because its extension to multiple classes is straightforward by applying the one-against-all or one-against-one methods. Let us first define the  $p$  dimensional feature vector  $\mathbf{x}_i = \{x_i^1, x_i^2, \dots, x_i^p\}$  and the associated class value  $y_i \in \{-1, 1\}$  for a given heart beat  $i$  with  $i$  ranging from 1 to  $N$ ,  $N$  being the total number of heart beats in the dataset.

SVMs are linear machines that rely on a preprocessing to represent the features in a higher dimension, typically much higher than the original feature space. With an appropriate non-linear mapping  $\varphi(\mathbf{x})$  to a sufficiently high-dimensional space, finite data from two categories can always be separated by a hyperplane. In SVMs, this hyperplane is chosen as the one with the largest margin.

Assume each observation  $\mathbf{x}_i$  has been transformed to  $\mathbf{z}_i = \varphi(\mathbf{x}_i)$ . The soft-margin formulation of the SVM allows examples to be misclassified or to lie inside the margin by the introduction of slack variables  $\xi_i$  in the problem constraints:

$$\text{Minimize } W(\mathbf{w}, b, \xi) = \frac{1}{2} \|\mathbf{w}\|^2 + C \cdot \Phi(\xi) \tag{1}$$

$$\text{s.t. } \begin{cases} y_i(\langle \mathbf{w}, \mathbf{z}_i \rangle + b) \geq 1 - \xi_i \quad \forall i = 1..N \\ \xi_i \geq 0 \quad \forall i = 1..N \end{cases} \tag{2}$$

where  $\mathbf{w}$  and  $b$  are the parameters of the hyperplane. The  $\Phi()$  term introduced in the objective function is used to penalize solutions with many training errors. For any feasible solution  $(\mathbf{w}, b, \xi)$ , misclassified examples have an associated slack value  $\xi$  greater than 1. Hence, it seems natural to chose a function counting the number of slacks greater or equal to one as penalization function  $\Phi()$  (the accuracy). Unfortunately, the optimization of such a function combined with the margin criterion is known to be NP hard. In practice, an approximation of the accuracy is thus commonly used:

$$\Phi(\xi) = \sum_{i=1}^n \xi_i \tag{3}$$

In this classical primal SVM formulation, the model is therefore actually optimizing a convex approximation of the accuracy [2].

However, in a heart beat classification task, around 90% of beats are normal beats and a dummy classifier which would always predict the normal class would get 90% accuracy. For this reason, the classification accuracy itself is a biased performance measure and the balanced classification rate (BCR) is to be preferred. In a two-class problem, the BCR is defined as the average between the sensitivity (SE) and the specificity (SP):

$$BCR = \frac{1}{2}(TP/(TP + FN) + TN/(TN + FP)) \tag{4}$$

$$= (SE + SP)/2 \tag{5}$$

where TP is the count of true positives, FN of false negatives, TN of true negatives and FP of false positives. In a multi-class problem, the BCR is equal to the average of class accuracies. Using this measure, a biased two-class classifier predicting always the normal class would only get 50% of BCR, correctly measuring that only one class has been correctly classified.

In the SVM formulation, different penalties for each class can be included in the objective function [10] so that a convex approximation of the BCR is optimized, rather than the accuray. Let  $N^+$  and  $N^-$  denote respectively the number of positive and

negative examples. To eliminate the effect of the unequal class sizes on the classification error rate, the objective function is rewritten as:

$$\Phi(\xi) = \frac{N}{N_+} \sum_{\{i|y_i=1\}} \xi_i + \frac{N}{N_-} \sum_{\{i|y_i=-1\}} \xi_i \quad (6)$$

By constructing the Lagrangian, this *primal* formulation can be rewritten in a so-called *dual* form. The optimization is then achieved by solving the system using quadratic programming. For this type of optimization, there exist many effective learning algorithms. A common method is Platt's Sequential Minimal Optimization (SMO) algorithm, which breaks the problem down into 2-dimensional sub-problems that may be solved analytically, eliminating the need for a numerical optimization algorithm such as conjugate gradient methods [16].

In the dual form, the explicit form of the mapping function  $\varphi$  must not be known as long as the kernel function  $K(\mathbf{x}_i, \mathbf{x}_j) = \varphi(\mathbf{x}_i)\varphi(\mathbf{x}_j)$  is defined. The kernel can for example be the linear kernel  $K(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_i^t \mathbf{x}_j$  or the radial basis function kernel  $K(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2)$  where  $\gamma$  is a kernel parameter to be tuned.

### 3.2 Hermite Basis Functions

The representation of the heart beat signal via Hermite basis functions (HBF) was first introduced by [11] for a clustering application and later by [14] for classification. This approach exploits similarities between the shapes of HBF and typical ECG waveforms. Let us denote the heart beat signal by  $x(t)$ . Its expansion into a Hermite series of order  $N$  is written as

$$x(t) = \sum_{n=0}^{N-1} c_n \phi_n(t, \sigma) \quad (7)$$

where  $c_n$  are the expansion coefficients and  $\sigma$  is the width parameter.  $\phi_n(t, \sigma)$  are the Hermite basis functions of the  $n$ th order defined as follows:

$$\phi_n(t, \sigma) = \frac{1}{\sqrt{\sigma 2^n n! \sqrt{\pi}}} e^{-t^2/2\sigma^2} H_n(t/\sigma) \quad (8)$$

where  $H_n(t/\sigma)$  is the Hermite polynomial of the  $n$ th order. The Hermite polynomials satisfy the following recurrence relation:

$$H_n(x) = 2xH_{n-1}(x) - 2(n-1)H_{n-2}(x) \quad (9)$$

with  $H_0(x) = 1$  and  $H_1(x) = 2x$ .

The higher the order of the Hermite polynomial, the higher its frequency of changes in the time domain, and the better the capability of the expansion in Eq. 7 to reconstruct the signal [6]. The width parameter  $\sigma$  can be tuned to provide a good representation of beats with large differences in durations. The coefficients  $c_n$  of the HBF expansion can be estimated by minimizing the sum of squared errors using singular value decomposition and the pseudo-inverse technique. These coefficients summarize the shape of the heart beat signal and can be treated as the features used in the classification process.

### 3.3 Higher Order Statistics

The statistical properties of the heart beat signal can be represented by its higher order statistics (HOS). The cumulants of order two, three and four are usually used [13]. Assuming the heart beat signal  $x(t)$  has a zero mean, its cumulant  $C_x^i$  of order  $i$  can be computed as follows:

$$\begin{aligned} C_x^2(\tau_1) &= E\{x(t)x(t + \tau_1)\} \\ C_x^3(\tau_1, \tau_2) &= E\{x(t)x(t + \tau_1)x(t + \tau_2)\} \\ C_x^4(\tau_1, \tau_2, \tau_3) &= E\{x(t)x(t + \tau_1)x(t + \tau_2)x(t + \tau_3)\} \\ &\quad - C_x^2(\tau_1)C_x^2(\tau_3 - \tau_2) \\ &\quad - C_x^2(\tau_2)C_x^2(\tau_3 - \tau_1) \\ &\quad - C_x^2(\tau_3)C_x^2(\tau_2 - \tau_1) \end{aligned}$$

where  $E$  is the expectation operator and  $\tau_1, \tau_2, \tau_3$  are the time lags.

## 4 ECG Database and Preprocessing

Data from the MIT-BIH arrhythmia database [8] are used in this work. The database contains 48 half-hour long ambulatory recordings obtained from 48 patients, for a total of approximately 110'000 heart beats manually labeled into 15 distinct types. Following the AAMI recommendations, the four recordings with paced beats are rejected.

The sampled ECG signals are first filtered to remove unwanted artifacts. The filtering procedure defined in [3] is used in this work. The ECG signal is filtered by two median filters. The first median filter is of 200 msec width and removes the QRS complexes and the P waves. The resulting signal is then processed with a second median filter of 600 msec width to remove the T waves. The signal resulting from the second filter operation contains the baseline wanderings and can be subtracted from the original signal. Powerline and other high frequency artifacts are then removed from the baseline corrected signal with a FIR filter.

The dataset configuration is the same as in [3][5]. The 44 available recordings are divided in two independent datasets of 22 recordings each with approximately the same ratio of heart beats classes. The first dataset is the training set, and is used to build the model. The second dataset is the test set, and is used to obtain an independent measure of the performances of the classifier.

The ECG characteristic points, corresponding to the onset and offset of P, QRS and T waves, are then annotated in each recording and the heart beats are separated using the standard *ecgpuwave*<sup>1</sup> segmentation software provided with the MIT-BIH arrhythmia database. The original R spike annotations are provided as input to the software to help the detection of the characteristic points. The beats whose Q and S points were not detected are considered as outliers and automatically rejected from our datasets.

The MIT-BIH heart beat labeled types are grouped according to the AAMI recommendations into four more clinically relevant heart beat classes [1] (see Table 2 for grouping details):

<sup>1</sup> See <http://www.physionet.org/physiotools/software-index.shtml>



**N-class** includes beats originating in the sinus node (normal and bundle branch block beat types);

**S-class** includes supraventricular ectopic beats;

**V-class** includes ventricular ectopic beats (VEBs);

**F-class** includes beats that result from fusing normal and VEBs.

Table 1 shows the number of beats in each class and their frequencies in the two datasets.

**Table 1.** Distribution of heart beat classes in the two independent datasets

	N	S	V	F	Total
Training	45801	938	3708	414	50861
	90.05%	1.84%	7.29%	0.81%	100%
Test	44202	1835	3204	388	49629
	89.06%	3.7%	6.46%	0.78%	100%

## 5 Feature Extraction

A large variety of features extracted from a heart beat signal have been proposed in the literature. They can be summarized in five feature groups: R-R intervals, segmentation intervals, morphological features, Hermite basis function expansion coefficients and higher order statistics. The features computed in this work for these different groups are now detailed.

1. R-R intervals: This group consists of three features built from the R-R interval series. The first three features are the R-R interval to the previous beat, the R-R interval to the next beat and the average of R-R intervals in a window of 10 surrounding beats.
2. Segmentation intervals: A large variety of 24 features are computed from the annotated characteristic points. These features include a boolean flag indicating the presence or absence of QRS, P and T waves. If the waves are present, their duration, maximum and minimum values, area, standard deviation, skewness and kurtosis are computed as features. The Q-T, S-T, Q-R, R-S intervals are also included. When the characteristic points needed to compute a feature failed to be detected in the heart beat annotation step, the feature value is set to the patient's mean feature value.
3. Morphological features: Ten features are derived by uniformly sampling the ECG amplitude in a window defined by the onset and offset of the QRS complex, and nine other features in a window defined by the QRS offset and the T-wave offset [3]. As the ECG signals were already sampled, linear interpolation was used to estimate the intermediate values of the ECG amplitude. Here again, when the onset or offset points needed to compute a feature were not detected, the feature value is set to the patient's mean feature value.
4. HBF coefficients: The parameters for computing the HBF expansion coefficients as defined in [15] are used. The order of the Hermite polynomial is set to 20, and the width parameter  $\sigma$  is estimated so as to minimize the reconstruction error for each beat.

5. Higher order statistics: The 2nd, 3rd and 4th order cumulant functions are computed. The parameters as defined in [14] are used: the lag parameters range from -250 msec to 250 msec centered on the R spike and 10 equally spaced sample points of each cumulant are used as features, for a total of 30 features.

Two additional normalized feature groups are also computed:

6. Normalized R-R intervals: These features correspond to the ratio between the previous three R-R values and their mean value for this patient. These last features are thus independent from the mean normal behavior of the heart of patients, which can naturally be very different between individuals, possibly misleading the classifier.
7. Normalized segmentation intervals: This group of features contains the same features as the previous segmentation group, but the values are normalized by their mean value for each patient. The normalization is obviously not applied to boolean segmentation features. Here again, the objective is to make each feature independent from the mean behavior of the heart of a patient, because it can naturally be very different between individuals.

## 6 Experiments and Results

In the following experiments, weighted SVMs with the *one-against-one* multi-class strategy are used. The training of the model is performed on the training dataset introduced in section 4, and this model is then applied to get a prediction of the class label of new heart beats from other patients in the test set.

Several types of kernels are evaluated in this study, and the linear kernel always outperforms the other kernels. As this is in accordance with previous works on heart beat classification using SVMs [15], only the results with the linear kernel are reported here. The regularization parameter  $C$  of the SVM model is optimized between  $10^1$  and  $10^{-5}$ .

### 6.1 Feature Relevance Assessment

The first experiment aims at assessing the relevance of each set of features (cfr. Section 5) in an exhaustive, “wrapper”, approach: every possible combination of the seven feature sets are built, and each of them is fed to a weighted SVM model, whose performances are taken as relevance measure of the feature combination. Table 3 holds the most interesting results out of the  $2^7 - 1 = 127$  configurations. Figure 1 shows how often (in percentage) each feature set is present in the top 20 best configurations.

Best overall performances are obtained with the combination of R-R features, normalized R-R features and higher order statistics with 83.55% of BCR. Using this selection, the lower class accuracy is of 78.53% for the V class. The addition of any other features to this selection always leads to a lower BCR. Also, these three feature sets are the three most selected sets in the top 20 configurations. In particular, the normalized R-R features are selected in every configuration of the top 15. It is interesting to note that when only R-R and normalized R-R feature sets are included in the model, a BCR of 80.07% can be reached. Also, when HOS are included alone, 76.94% of BCR is

**Table 2.** Grouping of the MIT-BIH labeled heart beat types according to the AAMI standards

Normal beats (N)	Supraventricular ectopic beats (S)	Ventricular beats (V)	ectopic	Fusion beats (F)
Normal beats	Atrial premature beat	Premature contraction	ventricular	Fusion of ventricular and normal beats
Left bundle branch block beats	Aberrated atrial premature beat	Ventricular beats	escape	
Right bundle branch block beats	Nodal (junctional) premature beats			
Atrial escape beats	Supraventricular premature beats			
Nodal (junctional) escape beats				

**Table 3.** Selection of the most interesting results out of the 127 configurations with the weighted SVM model (sorted in decreasing balanced classification rate)

	Feature sets						Results					
	RR	Norm-RR	Seg	Norm-Seg	HBF	HOS	Morph	N (%)	S (%)	V (%)	F (%)	BCR (%)
Top 6	•	•				•		80.00	88.07	78.53	87.63	83.55
						•		86.72	90.68	75.72	79.64	83.19
	•	•	•					77.68	89.70	86.99	77.84	83.05
					•	•		83.35	90.79	74.72	79.38	82.06
	•	•			•	•		80.58	79.24	78.21	89.43	81.87
Bottom 3	•	•						79.18	63.49	86.49	94.33	80.87
			•	•			•	80.00	17.28	72.78	3.35	43.35
					•		•	69.44	2.51	29.96	66.75	42.17
							•	42.76	72.97	37.73	12.63	41.52
	Individuals						•		51.57	93.13	75.16	87.89
			•					68.19	78.69	72.00	64.95	70.96
•								81.70	51.66	45.79	93.04	68.05
		•						83.25	64.96	84.46	1.29	58.49
				•				86.82	3.65	73.63	36.34	50.11
All					•			7.69	71.99	68.98	39.43	47.02
	•	•	•	•	•	•	•	91.57	19.56	84.58	57.73	63.36

obtained. These results indicate that R-R intervals, normalized R-R intervals and HOS features are the most important feature sets to include in the model.

The importance of feature selection is illustrated by the results obtained with the inclusion of all feature sets with a BCR of only 63.36% and an accuracy of 19.56% for the V class. Several feature sets do not seem to serve the classification performances, and some configurations can actually lead to a BCR below 50%, which is clearly not

appropriate. In particular, the normalization of the segmentation features with respect to each patient provides a lower accuracy than their non normalised version. Also, morphological features are not selected in any the top 20 configurations.

## 6.2 Performance Evaluation

The second experiment evaluates the performance of the weighted SVM with the best configuration in Table 3. It is compared to the results of previously reported work that also followed AAMI standards and inter-patient classification. In [3], the best results are obtained by the combination of two weighted linear discriminant (LD) classifier each built on 26 features extracted from one of the two leads of the holter recordings. In [15], a hierarchical model is built from three individual SVM classifiers. Each individual classifier is built on different features, including R-R intervals, HBF and HOS feature sets. The results of the raw SVM with the best configuration when no weights are defined, therefore optimizing an approximation of the accuracy rather than the BCR, are also displayed. Table 4 shows the results.

**Table 4.** Comparison of the weighted SVM using the best feature sets with previously reported works. The results when no weights are set in the SVM model (*raw SVM*) are also displayed.


	Results				
	N	S	V	F	BCR
Weighted SVM	80.0%	88.1%	78.5%	87.6%	83.6%
[3]	87.1%	76.0%	80.3%	89.4%	83.2%
[15]	86.2%	82.6%	80.8%	54.9%	76.1%
Raw SVM	96.5%	0.7%	77.84%	11.8%	46.7%

The weighted SVM model yields better overall results than the hierarchical SVM model. In particular, the results for the S and F pathological classes are significantly improved. The performances are slightly higher but very close the weighted LDA model. Nevertheless, the features included in our model only rely on R spike annotations and do not require the computation of segmentation interval features to reach the same level of performances. The detection of ECG characteristic points needed to extract these features can indeed be a difficult task in noisy signals.

The weights included in the SVM to take the class unbalance into account are also of major importance. If no weights are defined, the BCR obtained by the SVM model decreases to 46.7%, with an accuracy of only less than 1% for class S and of 11.8% for class F which is clearly unacceptable.

## 7 Conclusions

The classification of heart beats is of great importance for clinical applications involving the long-term monitoring of the cardiac function. The two main difficulties are the extraction of discriminative features from the heart beat time series and the class unbalance. For this purpose, this work has motivated and followed the use of:



**Fig. 1.** Percentage of feature set selection in the top 20 best configurations

- AAMI guidelines for the establishment of reliable classifiers and for the evaluation of their relative merits;
- inter-patient rather than intra-patient classification;
- weighted multi-class SVM models to address the class unbalance problem;
- assessment of the relevance of usual feature sets;
- and the balanced classification rate as performance measure.

For the feature set relevance assessment, seven feature sets are considered: R-R intervals, segmentation features, HBF coefficients, higher-order statistics, morphological features, patient-normalized RR intervals and patient-normalized segmentation features. Best results are obtained with the combination of R-R intervals, normalized R-R intervals and high-order statistics with an average class accuracy of 83.55%. Any addition of features to these two sets leads to a lower performance. These results confirm that the choice of the features is a task of major importance, as a bad selection or too many features can lead to unacceptable results.

The results obtained with the weighted SVM model and the best selection of feature sets yield better results than previously reported inter-patient classification models. Furthermore, these performances are obtained with features extracted only from the R spike annotations. Therefore, the model does not require the annotation of the ECG characteristic points which can be very difficult to detect in noisy or pathological beats.

Another important issue for classification of heart beats resides in the class unbalance, which is met with weights being included in the SVM model. Indeed, the average accuracy obtained by the model with our best feature selection decreases from 83.6% to 46.7%, with an accuracy below 12% for class S and F when these weights are removed, leading to rather useless models that are unable to grasp the importance of the pathological cases.

**Acknowledgements.** G. de Lannoy is funded by a Belgian F.R.I.A. grant. This work was partly supported by the Belgian “Région Wallonne” ADVENS and DEEP projects.

## References

1. Association for the Advancement of Medical Instrumentation, Testing and reporting performance results of cardiac rhythm and st segment measurement algorithms. ANSI/AAMI EC38:1998 (1998)
2. Callut, J., Dupont, P.:  $F_\beta$  Support Vector Machines. In: Proceeding of International Joint Conference on Neural Networks, Montreal, Canada, July 31-August 4, pp. 1443–1448 (2005)
3. Chazal, P.D., O'Dwyer, M., Reilly, R.B.: Automatic classification of heartbeats using ecg morphology and heartbeat interval features. *IEEE Transactions on Biomedical Engineering* 51, 1196–1206 (2004)
4. Cheng, W., Chan, K.: Classification of electrocardiogram using hidden markov models. In: Proceedings of the 20th Annual International Conference of the IEEE on Engineering in Medicine and Biology Society, vol. 1, pp. 143–146 (1998)
5. Christov, I., Gómez-Herrero, G., Krasteva, V., Jekova, I., Gotchev, A., Egiazarian, K.: Comparative study of morphological and time-frequency ecg descriptors for heartbeat classification. *Med. Eng. Phys.* 28(9), 876–887 (2006)
6. Clifford, G.D., Azuaje, F., McSharry, P.: *Advanced Methods And Tools for ECG Data Analysis*. Artech House, Inc., Norwood (2006)
7. François, D.: Feature selection. In: Wang, J. (ed.) *Encyclopedia of Data Mining and Warehousing*, 2nd edn. Information Science Reference (2008)
8. Goldberger, A., Amaral, L., Glass, L., Hausdorff, J., Ivanov, P.C., Mark, R., Mietus, J., Moody, G., Peng, C.-K., Stanley, H.: PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. *Circulation* 101(23), e215–e220 (2000)
9. Guyon, I., Gunn, S., Nikravesh, M., Zadeh, L.A.: *Feature Extraction: Foundations and Applications (Studies in Fuzziness and Soft Computing)*. Springer, New York (2006)
10. Huang, Y., Du, S.: Proceeding of the Fourth International Conference on Machine Learning and Cybernetics, Guangzhou, August 18-21, pp. 4365–4369 (2005)
11. Lagerholm, M., Peterson, C., Braccini, G., Edenbrandt, L., Sornmo, L.: Clustering ECG complexes using hermite functions and self-organizing maps. *IEEE Transactions on Biomedical Engineering* 47(7), 838–848 (2000)
12. Melgani, F., Bazi, Y.: Classification of electrocardiogram signals with support vector machines and particle swarm optimization. *IEEE Transactions on Information Technology in Biomedicine* 12(5), 667–677 (2008)
13. Osowski, S., Hoai, L.: Ecg beat recognition using fuzzy hybrid neural network. *IEEE Transactions on Biomedical Engineering* 48(11), 1265–1271 (2001)
14. Osowski, S., Hoai, L., Markiewicz, T.: Support vector machine-based expert system for reliable heartbeat recognition. *IEEE Transactions on Biomedical Engineering* 51(4), 582–589 (2004)
15. Park, K., Cho, B., Lee, D., Song, S., Lee, J., Chee, Y., Kim, I., Kim, S.: Hierarchical support vector machine based heartbeat classification using higher order statistics and hermite basis function. In: *Computers in Cardiology*, pp. 229–232 (2008)
16. Platt, J.C.: Fast training of support vector machines using sequential minimal optimization. In: Scholkopf, B., Burges, C., Smola, A. (eds.) *Advances in Kernel Methods*. MIT Press, Cambridge (1999)
17. Soria, M., Martinez, J.P.: An ECG classification model based on multilead wavelet transform features. In: *Computers in Cardiology*, pp. 105–108 (2007)
18. Vapnik, V.N.: *The Nature of Statistical Learning Theory (Information Science and Statistics)*. Springer, Heidelberg (1999)

# The Investigation and Clinical Significance of Resonance in the Heart Rate and Vascular Tone Baroreflexes

Evgeny G. Vaschillo, Bronya Vaschillo, Jennifer F. Buckman,  
Robert J. Pandina, and Marsha E. Bates

Center of Alcohol Studies, Rutgers, The State University of New Jersey  
607 Allison Road, Piscataway, New Jersey 08854, U.S.A.  
{evaschil,bvaschil}@rci.rutgers.edu

**Abstract.** The heart rate (HR) and vascular tone (VT) baroreflexes are control systems with negative feedback. As closed-loop control systems with delays, they possess resonance features at approximately 0.1 Hz and 0.03 Hz, which correspond to a ~5-s delay in the blood pressure (BP) response to HR changes and a ~15-s delay in VT response to BP changes, respectively. Unlike a single impact on the cardiovascular system (CVS) that elicits HR, BP, and VT oscillatory responses that fade over time, rhythmical stimulation of the CVS at 0.1 Hz or 0.03-Hz produces steady HR, BP, and VT oscillations with significantly higher amplitudes compared to stimulation at other frequencies. We posit that these baroreflex resonances are integral to the vast autonomic variability that underlies efficient and effective homeostatic reflexes. This paper characterizes the phenomenon of resonances in the CVS using a classic engineering approach and the ability to employ these resonances for clinical applications.

**Keywords:** Baroreflex, Closed-loop control system, Resonance frequency, HRV biofeedback.

## 1 Introduction

The arterial baroreflex is a mechanism that controls blood pressure (BP). It is triggered by a shift in BP and leads to changes in heart rate (HR) and vascular tone (VT) that counteract the BP shift. Typically, these HR and VT baroreflex systems are modelled using the classic “control system theory” approach, which conceptualizes the baroreflex as a closed-loop control system with negative feedback. These models are consistent with the premise that a critical function of the baroreflex system is to buffer BP oscillations [1, 2, 3, 4], but at first glance, are inconsistent with a system that demonstrates resonance property frequencies [5, 6, 7]. By definition, though, a closed-loop control system with a delay possesses resonance features and thus the classic control system theory can be used to model both the baroreflex’s buffering functions and its resonance properties.

A closed-loop system always possesses resonance properties because all biological or technical control systems have delays associated with inertia. When creating a stabilizing technical system with a closed-loop, the delay is manipulated so that the

resonance frequency falls far outside of the operating frequency range. In the case of the baroreflex system, there are two resonance frequencies within its very narrow ( $\sim 0.01$ - $0.5$  Hz) operating range.

The baroreflex system in humans demonstrates resonance properties at frequencies of about 0.1 Hz and 0.03 Hz [7, 8, 9]. In the HR baroreflex closed-loop, a shift in BP causes a compensatory HR response that is delayed for approximately 5 seconds. In the VT baroreflex closed-loop, the compensatory response of the vasculature is delayed for approximately 10-15 seconds [8, 5]. These delays of 5 and 15 seconds coincide with resonance oscillations at 0.1 and 0.03 Hz (oscillation periods are equal to twice the value of the delay).

Despite the current view that the main role of the baroreflex is to buffer BP oscillations and “that resonance is the price to be paid for effective buffering at other frequencies” [7], we consider resonance properties of the baroreflex as essential elements for the regulation of autonomic and central nervous system functions. We posit that the HR, BP, and VT resonances in the baroreflex systems are integral for the vast autonomic variability that underlies efficient and effective homeostatic reflexes. This is supported by evidence that baroreflex resonance properties can act to amplify adaptive responses to internal and external stimuli and buffer stress and emotional reactivity through the initiation of a cascade of neurobiological events that produces a generalized inhibitory effect on the brain [10, 11, 12].

This paper presents the results of our investigations of baroreflex resonance features using a classic engineering approach. In these studies, the baroreflex is considered the main source of resonance oscillation in cardiovascular functions. Accordingly, we examine resonance in the cardiovascular system (CVS) by stimulating it with paced breathing, rhythmical muscle tension, emotional pictures and auditory cues. Based on the importance of autonomic variability and the dependence of this variability on the frequency of stimulation, we propose that better understanding of the resonance properties of the HR and VT baroreflexes can significantly contribute to the development of novel therapeutic methods to treat physical and mental illnesses.

## 2 0.1 Hz Resonance in the CVS

Gatchel and Lang [13] and Lang et al. [14] found that HR responses to a single stimulus tended to last approximately 10 seconds and have a triphasic waveform, consisting of a small initial HR deceleration, a larger mid-interval acceleration, and a final deceleration. Onset of any stimulus - visual [14] or acoustical [15], long (a few seconds) [15] or brief (a few tens of milliseconds) [16] - caused the same HR waveform response. Thus, the triphasic waveform of the instantaneous HR response appears to result directly from the inherent resonance properties of the HR baroreflex closed-loop. We hypothesize that this basic feature of the HR baroreflex closed-loop can serve as the foundation for eliciting stable resonance oscillations in HR, BP, and VT using rhythmical 0.1 Hz stimulation.


### 2.1 Prior Studies: 0.1 Hz Resonance in the CVS Caused by Respiration

It is well known that respiratory activity continually perturbs the cardiovascular system and that breathing modulates HR due to respiratory sinus arrhythmia. Clynes



[17] showed that a single inhalation or exhalation elicits nearly identical triphasic HR waveform responses and that 0.1 Hz breathing caused high amplitude oscillations in HR. Angelone and Coulter [18] further explored this phenomenon by calculating the amplitude and phase transfer functions (with respiration as the input and HR as the output) for a single subject who performed paced breathing exercises at different frequencies (0.01 – 0.5 Hz range). They found that the 0.1 Hz breathing produced the highest HR oscillation and defined this phenomenon as resonance in the CVS. These early studies were used to guide the development of a heart rate variability (HRV) biofeedback procedure based on paced resonance frequency breathing as a novel approach for correcting abnormal autonomic regulation [19].

Our HRV biofeedback procedure teaches participants to breathe easily and naturally at a rate of ~ 6 times per minute; this breathing rate is in contrast to the usual respiratory rate of 12-18 breaths per minute (0.2-0.3 Hz). The procedure consists of in-laboratory training sessions, where the individual resonance frequency of each participant is defined and participants are taught to breathe slowly but not too deeply (i.e., to avoid hyperventilation and support normal minute ventilation) as well as home practice sessions, where participants are instructed to breathe at their resonance frequency for approximately 40 minutes per day over a 10-week period. HRV biofeedback has been demonstrated to normalize autonomic regulation, as measured by increased baroreflex gain and peak expiratory flow [20], and is efficacious for the treatment of asthma [21], major depression [22], fibromyalgia [23], neurosis [24], and hypertension [25]. The therapeutic effects were achieved by triggering high-amplitude oscillations in multiple cardiovascular systems that activate and train homeostatic reflexes [20, 21].



**Fig. 1.** Oscillations in physiological functions during HRV biofeedback. Paced breathing at 0.1 Hz triggers resonance oscillations in cardiovascular functions. Oscillations were observed in (A) respiration, (B) heart rate, and (C) vascular tone. The left panels show representative raw data traces. The right panels show the corresponding spectra generated from these traces.

We recently calculated the amplitude and phase transfer functions of the HR control system for eight participants involved in an HRV biofeedback study. We found that the  $0^\circ$  phase shift between respiration and HR curves occurred precisely at an individual's unique resonance frequency [26]. In this study, each participant's unique


resonance frequency was identified between 0.075 Hz and 0.107 Hz. Further, we noted that breathing at one's resonance frequency produces high-amplitude oscillations in HR, which, through the baroreflex, spread to other functions, such as BP and VT (Fig. 1). Thus, the therapeutic effects of HRV biofeedback appear to stem from the elicited high amplitude oscillations in HR, BP, VT, and other autonomic functions that retrain and tone homeostatic reflexes. Activation of the baroreceptors by these oscillations also activates inhibitory processes in the brain, thereby dampening stress. Taken together, these data suggest that daily "exercise" of autonomic functions associated with normalized autonomic regulation can restore sympathetic-vagal balance [21] and buffer patients from some negative influences of stress.

## 2.2 Study 1: 0.1 Hz Resonance in the CVS Caused by Rhythmical Muscle Tension

Respiration is a natural rhythmical stimulator of the CVS; however, there are other effective methods to stimulate the CVS at its resonance frequency. For example, the CVS functions adaptively to react to physical load. France, France, and Patterson [27] developed the Rhythmical Skeletal Muscle Tension (RSMT) technique to lower risk of a vasovagal reaction (fainting). RSMT has been successfully employed to avert fainting episodes that often occur during blood collection procedures and may discourage people from donating blood. It has also been successfully used to treat patients with blood and injury phobias. In their study, healthy young adults performed the RSMT task at a frequency of 0.1 Hz and demonstrated significant increases in HR, systolic and diastolic BP, and cerebral oxygen [27]. High amplitude 0.1 Hz oscillation in HR, BP, and VT were also observed. In the present study, we assessed the ability of rhythmical paced muscle tension (muscle tense-release cycles) at a frequency of 0.1 Hz to trigger resonance in the CVS.

**Method.** Sixteen young healthy participants between, 17 and 26 years old (9 female, 7 male), performed four 3.5-minute tasks in random order (30 second inter-task interval): a 0.1 Hz paced breathing task and then three paced muscle tension tasks at frequencies of 0.05, 0.1, and 0.2 Hz. Participants were seated in a comfortable armchair in front of a computer screen with their legs extended and supported parallel to the floor. During the paced muscle tension tasks, participant tensed their skeletal muscles when the computer screen turned red and relaxed their muscles when the screen changed to green. ECG and finger pulse were recorded during all tasks. Beat-to-beat HR, pulse transit time (PTT), and their respective Fourier spectra were calculated for each task. PTT was used to estimate vascular tone (shorter PTT corresponds to higher VT). HR and VT reactions in each task were estimated from the power of the spectra at the tested frequencies.

**Results.** The 0.05, 0.1, and 0.2 Hz muscle tension manipulations produced HR and VT oscillations at corresponding frequencies in all participants; however, only rhythmical 0.1 Hz muscle tension caused high amplitude HR oscillations like those observed with 0.1 Hz breathing (see Fig. 2). Averaged across all participants, HR and VT reactions to muscle tension at 0.1 Hz were 4-6 times higher than at 0.2 Hz or 0.05 Hz. Interestingly, although average HR oscillatory reaction to the 0.1 Hz muscle tension task was significantly lower than to the 0.1 Hz breathing task, average VT oscillatory reaction to the 0.1 Hz muscle tension task was significantly higher.



**Fig. 2.** Paced muscle tension at (A) 0.05 Hz [20 second period], (B) 0.1 Hz [10 second period], and (C) 0.2 Hz [5 second period] differentially affects heart rate. Tense-release cycles at 0.1 Hz produce maximal oscillations. Left panels show raw heart rate data (in beats per minute, bpm). Right panels show the corresponding spectral data.

**Discussion.** Results confirmed that rhythmical 0.1 Hz muscle tension tasks can trigger resonance in the CVS. Although oscillations were robust in both the HR and VT spectra, the muscle tension task generated smaller HR oscillations, but larger VT oscillations, than those elicited by the breathing task. These findings parallel those reported by Lehrer et al. [28]. The larger amplitude oscillations in the VT spectrum may be related to the physical load of muscle tension, which increases mean HR and consequently depresses HRV. These data suggest that BP and VT oscillatory activity act as compensatory reactions to dampened HRV.

**Clinical Utility.** The ability of the rhythmical muscle tension to produce high amplitude oscillations at this resonance frequency makes it potentially valuable as a clinical tool. A biofeedback procedure that includes systematic, every day use of paced rhythmical muscle tension may produce cumulative and long lasting effects on health in much the same way as HRV biofeedback. This idea builds on the RSMT technique developed by France, France, and Patterson [27], which was originally intended to be a single session performed immediately prior to an event that could induce a negative physiological reaction.

We propose that the mechanism by which the 0.1 Hz RSMT procedures prevent vasovagal reactions is high amplitude oscillations, which activate regulatory processes that, in turn, balance autonomic functioning, modulate neural inhibition, and buffer the body from stress. Accordingly, we speculate that other kinds of 0.1 Hz RSMT procedures, such as those that incorporate everyday use, can be exploited by researchers and clinicians for the development of novel approaches to correct abnormal autonomic regulation. Such procedures may prove especially useful in the rehabilitation process of patients following a cerebral stroke or myocardial infarction, in treatments where physical exercises are prescribed [29], or in sport medicine.


### 2.3 Study 2: 0.1 Hz Resonance in the CVS Caused by Emotional Pictures Cues

The CVS actively participates in emotional regulation, and emotions strongly affect cardiovascular functions. Picture cues that elicit emotional reactions have been reported to produce the common triphasic HR response [15]. The magnitude of this triphasic response, particularly the accelerative leg, appears to accurately discriminate picture valence [13, 14]. In accordance with our model, instigating rhythmical reactions at 0.1 Hz should produce resonance oscillation in HR, and the oscillation amplitude should discriminate the degree of emotional arousal. We previously showed that viewing emotional picture cues elicits strong 0.1 Hz HRV reactions in young adults who span the continuum of alcohol and drug use behavior [30, 31, 32]. In the present study, blocks of emotionally arousing picture cues were presented to healthy young adults at a frequency of 0.1 Hz and the relationship of resonance amplitude to the degree of emotional arousal was assessed.

**Method.** Seventy-six young participants, between 18 and 29 years old, were individually tested. Each participant viewed six categories of 30-picture blocks (negative emotional, positive emotional, and neutral, as well as alcohol, marijuana, and club drugs [ecstasy, cocaine]) [33, 34, 35]. Pictures were presented for 5 seconds with a 5-second inter-picture interval, resulting in a 0.1 Hz picture presentation frequency. The interval between each picture block was 30 seconds. Pictures were presented on a 75-cm LCD TV (View Sonic N3000W). ECG and finger pulse were recorded during all tasks. Beat-to-beat HR and PTT, and their Fourier spectra, were calculated for each picture cue block. Reaction of the CVS to the picture cue blocks was estimated by the power of HR and VT spectra at frequency of 0.1 Hz (i.e., the 0.1 Hz HR index and the 0.1 Hz VT index, respectively). To evaluate the sensitivity of the 0.1 Hz HR index for estimating emotional valence, common HRV indices (total HRV, high frequency (HF) HRV, and low frequency (LF) HRV) were also calculated.

**Results.** Picture presentation at a frequency of 0.1 Hz caused high amplitude HR and VT oscillations at the 0.1 Hz resonance frequency of the CVS. The resonance in VT, however, was less prominent than that in HR. Moreover, as we previously reported [31], the 0.1 Hz HR index detected reactivity to the picture cue blocks more sensitively than the other, commonly used HRV indices (total HRV, HF HRV, LF HRV). For example, a participant's reactions to negative versus neutral picture cue blocks often appeared different using the 0.1 Hz HRV index, but not the common HRV indices.

In general, participants' average 0.1 Hz HR index response was lowest for neutral picture cues and highest for negative emotional picture cues. The average 0.1 Hz HR index responses to positive emotional, alcohol, marijuana, and club drugs picture cues did not generally appear different. There were individual differences in the response patterns of 0.1 Hz HR indices across the picture blocks. For example, one participant strongly reacted to negative emotional and alcohol picture stimuli, but weakly reacted to positive emotional and marijuana cues. Conversely, another participant showed high 0.1 Hz HR index responses to club drug (ecstasy and cocaine) cues, a moderate response to negative and positive emotional cues, and a weak response to all others cue blocks (see Fig. 3).



**Fig. 3.** The RRI spectra from one participant in response to pictures cues presented at 0.1 Hz. Responses to emotional picture cues (neutral, negative, positive) as well as to appetitive cues (alcohol, marijuana, club drugs) were assessed. The power of the spectra at 0.1 Hz likely reflects the strength of emotions caused by picture cue blocks.

**Discussion.** Emotional arousal modulated by rhythmical picture cue exposure at a frequency of 0.1 Hz was reflected by high amplitude resonance oscillations in HR and other cardiovascular functions. Because it seems unlikely that, by itself, picture cue presentation rate affects emotional arousal, we posit that the amplified reactivity observed when pictures were presented at 0.1 Hz is directly related to pairing the presentation frequency with the resonance frequency of the CVS.

Our approach of using the CVS' 0.1 Hz resonance property to evaluate the degree of emotional arousal in response to stimuli is similar to the engineering approach of measuring weak oscillatory signals. An engineer tunes the resonance frequency of a measurement device to the main frequency of a weak signal. Tuning to the resonance frequency does not change the value of the signal, but rather enhances the sensitivity of the measurement device. Here, we enhance the sensitivity of the CVS as a measurement device by modulating emotional arousal at a frequency of 0.1 Hz. In this study, some individuals reacted more strongly to emotional picture cues, whereas others reacted more strongly to drug-related picture cues. This provides indirect evidence that the salience of the stimuli has a significant impact on the amplitude of resonance oscillations. Future research linking other components of an individual's emotional state and substance use histories to their resonance oscillations represents an important step in understanding the dynamic interplay between physiological and cognitive indices of emotional regulation.

**Clinical Utility.** Investigating the resonance properties of the CVS and using an engineering approach to study weak oscillatory signals may be useful for developing a method for estimating emotional arousal, particularly, in relation to the diagnosis of psycho-emotional disorders. In addition, the use of paced visual stimulation may prove useful for the treatment and rehabilitation of a variety of disorders since visual cues presented at a frequency of 0.1 Hz appear to cause high-amplitude HR

oscillation similar to those caused by resonance-paced breathing and muscle tension. An advantage of picture cue stimulation is that the content of the stimuli can be easily manipulated; thus, it is conceivable that rhythmical visual stimulation using cues with specific cognitive content (e.g., emotional cues to induce altered mood states or drug-related cues to manipulate craving reactions) may open new doors to treatment applications in the mental health and addictions field.

### 3 0.03 Hz Resonance in the CVS


There are two interconnected branches of the baroreflex system. The 0.1 Hz resonance frequency explored in the aforementioned studies reflects the functioning of the HR baroreflex. The VT baroreflex closed-loop controls BP and participates in modulating the coordinated actions of the central and autonomic nervous systems in a manner similar to the HR baroreflex, but is much less studied. An important difference between the HR and VT baroreflexes is that the latter functions by modulating the stretch of blood vessel walls and operates in a lower frequency range [36]. Based on our prior research [8, 37], we hypothesize that the VT baroreflex closed-loop possesses a resonance property at a frequency of  $\sim 0.03$  Hz.

As mentioned above, we linked the 10-s HR triphasic waveform stimulus response [14] with resonance in HR baroreflex closed loop. In the following studies, we sought to link the  $\sim 30$ -s HR triphasic waveform response to intense auditory stimulation [38] to resonance in the VT baroreflex closed loop. This is in keeping with studies that described the  $\sim 30$ -s HR triphasic waveform response as a defensive reflex [39, 40] that unifies vascular responding, causing reactions such as concomitant vasoconstriction in the forehead and hands [41]. We performed two studies to investigate whether external stimuli can elicit a 0.03 Hz oscillatory HR response in addition to the 0.1 Hz response. In order to detect a clear HR response at this frequency, very strong stimuli were used.

#### 3.1 Study 1: 0.03 Hz Resonance in the CVS Caused by Aversive Sound

Intense sound has been shown to cause a  $\sim 30$ -s triphasic response in HR [39, 40]. However, these studies have not characterized the underlying mechanisms that drive these low frequency CVS oscillations. We hypothesize that the VT baroreflex drives this  $\sim 30$ -s triphasic response. Identification of 0.03 Hz oscillations that fade over time would suggest a resonance system consistent with the VT baroreflex, and that the first period of these oscillations is the  $\sim 30$ -s triphasic response. In this study, HR response to aversive sound was assessed.


**Method.** Seventeen adult participants were exposed to 8 synthetic sounds, which were chosen as the most unpleasant from 362 sounds created in the lab. Participants sat in a sound attenuated room and listened to each sound at 82 dB(A) or 92 dB(A) for 2 minutes (with a 30-s inter-stimulus interval) from four equidistant speakers. ECG was collected and beat-to-beat HR curves were calculated (see Fig. 4).



**Fig. 4.** Exposure to aversive sounds (A) causes changes in heart rate (B). The onset of a 82 db(A) or 92 db(A) sound causes oscillations at 0.03 Hz.

**Results.** Onset of unpleasant sounds typically caused a long-duration (~0.03 Hz) oscillatory HR response. Sound offset occasionally caused the same response. These responses appeared to contain overlapping 10-s triphasic waveforms which lasted for more than 30 seconds, and then gradually faded (see Fig. 5). Two types of 30-s triphasic responses were found: one with an initial HR acceleration and one with an initial HR deceleration (Fig. 5, Table 1).

**Discussion.** We were able to replicate earlier studies that demonstrated a ~30-s triphasic HR response to aversive auditory stimuli. As hypothesized, this waveform was a component of a longer duration oscillation that reacted to the onset (and sometimes to offset) of unpleasant sounds and faded over time. These results offer support for the involvement of VT baroreflex in controlling low frequency HR oscillations. In addition, we observed two distinct types of triphasic responses. In most cases, the HR response began with an acceleration (Fig. 5, left panel); however, in some cases, the HR response began with a deceleration (Fig. 5, right panel). Although the reason for the waveform differences is not clear, it may be related to the level of aversion experienced by the participant. Future studies that assess waveform characteristics following stimuli that range in aversiveness are warranted.



**Fig. 5.** Examples of ~0.03 Hz oscillatory HR reactions to aversive sound (30-s triphasic waveform HR response). The left panel illustrates an HR response with an initial HR acceleration. The right panel shows an HR response with an initial HR deceleration. In both panels, extremum 1 occurred approximately 15 seconds after sound onset and extremum 2 occurred approximately 30 seconds after sound onset.

**Table 1.** Two types of ~30 second triphasic waveform HR responses averaged across all participants and sound characteristics

Response Type	Extremum 1 (Beats/Minute)	Extremum 2 (Beats/Minute)	Oscillatory Period (Seconds)	Oscillatory Frequency (Hz)
TYPE 1	7.86±0.58	-5.03±0.79	28.62 ±1.23	0.035±0.006
TYPE 2	- 8.34±0.78	4.14±1.1	26.74 ±1.86	0.037±0.007

*Note:* Data are presented as mean ± standard error.

**Clinical Utility.** Based on the utility of the 0.1 Hz resonance in developing novel medical applications for treating various physical and mental disorders, it is likely that exploration of the 0.03 Hz resonance could prove valuable in much the same way. However, little is known about this resonance frequency. This study is an important step in understanding the significance of the 0.03 Hz resonance in the VT baroreflex and its relation to the ~30-s triphasic waveform response. It also offers insight into the capacity of the CVS for long-period oscillatory responses to acoustic stimuli.

### 3.2 Study 2: 0.03 Hz Resonance in the CVS Caused by Negative Emotional Pictures

We hypothesized that if the ~0.03 Hz resonance oscillations observed in response to acoustic stimulation reflected a CVS resonance frequency, then these oscillations should be reproducible using other stimulus modalities. Viewing picture stimuli that carry a high negative emotional valence was expected to provoke a long-period oscillatory HR response that was similar to that observed with aversive sound. In this study, cardiovascular responses were measured in young, healthy individuals while they viewed highly negative pictures.

**Method.** See section 2.3.


**Results.** Similar to listening to aversive auditory stimuli, viewing very unpleasant, negatively-valenced pictures (e.g., plane crashes, blood, violence) caused a strong ~0.03 Hz oscillatory HR response (see Fig 7).

**Discussion.** Strong stimuli from various modalities elicited oscillatory HR responses at a frequency of about 0.03 Hz, which overlapped with the 0.1 Hz response. This effect may be the result of resonance in the VT baroreflex closed-loop. In these studies, we employed strong visual and auditory stimuli and observed 0.03 Hz oscillatory responses. Less aversive stimuli should also be capable of eliciting these slower oscillations; however, these oscillations may be masked by those associated with HRV. Future studies are needed to assess the malleability of the 0.03 Hz resonance and its utility in clinical applications.

**Clinical Utility.** To be of value for clinical application, it is necessary to develop experimental methods that can reliably and easily produce stable high amplitude 0.03 Hz oscillation in the CVS. The use of very strong negative stimuli, while capable of producing such oscillations, may not prove clinically useful because of the possibility



of negative psychological side effects. However, the use of highly negative stimuli was necessary to first demonstrate that a 0.03 Hz resonance can be triggered in the CVS. We speculate that pleasant auditory stimuli, such as music, that includes 0.03 Hz rhythmical components or the presentation of positively-valenced visual cues at 0.03 Hz should also cause stable oscillations. In addition, exploration of paced breathing or rhythmical muscle tension techniques at  $\sim 0.03$  Hz warrants further study as these procedures may also be capable of triggering therapeutic oscillation in the same way as 0.1 Hz stimulation triggers them. Elements of such stimulation can be found in Eastern health procedures (e.g., Yoga, Tai Chi).



**Fig. 6.** Examples of  $\sim 0.03$  Hz oscillatory HR reactions to negatively-valenced picture cues. The top panel shows the beat-to-beat intervals (RRI) in seconds. The bottom panel shows the corresponding spectral data. Note the large amplitude peak at 0.027 Hz that reflects the VT baroreflex.

## 4 Discussion

Two resonance frequencies, 0.1 and 0.03 Hz, have been identified in the CVS. These resonances are thought to reflect two interdependent closed-loop systems with delays, namely the HR and VT baroreflex systems, and contribute independently to the overall resonance properties of the CVS. The baroreflex system is well appreciated for its ability to buffer perturbations in BP [1, 2, 3, 4, 5, 7]. Nonetheless, this system has a very narrow operative frequency range ( $\sim 0.01$ -0.5 Hz), with two resonance frequencies inside this range. This suggests that the ability of the baroreflex system to stabilize the CVS is limited, and that the primary functions of the HR and VT baroreflex closed-loop systems may be more related to their capacity for *variability* rather than their efficiency at stabilizing the system. The high frequency boundary is defined by inertia of blood mass and slow changes in vessel tone, while the low frequency boundary is defined by the differentiative property of the baroreceptors (i.e., they react only to the speed of BP changes).

The long-term aim of investigating the dynamic properties of the HR and VT baroreflexes is to develop new therapeutic methods for treating disorders associated with the dysregulation of the autonomic and central nervous systems. HRV biofeedback is capable of harnessing the resonances within the CVS to promote health benefits. These therapeutic effects have been linked to the generation of generalized high-amplitude oscillations in autonomic functions that are elicited by biofeedback procedures [24, 20, 21] that act to restrain autonomic reflexes. The systematic retraining of autonomic reflexes normalizes and improves autonomic regulation. The studies described here show that breathing, visual cues, and muscle tension management of the 0.1 Hz resonance may be useful for the treatment of various unhealthy physical and mental conditions. We believe that novel therapeutic interventions involving the VT baroreflex and its resonance at 0.03 Hz through passive (e.g., viewing rhythmical picture stimuli) or active (e.g., paced muscle tension) tasks may also be beneficial, but additional investigations are needed.

## 5 Conclusions

Classic control system theory applied to the investigation of physiological systems can be a useful tool for medical practice. An engineering approach offers the opportunity to create simple, clinically useful stimulation procedures, which may be used to enhance an individual's regulatory capacity and thus open new doors to treatment applications in the mental health and addictions field.

**Acknowledgements.** This research was supported by grants and contracts from the National Institutes of Health (R01 AA015248, K02 AA00325, K01 AA017473, P20 DA017552, and HHSN275201000003C).

## References

1. Just, A., Wittmann, U., Nafz, B., Wagner, C.D., Ehmke, H., Kirchheim, H.R., Persson, P.B.: The blood pressure buffering capacity of nitric oxide by comparison to the baroreceptor reflex. *Am. J. Physiol. Heart Circ. Physiol.* 527, H521–H527 (1994)
2. Christou, D.D., Jones, P.P., Seals, D.R.: Baroreflex buffering in sedentary and endurance exercise-trained healthy men. *Hypertension* 41, 1219–1222 (2003)
3. Jones, P.P., Christou, D.D., Jordan, J., Seals, D.R.: Baroreflex buffering is reduced with age in healthy men. *Circulation* 107, 1770–1774 (2003)
4. Jordan, J., Tank, J., Shannon, J.R., Diedrich, A., Lipp, A., Schroder, C., Arnold, G., Sharma, A.M., Biaggioni, I., Robertson, D., Luft, F.C.: Baroreflex buffering and susceptibility to vasoactive drugs. *Circulation* 105, 1459–1464 (2002)
5. Magosso, E., Biavati, V., Ursino, N.: Analysis of cardiovascular instability by mathematical model of baroreflex control. In: 2001 Proceeding of the 23rd Annual EMBS International Conference, pp. 596–599 (2001)
6. Julien, C.: The enigma of Mayer waves: Facts and models. *Cardiovasc. Res.* 70, 12–21 (2006)
7. van de Vooren, H., Gademan, M.G.J., Swenne, C.A., TenVoorde, B.J., Schalijs, N.J., van der Wall, E.E.: Baroreflex sensitivity, blood pressure buffering, and resonance: what are the links? Computer simulation of healthy subjects and heart failure patients. *J. Appl. Physiol.* 102, 1348–1356 (2007)

8. Vaschillo, E., Lehrer, P., Rische, N., Konstantinov, M.: Heart rate variability biofeedback as a method for assessing baroreflex function: a preliminary study of resonance in the cardiovascular system. *Appl. Psychophysiol. Biofeedback* 27, 1–27 (2002)
9. Vaschillo, E., Vaschillo, B., Lehrer, P.: Characteristics of resonance in heart variability stimulated by biofeedback. *Appl. Psychophysiol. Biofeedback* 31, 129–142 (2006)
10. Dworkin, B.R., Elbert, T., Rau, H., Birbaumer, N., Pauli, P., Droste, C., Brunia, C.H.: Central effects of baroreceptor activation in humans: attenuation of skeletal reflexes and pain perception. *Proc. Natl. Acad. Sci. USA* 91, 6329–6333 (1994)
11. Nyklicek, I., Wijnen, V., Rau, H.: Effects of baroreceptor stimulation and opioids on the auditory startle reflex. *Psychophysiology* 42, 213–222 (2005)
12. Yasumasu, T., Reyes del Paso, G., Takahara, K., Nakashima, Y.: Reduced baroreflex cardiac sensitivity predicts increased cognitive performance. *Psychophysiology* 43, 41–45 (2006)
13. Gatchel, R.J., Lang, P.J.: Accuracy of psychological judgments and physiological response amplitude. *J. Exp. Psychol.* 98, 175–183 (1973)
14. Lang, P.J., Greenwald, M.K., Bradley, M.M.: Looking at pictures: Affective, facial, visceral, and behavioral reactions. *Psychophysiology* 30, 261–273 (1993)
15. Bradley, M.M., Lang, P.J.: Affective reactions to acoustic stimuli. *Psychophysiology* 37, 204–215 (2000)
16. Codispoti, M., Bradley, M.M., Lang, P.: Affective reactions to briefly presented pictures. *Psychophysiology* 38, 474–678 (2001)
17. Clynes, M.: Respiratory sinus arrhythmia: laws derived from computer simulation. *J. Appl. Physiol.* 15, 863–874 (1960)
18. Angelone, A., Coulter Jr., N.A.: Respiratory sinus arrhythmia: A frequency depended phenomenon. *J. Appl. Physiol.* 19, 479–482 (1964)
19. Lehrer, P.M., Vaschillo, E., Vaschillo, B.: Resonant frequency biofeedback training to increase cardiac variability: Rationale and manual for training. *Appl. Psychophysiol. Biofeedback* 25, 181–192 (2000)
20. Lehrer, P.M., Vaschillo, E., Vaschillo, B., Lu, S.E., Eckberg, D.L., Edelberg, R., Shih, W.J., Lin, Y., Kuusela, T.A., Tahvanainen, K.U.O., Hamer, R.: Heart rate variability biofeedback increases baroreflex gain and peak expiratory flow. *Psychosomatic Med.* 65, 796–805 (2003)
21. Lehrer, P., Vaschillo, E., Vaschillo, B., Lu, S., Scardella, A., Siddique, M., Habib, R.: Biofeedback treatment for asthma. *Chest* 126, 352–361 (2004)
22. Karavidas, M.K., Lehrer, P.M., Vaschillo, E., Vaschillo, B., Marin, H., Buyske, S., Radvanski, D., Hasset, A.: Preliminary results of an open label study of heart rate variability for the treatment of major depression. *Appl. Psychophysiol. Biofeedback* 32, 19–30 (2007)
23. Hasset, A.L., Radvanski, D.C., Vaschillo, E., Vaschillo, B., Sigal, L., Karavidas, M., Buyske, S., Lehrer, P.M.: A pilot study of the efficacy of heart rate variability biofeedback in patients with fibromyalgia syndrome. *Appl. Psychophysiol. Biofeedback* 32, 1–10 (2007)
24. Chernigovskaya, N.V., Vaschillo, E.G., Rusanovsky, V.V., Kashkova, O.E.: Instrumental autotraining of mechanisms for cardiovascular function regulation in treatment of neurotics. *The SS Korsakov's J. Neuropathol. Psychia.* 90, 24–28 (1990)
25. McCraty, R., Atkinson, M., Tomasino, D.: Impact of a workplace stress reduction program on blood pressure and emotional health in hypertensive employees. *J. Altern. Complement. Med.* 9, 355–369 (2003)
26. Vaschillo, E.G., Vaschillo, B.: Transfer function of the heart rate control system with respiratory input - The classical engineering approach. *BIOSIGNALS*, 233–238 (2009)

27. France, C.R., France, J.L., Patterson, S.M.: Blood pressure and cerebral oxygenation responses to skeletal muscle tension: a comparison of two physical maneuvers to prevent vasovagal reactions. *Clin. Physiol. Functional Imaging* 26, 21–25 (2006)
28. Lehrer, P., Vaschillo, E., Trost, Z., France, C.: Effects of rhythmic muscle tension at 0.1 Hz on cardiovascular resonance and the baroreflex. *Biol. Psychol.* 81, 24–36 (2009)
29. Buch, A.N., Coote, J.H., Townend, J.N.: Mortality, cardiac vagal control and physical training—what’s the link? *Exp. Physiol.* 87, 423–435 (2002)
30. Mun, E.Y., von Eye, A., Bates, M.E., Vaschillo, E.: Finding groups using model-based cluster analysis: Heterogeneous emotional self-regulatory processes and heavy alcohol use risk. *Dev. Psychol.* 44, 481–495 (2008)
31. Vaschillo, E.G., Bates, M.E., Vaschillo, B., Lehrer, P., Udo, T., Mun, E.Y., Ray, S.: Heart rate variability response to alcohol, placebo, and emotional picture cue challenges: Effects of 0.1 Hz stimulation. *Psychophysiology* 45, 847–858 (2008)
32. Buckman, J. F., White, H. R., Bates, M. E.: Psychophysiological reactivity to emotional picture cues two years after college students were mandated for alcohol interventions. *Addict. Behav.* (in press)
33. Lang, P.J., Bradley, M.M., Cuthbert, B.N.: International affective picture system (IAPS): Instruction manual and affective ratings (Technical Report A-4). The Center for Research in Psychophysiology, University of Florida (1999)
34. Stritzke, W.G., Breiner, M.J., Curtin, J.J., Lang, A.R.: Assessment of substance cue reactivity: Advances in reliability, specificity, and validity. *Psychol. Addict. Behav.* 18, 148–159 (2004)
35. Tapert, S.F., Cheung, E.H., Brown, G.G., Frank, L.R., Paulus, M.P., Schweinsburg, A.D., Meloy, M.J., Brown, S.A.: Neural responses to alcohol stimuli in adolescents with alcohol use disorder. *Arch. Gen. Psychiatry* 60, 727–735 (2003)
36. Aljuri, N., Marini, R., Cohen, R.J.: Test of dynamic closed-loop baroreflex and autoregulatory control of total peripheral resistance in intact and conscious sheep. *Am. J. Physiol. Heart Circ. Physiol.* 287, H2274–H2286 (2004)
37. Vashchillo, E.G., Zingerman, A.M., Konstantinov, M.A., Menitsky, D.N.: An investigation of the resonance characteristics of the cardiovascular system. *Human Physiol.* 9, 257–265 (1983)
38. Turpin, G., Siddle, D.A.: Cardiac and forearm plethysmographic responses to high intensity auditory stimulation. *Biol. Psychol.* 6, 267–281 (1978)
39. Vila, J., Guerra, P., Muñoz, M.A., Vico, C., Viedma-del Jesús, M.I., Delgado, L.C., Perakakis, P., Kley, E., Mata, J.L., Rodríguez, S.: Cardiac defense: From attention to action. *Int. J. Psychophysiol.* 66, 169–182 (2007)
40. Ramirez, I., Sanchez, M.B., Fernandez, M.C., Lipp, O.V., Vila, J.: Differentiation between protective reflexes: Cardiac defense and startle. *Psychophysiology* 42, 732–739 (2005)
41. Sokolov, E.N.: Perception and the conditioned reflex. Pergamon Press, Elmsford (1963)

# Glottal Inverse Filtering of Speech Based on Homomorphic Prediction: A Cepstrum-Based Algorithm Not Requiring Prior Detection of Either Pitch or Glottal Closure

Rubén Fraile<sup>1,\*</sup>, Malte Kob<sup>2</sup>, Juana M. Gutiérrez-Arriola<sup>1</sup>, Nicolás Sáenz-Lechón<sup>1</sup>, J. Ignacio Godino-Llorente<sup>1</sup>, and Víctor Osma-Ruiz<sup>1</sup>

<sup>1</sup> Department of Circuits & Systems Engineering, Universidad Politécnica de Madrid  
Carretera de Valencia Km 7, 28031 Madrid, Spain

<sup>2</sup> Erich Thienhaus Institute, Hochschule für Musik Detmold  
Neustadt 22, D32756 Detmold, Germany  
{rfraile, jmga, igodino, vosma}@ics.upm.es,  
nicolas.saenz@upm.es, kob@hfm-detmold.de

**Abstract.** Pathological voices have features that make them distinct from normal phonetic voices. In fact, the instability of phonation associated to some voice disorders has a big impact on the spectral envelope of the speech signal and also on the feasibility of reliable pitch detection. These two issues (characteristics of the spectral envelope and pitch detection) and corresponding assumptions play a key role in many current inverse filtering algorithms. Thus, the inverse filtering of disordered or special voices is not a solved problem yet. Nevertheless, the assessment of glottal function is expected to be useful in voice function evaluation. This paper approaches the problem of inverse filtering by homomorphic prediction. While not favoured much by researchers in recent literature, such an approach offers two potential advantages: it does not require previous pitch detection and it does not rely on any assumptions about the spectral envelope of the glottal signal. Its performance is herein assessed and compared to that of an adaptive inverse filtering method making use of synthetic voices produced with a biomechanical voice production model. Results indicate that the performance of the inverse filtering based on homomorphic prediction is within the range of that of adaptive inverse filtering and, at the same time, it has a better behaviour when the spectral envelope of the glottal signal does not suit an all-pole model of predefined order.

## 1 Introduction

In a discrete signal processing framework, the most basic, yet widely used, voice production models map the human phonatory system to a set of linear systems connected in cascade. In their simplest form, such models consist of three blocks [1]: (i) the voice

---

\* This research was carried out within a project funded by the Ministry of Science and Technology of Spain (TEC2006-12887-C02) and within the “José Castillejo” program for the mobility of researchers. The work has also been done in the context of European COST action 2103.

source (glottal source)  $G(z)$ , which may be either a quasi-harmonic or a noise-like signal, depending on the type of sound being modelled (voiced or unvoiced), (ii) a vocal-tract filter  $V(z)$  that adds resonances to the signal and (iii) a radiation filter  $R(z)$  that accounts for flow-to-pressure conversion in the lips.

Actually, using non-invasive measuring techniques only the output of the composite system  $S(z) = G(z) \cdot V(z) \cdot R(z)$  can be measured as an audio signal. In this context, glottal inverse filtering is a blind deconvolution problem consisting in estimating the glottal source  $G(z)$  from the recorded voice  $S(z)$  without any *a priori* knowledge on the precise form of  $V(z) \cdot R(z)$ . While the extraction of the glottal source from recorded voice is of interest for many applications and several algorithms for such inverse filtering have been proposed, to present a standard automatic procedure is still lacking [2].

The highest difficulty in glottal inverse filtering (as in other blind deconvolution problems) is the discrimination among the effects of individual systems of the overall output. In the specific case of glottal inverse filtering, this problem may be approached in a number of different ways [2]:

- *Pitch-synchronous Approaches*. They require identification of glottal closure instants. The reason for this need may either be for analysing glottal closed phase [3], which is a time interval in which the vocal tract receives no input and, therefore, output voice is assumed to be independent of glottal source, or for isolating glottal source periods whose waveform may be estimated either adaptively [4] or jointly with the vocal tract [5].
- *Pitch-asynchronous Approaches*. These approaches do not necessarily require either detection of specific time instants or pitch-period calculation, though the performance of some of them may be increased including pitch synchronism. The most known scheme in this group is the iterative adaptive inverse filtering (IAIF) algorithm [6]. It assumes a two-pole model for the glottal source and uses such assumption to refine the all-pole vocal tract estimation. A similar procedure based on lattice filters has been reported in [7]. An alternative asynchronous approach consists in making use of the deconvolution capabilities of cepstrum for discriminating between glottal source and vocal tract plus radiation. Such approach was firstly proposed in [8] and was refined afterwards with the addition of pole-zero modelling [9]. Recently, its use for the estimation of vocal tract resonances has been reported [10].

The interest of the glottal source waveform for the assessment of laryngeal pathologies comes from the close expected relationship between laryngeal function and the glottal waveform itself. Some results on this application have been reported, for instance, in [11], [12] and [7]. However, the application of glottal inverse filtering techniques to pathological voices has a number of difficulties that should not be disregarded. In the first place, pathological voices may not have a clear harmonic or quasi-harmonic structure (see type 3 voice segments in Chap. 4 of [13]) and some pathologies may prevent complete glottal closure [13] (Chap. 5). Therefore, the implementation of pitch-synchronous approaches may be problematic in such cases. In the second place, assumptions about the spectral envelope of the glottal waveform (e.g. a 12 dB/oct decay [2]) that are inherent to some approaches, for instance IAIF, may not be valid for

pathological voices. In addition, other not yet solved issues of inverse filtering, no matter its application, have to be considered too. One of the most remarkable of such issues is the evaluation of the inverse filtering algorithms themselves. Although a set of objective measures for this evaluation has been proposed [14], these rely on the expected characteristics of the glottal source waveform, not on the measured characteristics, since the glottal source is commonly unknown. One way to solve that problem is the usage of synthetic voices for the assessment of the algorithms [2], but the validity of this approach depends on the realism of the used voice synthesisers.


In the previously described context, the authors herein report on the evaluation of two inverse filtering approaches for pathological voice signal analysis. Due to the above-mentioned potential characteristics of pathological voices, pitch-asynchronous approaches have been preferred. Among these, the performance of IAIF [6] has been compared to that of a variant of the homomorphic prediction (HoP) proposed in [9]. The performance has been evaluated using synthetic voice signals produced with a physical voice model [15] [16]. The usage of synthetic voices has allowed an objective and quantitative performance evaluation that has been carried out both in temporal and spectral domains. The rest of the paper is organised as follows: in Sect. 2 a description of the voice simulator and the voices produced with it is provided, Sect. 3 contains a description of the analysed inverse filtering algorithms, Sect. 4 presents the results of applying these algorithms to the synthetic voices and, last, Sect. 5 is dedicated to the conclusions.

## 2 Simulated Voice Signals

### 2.1 Simulation Model

The materials used for the herein reported experiments have been synthetic voice signals generated with the VOX simulator [17]. An overview of the simulation model can be found in [15] and a more thorough description in [16]. As far as this paper is affected, the simulation model consists roughly of two blocks: glottals model and vocal tract model. The glottal model is formed by a set of vocalis-mucosa pairs connected among them and with the larynx walls by means of springs and dampers of tunable parameters. Within each pair, both the vocalis and the mucosa are represented by one mass each, the mucosa above the vocalis, and connected between them also by a spring and a damper. For the work reported in this paper, each vocal fold has been modelled by a series of 15 vocalis-mucosa pairs. Two types of glottis have been simulated: a normal glottis with the vocal folds having straight edges and uniform stiffness and mass distribution and a glottis with one pair of nodules simulated by a localised concentration of mass and irregular vocal-fold edges. The specific form of the vocal folds has been chosen so as to mimic the movement of one of the nodular glottis examples provided in [13]. As for the rest of parameters, default values suggested in [16] have been used. Figure 1 shows a schematic representation of the glottal model.

As for the vocal tract, among the possible modelling approaches offered by VOX, the “stairs” model has been chosen. This corresponds to the Kelly-Lochbaum model. It considers the vocal tract as a series of 45 concatenated tubes with different diameters and the pressure wave propagation along the inside of such structure is simulated.




**Fig. 1.** Schematic representation of the vocal fold model:

*Top left:* Transverse view of the glottal model; each vocal fold is split into 15 elements

*Top right:* Arrangement in the coronal plane of each pair out of the 15 of elements; each element consists of two masses modelling *vocalis* and *mucosa*


*Bottom left & right:* Details of the links among elements via springs and dampers that simulate tissue behaviour



**Fig. 2.** Schematic representation of the vocal tract model

The simulation approach is similar to that of [18]. Figure 2 depicts one simulated configuration of the vocal tract. Specifically, three different vocal tract shapes have been simulated which intend to mimic the articulation of the vowels /a/ (as in “hat”), /o/ (as in “toe”) and /u/ (as in “foot”). These shapes have been obtained from [19]. Therefore, on the whole the experiments have been realised with six signals which come from the combination of the two glottal models with the three vocal tract models. For this particular experiment, uncoupled simulation of vocal folds and vocal tract has been chosen. Indeed, this eases the task of inverse filtering algorithms, since they are fed with signals






**Fig. 3.** Simulated normalised glottal flow for normal (up) and nodular (down) voices

for which the simplified voice production model mentioned in Sect. 1 is fully valid, which is not the case of real voices. However, the choice is justified because such a model has proven to be valid for a wide variety of applications and the purpose of the herein reported work is to provide a comparative analysis of algorithms and not absolute performance measures.

## 2.2 Characteristics of Simulated Glottal Signals

Figure 3 shows part of the simulated glottal flow signals for normal and nodular voices. In both plots, the signals have been normalised to be zero-averaged and with mean square value equal to one ( $\mu_{x_g} = 0$ ,  $\sigma_{x_g} = 1$ ). The sampling frequency of both signals is  $f_s = 17640$  Hz. While both signals have a fundamental frequency of 206 Hz –see the modulus of the corresponding Discrete Fourier Transforms (DFT) plotted in Fig. 4–, the normal voice presents a smooth glottal cycle, while the nodular voice has an irregular glottal cycle with oscillations shorter than one cycle. These correspond to a disjoint oscillation of the anterior and posterior parts of the glottis, with the nodules being the limit between those parts. In spectral domain, this is reflected by a non-harmonic structure (peaks are not repeated at regular frequency intervals) with spectral peaks higher than that corresponding to the fundamental frequency (Fig. 4, down). In Fig. 4, the spectrum of a linear prediction (LP) model fitted to the glottal signals has also been depicted. The spectrum of the glottal flow is typically assumed to have an envelope with a 12 dB/dec decay, which can be modelled with a 2-pole LP model [2]. In fact, a 2-to-4 pole LP model for the glottal flow is proposed in [6] and a similar number of poles (1 to 3) is suggested in [7]. The discrete all-pole (DAP) modelling algorithm [20] has been used for the fitting of the LP model. For the normal voice in Fig. 4 the LP model approximately corresponds to the simulated signal as for the placement of the highest spectral peak and the smooth decay of the spectral envelope above 1 kHz. In contrast the slope of the decay does not coincide. However, in the case of the nodular signal, the decay slope of both the model and the signal is roughly similar, but the highest peak of the spectral envelope is clearly different and the behaviour of the signal for high frequencies (above 3 kHz) does not match the model. These divergencies between the signals



**Fig. 4.** DFT of the simulated normalised glottal flow for normal (up) and nodular (down) voices. In both cases, the spectrum of an fitted 2-pole LP model has been plotted.

and the all-pole LP models pose the interest of using inverse filtering algorithms that make as less previous assumptions as possible regarding the spectral envelope of the vocal tract.

### 3 Inverse Filtering Algorithms

#### 3.1 Iterative Adaptive Inverse Filtering (IAIF)

As mentioned in Sect. 1, the IAIF algorithm [6] is a pitch-asynchronous scheme that has become prototypical for adaptive approaches to inverse filtering. The flow diagram of the IAIF algorithm is represented in Fig. 5. Within this algorithm, the voice signal is processed by three blocks:

1. A high-pass filter with cut-off frequency near 30 Hz that removes slow variations of the glottal flow.
2. A FIR filter that removes the effect of the vocal tract out of the voice signal. This block is called inverse filtering because its response  $H_{IF}(z)$  is sought to be the inverse of the vocal tract response:

$$H_{IF}(z) = \frac{1}{V(z)}. \tag{1}$$

Since  $V(z)$  corresponds an all-pole LP model:

$$V(z) = \frac{1}{1 + \sum_{k=1}^t a_k \cdot z^{-k}} \Rightarrow H_{IF}(z) = 1 + \sum_{k=1}^t a_k \cdot z^{-k}. \tag{2}$$

3. An integrator whose response  $H_i(z)$  compensates for the radiation effect of the lips. Such radiation effect is, approximately, a differentiation, thus:

$$H_i(z) = \frac{1}{1 - \rho \cdot z^{-1}} \approx \frac{1}{1 - z^{-1}} = \frac{1}{R(z)}. \tag{3}$$

where  $\rho$  is a positive real number close to 1 but lower, so as to ensure stability.

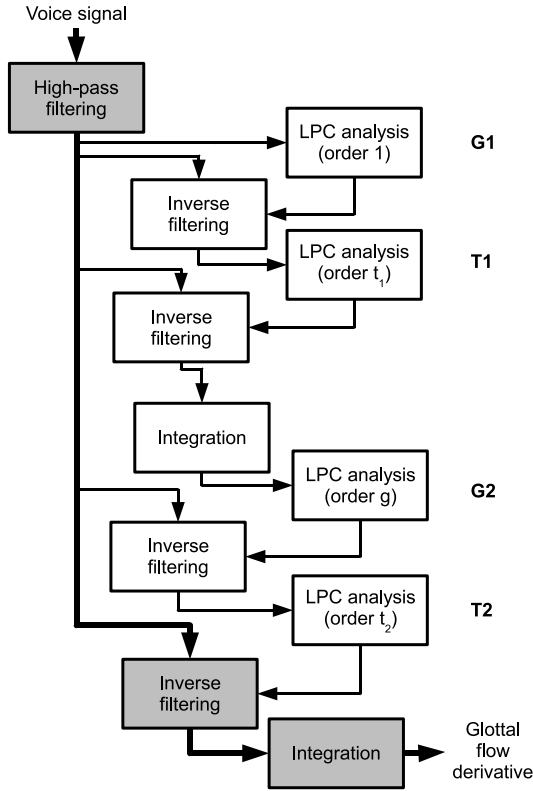


Fig. 5. Flow diagram of the IAIF algorithm

The purpose of the white blocks in figure 5 is to provide a good estimate for  $V(z)$ . For this purpose, a two-step procedure is repeated twice. In the first pass, a single-pole LP model is assumed for the glottal waveform  $G(z)$  ( $G1$ ) and, based on this assumption a  $t_1$ -pole LP model is fitted (LPC analysis) to the result of inverse-filtering the voice signal with that glottal model ( $T1$ ). Such  $t_1$ -pole LP model is a first estimate of  $V(z)$  which is used to produce a refined 2-to-4 pole LP model of  $G(z)$  ( $G2$ ). This allows, again, an improved estimation of  $V(z)$  using a LP model with  $t_2$  poles ( $T2$ ). The LPC analysis is done with the DAP algorithm [20], which has shown to provide better model estimates than classical autocorrelation methods. A free implementation of the IAIF algorithm can be found in [21].

### 3.2 Homomorphic Prediction (HoP)

An alternative, pitch-asynchronous approach for inverse filtering is that of homomorphic (or cepstral) analysis. This was first proposed in [8] and it is based on two ideas:

- The introduction of a logarithm operation allows to transform the convolution operation in a sum. That is, in the speech production model could be expressed in

“log- $z$ ” domain as:

$$\log S(z) = \log G(z) + \log V(z) + \log R(z) \quad (4)$$

- The vocal tract affects the spectral envelope of voice, while the glottal source is responsible for the fine quasi-harmonic structure; therefore, the inverse transform of (4) should result in a signal in which  $G(z)$  and  $V(z)$  are highly separable.

A combination of these two ideas with the LP modelling approach was introduced in [9] under the name of homomorphic prediction (HoP). This combination consisted in calculating the real cepstrum of the voice signal:

$$\tilde{S}[q] = \mathcal{F}^{-1} \{ \log |S(e^{j\omega})| \} = \tilde{G}[q] + \tilde{V}[q] + \tilde{R}[q] , \quad (5)$$

eliminating the part of the cepstrum corresponding to the harmonic structure of the voice signal by short-pass liftering:

$$\tilde{S}^{\text{SP}}[q] = \begin{cases} \tilde{S}[q] & \text{if } q < q_{\text{th}} \\ 0 & \text{if } q \geq q_{\text{th}} \end{cases} \quad (6)$$

and fitting an all-pole LP model to the remaining minimum-phase signal that can be recovered as:

$$V(z)|_{z=e^{j\omega}} \approx e^{\mathcal{F}\{\tilde{S}^{\text{SP}}[q]\}} . \quad (7)$$

The Fourier transform  $\mathcal{F}$  and its inverse  $\mathcal{F}^{-1}$  can be implemented with the DFT if a sufficient number of zeroes is appended to the original voice signal [22]. A very similar approach has more recently been used in [10] to estimate the resonances of the vocal tract for high-pitched voices.

While such an approach for separating the effects of glottis and vocal tract has been applied to the identification of the vocal tract response both in [9] and [10], it has not been tested yet for glottal-source recovery. In fact, the liftering operation in (6) separates between spectral envelope and harmonic structure, provided that  $q_{\text{th}}$  is chosen to be less but close to the fundamental period of the signal. However, it does not differentiate between the component of the spectral envelope of the signal due to the vocal tract and that due to the glottal waveform itself. The glottal waveform is expected to provide a smooth decay of the spectral envelope (as illustrated in figure 4), while the vocal tract is expected to produce peaks associated to its resonant frequencies. Therefore, in the cepstrum, the lowest values of  $q$  in (5) should be associated to the glottal waveform and the highest values below the fundamental period should be associated to the effect of the vocal tract.

Thus, the following algorithm based on HoP is proposed for inverse filtering:

1. Apply an integrating filter (3) to the voice signal so as to compensate for the lip radiation  $R(z)$ . This produces a radiation-compensated signal  $S_R(z)$ .
2. Calculate the cepstrum of  $S_R(z)$  as in (5):

$$\tilde{S}_R[q] = \mathcal{F}^{-1} \{ \log |S_R(e^{j\omega})| \} \approx \tilde{G}[q] + \tilde{V}[q] . \quad (8)$$

3. Choose and appropriate value of  $q_{\text{th}}$  to perform long-pass liftering:

$$\tilde{S}_R^{\text{LP}} [q] = \begin{cases} 0 & \text{if } q \leq q_{\text{th}} \\ \tilde{S}_R [q] & \text{if } q > q_{\text{th}} \end{cases} . \quad (9)$$

4. Recover the minimum-phase signal:

$$S_R^{\text{LP}} (z) \Big|_{z=e^{j\omega}} = e^{\mathcal{F}\{\tilde{S}_R^{\text{LP}} [q]\}} . \quad (10)$$

5. Fit a discrete all-pole LP model to  $S_R^{\text{LP}} (z)$ . Such a model should be an estimate of  $V (z)$ .
6. Use the estimate of  $V (z)$  to perform inverse filtering on  $S_R (z)$ , hence obtaining an estimate of the glottal flow derivative  $G (z)$ .

## 4 Results

The two inverse filtering algorithms described in Sect. 3, namely IAIF and HoP, have been applied to the six synthetic voice signals mentioned in Sect. 2. In all six cases, signal interval with lengths equal to 717 ms have been chosen, skipping the initial transient phase of the simulation. The outputs of both IAIF and HoP have undergone an additional integration to obtain the normalised glottal flow from its derivative. Such integration has been performed with filter  $H_i (z)$  in (3). The value of  $\rho$  has been tuned manually to be 0.92. The values of the vocal-tract model order  $t$  and the glottal model order  $g$  are as recommended in [6] and [7]:  $t = \lceil f_s (\text{kHz}) \rceil = 18$  and  $g = 2$ . The value of  $q_{\text{th}}$  has been chosen considering the widest vocal tract formant bandwidths (around 300 Hz) reported in [4]:  $q_{\text{th}} = \lceil f_s / 300 \rceil = 59$ . As for the input signals, all of them have been normalized to have zero mean and unit standard deviation.

Figure 6 show the amplitude of the frequency response of the simulated vocal tract corresponding to the vowel /a/ and its estimates provided by the IAIF and HoP algorithms for both normal and nodular voices. Qualitatively, it can be noticed that while the IAIF algorithm allows a better identification of the resonant frequencies of the vocal tract, the HoP provides a better fit to its locally averaged shape. This is specially true for frequencies below 2000 Hz.


The simulated and estimated glottal flows for 30 ms segments of the signals corresponding to the vocal tract /o/ are plotted in Fig. 7. Qualitatively, both for normal and nodular voices, HoP tends to provide a better fit to the original signal than IAIF. In order to provide a quantitative comparison of the performance of both approaches, the available signals have been randomly split in segments with durations uniformly distributed between 70ms and 210ms. For each segment, both the IAIF and HoP algorithms have been applied to produce glottal flow estimates  $\hat{g}_k [n]$ , where  $k = 1 \dots K$  stands for the processed voice segment. Afterwards, such estimates have also been normalized and time aligned with the corresponding simulated glottal flows  $g_k [n]$ . The experiment has been repeated 60 times with each signal with different starting points for the voice segments and an average signal-to-error measure has then been computed as follows:

$$SER = -20 \cdot \log \left[ \frac{1}{K} \sum_{k=1}^K \sqrt{\frac{E \{ (\hat{g}_k [n] - g_k [n])^2 \}}{E \{ g_k^2 [n] \}}} \right], \quad (11)$$


where  $E\{\cdot\}$  means time averaging and  $K$  is the total number of segments along the 60 iterations. The obtained values of SER for the six test signals are compiled in Tab. I. The performance of HoP for this particular set-up and the used test signals is significantly better than that of IAIF for the normal voices and it is roughly similar for the pathological voices.

**Table 1.** Values in dB of the SER, as defined in (11)

Vocal tract model	/a/	/o/	/u/
<i>Normal Voice</i>			
HoP	13.09	11.97	12.29
IAIF	10.69	7.40	9.29
<i>Nodular Voice</i>			
HoP	1.25	5.44	5.44
IAIF	2.37	4.37	3.72



**Fig. 6.** Simulated and estimated vocal tract frequency responses for normal (up) and nodular (down) voices




**Fig. 7.** Simulated and estimated glottal flows for normal (up) and nodular (down) voices

An insight into the reasons for the reported results can be found by analysing the spectral behaviour of the glottal waveform estimation error. Specifically, figure 8 shows the plot of:

$$E(\omega) = 20 \cdot \log \left[ \frac{1}{K} \sum_{k=1}^K \left| \hat{G}_k(e^{j\omega}) - G_k(e^{j\omega}) \right| \right] \quad (12)$$

for both normal and nodular voices and for the vocal tract /a/, for which HoP has the worst comparative performance. For normal voice, the HoP estimation error is lower than that of IAIF for all frequencies except for an interval around the first harmonic, which is very close to the peak provided by the 2-pole LP model of the glottal signal (recall Fig. 3). It should be remembered that the 2-pole LP model in this case was not able to match the decay of the spectral envelope of the signal, hence the wide interval over which HoP provides a better estimate of the glottal flow. In contrast, the fit between the decay of the 2-pole LP model and the signal was much better for the simulated pathological signal (recall Fig. 3) again). This results in a better performance of IAIF except around the position of the spectral peak of the 2-pole model (500 to 1100 Hz) and for very low frequencies (less than 200 Hz). Therefore, the comparative performance of HoP versus IAIF seems to be directly affected by the goodness of the fit between

the spectral envelope of the glottal flow and the all-pole LP model used to model that flow. In this case, a good fit, as is the case for the nodular voice, results in a similar performance of IAIF and HoP, while a divergence results in HoP outperforming IAIF.



**Fig. 8.** Estimation error in spectral domain for normal (up) and nodular (down) voices

## 5 Conclusions

Within this paper, the applicability of homomorphic prediction to inverse filtering has been analysed. This approach, while it has not been favoured much by researchers during the last years was firstly proposed in [9] for the estimation of the resonance structure of the vocal tract. Herein, it has been shown that a modification of such an algorithm by using long-pass liftering in cepstral domain instead of short-pass liftering provides an alternative method for the identification of the glottal source in voice signals. This method shares with iterative adaptive inverse filtering its capability for working pitch-asynchronously, which is a relevant issue for the processing of pathological voices. Moreover, though not directly studied in this paper, previously reported analyses indicate that HoP may outperform IAIF for high-pitched voices [10] [2] in pitch-asynchronous schemes.



The performance of both HoP and IAIF has been analysed quantitatively by using synthetic voice signals provided by a biomechanical simulator of voice production. The reported results show that both algorithms provide a similar performance when the glottal signal suits well the all-pole LP model assumed in IAIF, while HoP gives better figures when this does not happen. Therefore, it can be concluded that HoP provides a scheme for inverse filtering that is adequate for pathological voices since it pitch-asynchronously and that is more robust against variability in the spectral envelope of the glottal signal, since it does not impose assumptions related to its spectral decay.

## References

1. Rabiner, L.R., Schafer, R.W.: *Digital processing of speech signals*. Prentice-Hall, Englewood Cliffs (1978)
2. Walker, J., Murphy, P.: A review of glottal waveform analysis. In: Stylianou, Y., Faundez-Zanuy, M., Esposito, A. (eds.) *COST 277. LNCS*, vol. 4391, pp. 1–21. Springer, Heidelberg (2007)
3. Wong, D., Markel, J.D., Gray Jr., A.H.: Least squares glottal inverse filtering from the acoustic speech waveform. *IEEE Transactions Acoustics, Speech and Signal Processing* 27, 350–355 (1979)
4. Akande, O.O., Murphy, P.J.: Estimation of the vocal tract transfer function with application to glottal wave analysis. *Speech Communication* 46, 15–36 (2005)
5. Fu, Q., Murphy, P.: Robust glottal source estimation based on joint source-filter model optimization. *IEEE Transactions on Audio, Speech and Language Processing* 14, 492–501 (2006)
6. Alku, P.: An automatic method to estimate the time-based parameters of the glottal pulse-form. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 2, pp. 29–32 (1992)
7. Gómez-Vilda, P., Fernández-Baillo, R., Rodellar-Biarge, V., Nieto-Lluis, V., Álvarez-Marquina, A., Mazaira-Fernández, L.M., Martínez-Olalla, R., Godino-Llorente, J.I.: Glottal source biometrical signature for voice pathology detection. *Speech Communication* 51, 759–781 (2009)
8. Oppenheim, A., Schafer, R.W.: Homomorphic analysis of speech. *IEEE Transactions on Audio and Electroacoustics* 16, 221–226 (1968)
9. Kopec, G., Oppenheim, A., Tribolet, J.: Speech analysis homomorphic prediction. *IEEE Transactions on Acoustics, Speech and Signal Processing* 25, 40–49 (1977)
10. Rahman, M.S., Shimamura, T.: Formant frequency estimation of high-pitched speech by homomorphic prediction. *Acoustical Science and Technology* 26, 502–510 (2005)
11. de Oliveira-Rosa, M., Pereira, J., Grellet, M.: Adaptive estimation of residue signal for voice pathology diagnosis. *IEEE Transactions on Biomedical Engineering* 47, 96–104 (2000)
12. Gómez-Vilda, P., Fernández-Baillo, R., Nieto, A., Díaz, F., Fernández-Camacho, F.J., Rodellar, V., Álvarez, A., Martínez, R.: Evaluation of voice pathology based on the estimation of vocal fold biomechanical parameters. *Journal of Voice* 21, 450–476 (2007)
13. Sapienza, C., Hoffman-Ruddy, B.: *Voice Disorders*. Plural Publishing (2009)
14. Moore, E., Torres, J.: A performance assessment of objective measures for evaluating the quality of glottal waveform estimates. *Speech Communication* 50, 56–66 (2008)
15. Kob, M., Alhuser, N., Reiter, U.: Time-domain model of the singing voice. In: *Proceedings of the 2<sup>nd</sup> COST G-6 Workshop on Digital Audio Effects, Trodheim, Norway* (1999)
16. Kob, M.: *Physical Modeling of the Singing Voice*. PhD thesis, Fakultät für Elektrotechnik und Informationstechnik - RWTH Aachen, Logos-Verlag (2002)

17. Kob, M.: Vox - a time-domain model for the singing voice (2002), <http://www.akustik.rwth-aachen.de/~malte/vox/index.html.en> (visited May 2009) Computer software
18. Mathur, S., Story, B.H., Rodriguez, J.J.: Vocal-tract modeling: fractional elongation of segment lengths in a waveguide model with half-sample delays. *IEEE Transactions on Audio Speech and Language Processing* 14, 1754–1762 (2006)
19. Story, B.H., Titze, I.R.: Parameterization of vocal tract area functions by empirical orthogonal modes. *Journal of Phonetics* 26, 223–260 (1998)
20. El-Jaroudi, A., Makhoul, J.: Discrete all-pole modeling. *IEEE Transactions on Signal Processing* 39, 411–423 (1991)
21. Arias, M., Bäckström, T.: TTK aparat (2008), <http://aparat.sourceforge.net> (visited May 2009)
22. Childers, D., Skinner, D., Kemerait, R.: The cepstrum: A guide to processing. *Proceedings of the IEEE* 65, 1428–1443 (1977)

# Texture Discrimination with Artificial Whiskers in the Robot-Rat Psikharpax

Steve N'Guyen<sup>1,2</sup>, Patrick Pirim<sup>2</sup>, and Jean-Arcady Meyer<sup>1</sup>

<sup>1</sup> Institut des Systèmes Intelligents et de Robotique  
Université Pierre et Marie Curie-Paris 6, CNRS UMR 7222

4, place Jussieu, 75252 Paris Cedex 05, France

<sup>2</sup> Brain Vision Systems, 75013 Paris, France  
steve.nguyen@isir.upmc.fr

**Abstract.** We describe a novel algorithm for texture discrimination which we tested on a robot using an artificial whisker system. Experiments on both fixed head and mobile platform have shown that this system is efficient and robust, with greater behavioral capacities than previous similar approaches, thus, demonstrating capabilities to complement or supply vision. Moreover, results tends to show that the length and number of whiskers may be an important parameter for texture discrimination. From a more fundamental point of view these results suggest that two currently opposing hypotheses to explain texture recognition in rats, namely the “kinetic signature hypothesis” and the “resonance hypothesis”, may be, in fact, complementary.

## 1 Introduction

Touch is a very important sensory modality for many species of insects and mammals. For example, the whiskers of a rat are often compared to human fingertips in terms of their tactile - or haptic - ability. In particular, they make it possible to finely discriminate textures [1,2] or objects [3] and even to precisely determine an aperture width [4]. Biologists have studied rat whiskers for decades and know quite precisely the pathway from an individual vibrissa to the somatosensory cortex. One remarkable property of this haptic system is that whiskers project somatotopically to this part of the cortex, into a structure named “barrel cortex”. A “barrel” is a discrete neural structure that receives an input principally from a given whisker, with little influence from neighboring whiskers [5]. This relatively simple system, as compared to vision for example, facilitates the study of the neural coding scheme, as well as its use for perception and higher-level cognition.

Being simple, efficient and robust, whiskers should become popular in robotics [6] although few robots have been equipped with such devices in the past. The corresponding implementations were calling on various sensors ranging from the simplest binary switch to a very accurate bi-dimensional torque sensor. Brooks (1989), for example, used a simple sensor made of a metal shaft fixed on a push button, providing a very robust security sensor for a walking robot. Another implementation called upon probe whiskers made of a stem glued to a potentiometer with return springs and was used to evaluate the contour of an object [8]. Even wind sensitive sensors have been designed

[9] allowing a robot to navigate through a labyrinth. Basically, this sensor was made of small springs surrounded by electric contacts and was able to detect the direction of the wind.

Recently, several artificial whisker systems have been used in robotics to discriminate textures. Whisker hairs of real rats, glued to capacitive sensors (electret microphone), served [10,11] to produce very precise haptic sensors, with an uni-dimensional measurement of dynamic signals. Using an active whisker array of such sensors mounted on a mobile robot, Fend et al. (2003) successfully discriminated a set of 11 textures. Kim and Möller (2004) tried both piezo and hall-effect sensors which, mounted in orthogonal pairs, provided a bi-dimensional measure of vibrissa deflection. Like capacitive sensors, piezo sensors cannot deliver static signals, but this can be achieved using an extra integrator circuit. With a data processing based on spectrum density, these authors were able to discriminate a set of 7 sandpapers. Likewise, Seth et al. (2004) performed texture discrimination using arrays of Flex sensors, which provided an unidimensional measure of curvature. Here, temporal differences between pairs of vibrissæ were fed into a barreloid system with spiking neurons. Finally, Fox et al. (2009) used two active whiskers with strain gage-based sensors mounted on a mobile robot. They explored different bioinspired methods of feature extraction and the implication of unconstrained whisker-texture contact on classification performance.

The work described herein contributes to the Psikharpax project [15] which aims at designing a biomimetic artificial rat. Besides visual, auditory and vestibular sensors, the corresponding robotic platform is equipped with an original whisker system described elsewhere [16]. This system is intended to be used for texture discrimination and object recognition and, more generally, as an auxiliary or a substitute to vision. Its performance in texture discrimination is the subject of the present article.

## 2 System Description

Insofar as the impact of the specific implementation of a rat's whisker system on its functionalities is currently unknown, we tried to design an artificial whisker system mimicking as much as possible the natural organization. Accordingly, our system [16] is based on a simple, elastomer-based, artificial skin with two arrays of 33 vibrissæ and an arc/row organization (cf. Fig 1) and a whisker-length gradient (cf. Table 1) quite similar to those encountered in the rat.

The deflection of each vibrissa is sampled in both its anteroposterior and dorsoventral axes, providing two 8-bit measurements at 1157Hz. However, orientation information being not necessary for texture discrimination, the two measures are normed ( $\sqrt{x^2 + y^2}$ ).


## 3 Texture Discrimination

### 3.1 Feature Extraction

Neither the details of how a rat's brain actually encodes texture features, nor the exact nature of these features, are yet known. Arabzadeh et al. (2004) experienced different

feature codings on both artificial and natural (in vivo) whiskers. Starting from the principle that a pure sinusoidal signal can be fully described by its amplitude  $A$  and its frequency  $f$ , they stimulated a rat's whiskers with various signals varying in amplitude and frequency. Then, recording the induced neural activity in the barrel cortex, they deduced that the neural activity most probably encodes a quantity homogeneous to the product  $Af$ . The generalized expression of this quantity to any natural signal is known as the "equivalent noise level (ENL)" (for more details see: [18]), which is usually used to measure sound power. This quantity can also be related to the more common "spectral centroid" [14].


To compute the latter, instead of using a Fast Fourier Transform algorithm - of which no natural equivalent is known - we simply called upon a time domain "on-the-fly" estimate of the quantity  $X\omega$ . The corresponding algorithm (cf. Fig 2) can be compared to those used in auditory feature extraction, like ZCPA that is used for speech recognition [19,20,21]. It provides a quantity homogeneous to the ENL which we call the "Instantaneous Mean Power" or IMP feature.



**Fig. 1.** (Top) Psikharpax robot. (Bottom) Comparison of whisker pads in a real rat and in our robot.

**Table 1.** Vibrissæ arcs, with mean measured lengths in mm (in one adult *rattus norvegicus* specimen). Compared to those of a real rat, the robot’s whiskers are approximately 6 times longer.

Arc	vibrissæ	rat	robot
1	$\alpha, \beta, \gamma, \delta, E1$	41.8	250
2	A1, B1, C1, D1, E2	37.2	200
3	A2, B2, C2, D2, E3	27.6	150
4	A3, B3, C3, D3, E4	20.6	120
5	A4, B4, C4, D4, E5	12.6	100
6	C5, D5, E6	8.33	90
7	C6, D6, E7		70
8	D7, E8		55



**Fig. 2.** Feature extraction algorithm. "Peaks" are detected through the monitoring of the signal’s derivative and frequencies are estimated through the inverses of the time intervals between successive peaks. Then, the peak amplitude is multiplied by the peak frequency, averaged within a time window.

This approach relies on the strong hypothesis that the peaks thus characterized provide enough information to describe a texture. Such hypothesis is reinforced by the fact that, when Licklider and Pollack (1948) assessed the effects of various signal distortions in human speech recognition, they found that “infinite clipping” - a treatment that only kept a signal’s periodicity - did not prevent speech recognition in humans.

Be that as it may, the corresponding algorithm is very simple and computationally very cheap as it necessitates only one division per peak detection ( $\frac{\text{Peak amplitude}}{\text{Peak period}}$ ) plus one addition (to compute the peak’s period) at each time step. As for peak detection proper, it only entails one subtraction ( $x_t - x_{t-1}$ ) and a comparison.

There are however possible limitations to the proposed algorithm. In particular, input data are drastically reduced by this procedure according to which a pure sinus input of frequency  $F$  and a triangle input of fundamental frequency  $F$  will lead to the same feature value although they obviously don’t have the same spectrum. Likewise, turns out that amplitude modulations cannot be detected by a single vibrissa. Our hypothesis is that such limitations are alleviated by the fact that the natural filtering of vibrissæ, due to their intrinsic mechanical characteristics, decomposes complex signals along the pad in a manner similar to how the cochlea decomposes auditory signals.

### 3.2 Fixed Head Experiment

**Experimental Apparatus.** At first, we tested this haptic system according to a relatively constrained fixed head experiment that consisted in sweeping a whisker pad over a set of eight sandpapers whose grits varied from P180 to P50 (cf. Fig 3). Sandpaper provide a complex random texture appropriate for this task and has been used on



Fig. 3. The texture set used for discrimination




Fig. 4. Schematic of the experimental protocol. A: start position, B: mid position, C: end position

various experiment with real rats. Using this material, we performed qualitative experiments with humans that clearly indicated that the task of discriminating such textures by tactile contact only is a very difficult one, an observation also made by Hipp and coll. (2005).


A vibrissa pad was fixed on the robot's head which could move in pan-tilt directions. The pan axis was at a fixed distance from the texture sample (cf. Fig 4) that was presented with a small amount of variability in position ( $\sim \pm 1cm$ ) between each trial, with an appropriate angle to provide contact with a maximum number of whiskers.

For each texture, 400 experiments were made, 300 for learning and 100 for testing. The raw data ( $x$  and  $y$  deflections, 8 bits resolution sampled at 1157Hz) were normed ( $\sqrt{x^2 + y^2}$ ). For each vibrissa, this measure was fed into the feature extraction algorithm that output the IMP as one float value. Finally we summed these IMP values for each vibrissa during the sweep. Having thus obtained an input vector of 33 floats for each trial, we fed it into a simple multi layer perceptron (MLP) with 33 input neurons, 33 hidden layer neurons and 8 output neurons, to perform supervised learning. We used the FANN library [24] with the iRPROP training algorithm [25]. The final classification was done by a winner-take-all (WTA) on the 8 output neurons.

**Results.** Table 2 gives the confusion matrices obtained on 100 test data. The mean performance is clearly above 90% (here the chance level is  $\frac{1}{8} = 12.5\%$ ), which greatly improves the human aptitude at solving the same task.

**Table 2.** Confusion matrix obtained for the 8 textures using IMP

IMP	1	2	3	4	5	6	7	8
1	100	0	0	0	0	0	0	0
2	0	99	1	0	0	0	0	0
3	0	2	95	0	3	0	0	0
4	0	0	0	96	4	0	0	0
5	0	0	1	0	99	0	0	0
6	0	0	0	0	1	93	6	0
7	2	0	6	0	3	9	80	0
8	2	0	0	0	0	1	0	97


**Fig. 5.** Top view of the experimental protocol

Using the data thus acquired, we tested the influence of the number of vibrissæ on the classification performance. Starting with data obtained with one arc (Arc 1, 5 vibrissæ cf. Table 1), then with two arcs (Arc 1 + Arc 2, 10 vibrissæ) etc, up to the whole whisker pad, we assessed at each stage the quality of the discrimination. Results are summarized on Figure 6. The percentage of successful discriminations is quickly rising with the number of vibrissæ involved and reaches values comprised between 90 and 95% when three or more arcs (15 vibrissæ) are concerned. This result confirms previously obtained ones in [10,23].


When analyzing the performances obtained with a single arc (cf. Fig 7), one observes a great variability. Indeed, it seems that each arc separately performs better on a subset of the textures. For example, arc 2 is very bad at recognizing texture 5, but quite good with texture 6. This suggests that iso-length arcs complement each other and probably explains the increase in performance with the number of arcs involved.

The relative quality of these results demonstrates that the IMP is a suitable feature for texture recognition. However, as Fox et al. (2009) pointed out, the kind of fixed head task used so far is very different from that of a robot moving in an environment,





**Fig. 6.** Mean performance (% of successful discriminations) obtained with IMP, over the number of vibrissae involved




**Fig. 7.** Mean performance obtained with IMP for the 5 longer arcs across the 8 textures


where the distances and angles with which whiskers touch any texture are constantly varying. Therefore, to assess the robustness of the IMP, we also performed such a complementary experiment.

### 3.3 Mobile Robot Experimentation

**Experimental Apparatus.** In this experiment, a set of complex textures (cf. Fig 8) made of plexiglass were fixed on the sides of two small corridors (1m long). A different texture was assigned to each side of each corridor. The robot’s task was to follow the walls in its environment, to enter a corridor if encountered, to recognize the textures on its sides, and to learn to turn left or right at the end of the corridor, depending on the left/right combination of the textures thus recognized. A main difference with the previous experiment was the “touch strategy”. We previously swept



**Fig. 8.** The four textures used in the mobile robot experiment. These textures were made of relief decorated plexiglass.



**Fig. 9.** Robot environment showing the 2 corridors and direction convention used (Top). Robot inside a corridor, with dimensions and command vector (Bottom).

whiskers on a texture by rotating the head, trying to maximize the number of whiskers in contact with the texture. Here, the whole robot was moving, the head didn't rotate and only a subset ( $\sim 10$  vibrissae, the two longer arcs) of whiskers were actually touching the textures, from a varying distance.

To allow the robot to navigate in its environment using only its whiskers as sensory input, we developed a simple obstacle avoidance strategy. A distance information was first computed by taking into account the iso-length arcs. One minus the mean arc deflection was weighted by the mean within-arc vibrissa size. Thus, the more vibrissæ were bent, the smaller was the output distance. Repeating this computation for each arc, we obtained a value that decreased with the contact distance,

$$D = \frac{1}{N} \sum (1 - V_i) \times L_i \quad (1)$$

$V_i$  being the mean deflection of the  $i^{th}$  arc and  $L_i$  the mean length of the  $i^{th}$  arc. One can notice that the smaller whiskers - the most frontal ones - contribute less to this measure than the longer ones. This may seem counter intuitive as, when an object touches the small whiskers, it is probably closer than if it only touches the longer ones. But generally in the described task, when an object touches the smaller whiskers, it also touches the longer ones and the above weighting prevents an over-reaction. Moreover this method has shown a better stability in corridors, where a small variation of vibrissa deflection should provoke a small orientation reaction in order to make the corresponding trajectory as straight as possible. This centering strategy was an important component of the robot's behavior since the corridors were 35cm wide, while the robot's width was 25cm (including wheels) and the maximum whisker range was 50cm, leaving a small error margin for steady forward movement and whisker crack avoidance (cf. Fig 9). We controlled the robot through a speed vector  $V$  applied at the axis of the neck whose orientation component  $V_y$  was given by:

$$V_y = (D_{left} - D_{right}) \times G \quad (2)$$


with the gain  $G = 0.01$ .  $V_x$ , the translation speed, was fixed to 10cm/s. This simple control produced an obstacle avoidance behavior, with a tendency to wall following. Additionally, this control produced a relatively stable corridor centering behavior - which was its principal objective. Finally, using  $D_{left}$  and  $D_{right}$  values, we could roughly determine the corridor's aperture size and trigger the learning/recognition procedure only when the robot was inside a corridor as determined by a distance threshold.

We first ran a series of 10 experiments for each corridor and each direction. We simply positioned the robot approximately in front of the corridor and recorded the IMP feature output at each time step within the corridor. We then fed 7 data runs to a MLP ( $2 \times 33$  neurons in the input layer,  $2 \times 33$  neurons in the hidden layer and 4 neurons in the output layer), keeping the 3 others runs to test the learning result. A typical data run length was  $\sim 7000$ .

**Results.** Once the learning was completed, we ran four series of 20 additional experiments to evaluate the capacity of the robot to turn in the right direction at each end of each corridor. While the robot moved in a given corridor, in a given direction, we fed the smoothed (low pass filter)  $2 \times 33$  IMP output to the previously learned MLP and computed the WTA on the output layer. By accumulating this winner value through the whole corridor, we obtained a mean decision vector which served to take the final decision (once again by a WTA). The corresponding results are summarized in Table 3. As expected, the trajectory stability played a role in performances as dithering in

**Table 3.** Confusion matrix obtained for 20 runs for each corridor, in each direction

Corridor-direction	1-1	1-2	2-1	2-2
1-1	75	15	10	0
1-2	0	100	0	0
2-1	15	0	85	0
2-2	0	15	0	85

**Fig. 10.** Typical trajectory of the robot into the maze. Top: left oriented start. Bottom: right oriented start.

the corridor induced variations in the perceived vibrations. Most of the errors occurred when the robot's trajectories were unstable (lot of dithering).

We finally conducted qualitative experiments in the whole maze using the above described navigation rules. The maze was a round corner rectangle of 2.20m by 4.10m made of cardboard boxes with 2 corridors (cf. Fig 9). We added a simple hand cabled behavior consisting of turning left or right at the end of a corridor, depending on the

recognized textures. The robot was initially positioned near the wall on the top of the maze with left or right orientations. Any other starting position could have been used with the limitation of avoiding a direct wall facing, as no “reversal” behavior was implemented. In these conditions, the robot succeeded to autonomously circulate around the maze, following either direction indicated by the textures on the corridor’s sides. Several tours could be completed in a row thus demonstrating the efficiency of the robot’s haptic system.

Figure 10 shows an example of the kind of trajectories obtained.

## 4 Discussion

If some research efforts have been devoted to texture discrimination in “fixed head” tasks [10,11,12,14], very few robots have been able to navigate and recognize tactile cues in a less constrained environment using whiskers. One such work was done using curvature sensors with two different types of surfaces that one may consider more as a “shapes” than as a “textures”, as they seem to induce a mere deflection sequence rather than a complex vibration [13]. This robot could be conditioned to associate an aversive response with a given texture. A related work concerned a smooth versus rough discrimination task in an open arena and involved an active microphone-based whisker sensor with a natural rat’s hair [26]. Feature extraction called upon spectral analysis and lead to qualitatively good results. However, as the author concludes, such system could not be used to perform a more complex task without an improvement of its discriminatory capability and reliability. Finally, Fox et al. (2009) also obtained good results in a smooth/rough discrimination task on a mobile robot equipped with active whiskers using an “onset” feature. This “onset” feature is roughly the FFT magnitude within a short frequency band (2-3kHz) during the onset period of the whisker-texture contact (the first 12.8ms of the contact). Moreover, this feature is invariant to the relative position and orientation of whiskers and textures. Experimental conditions were slightly different from ours, as the texture position was chosen randomly and the robot didn’t move while touching a texture.

None of these related approaches seems suitable for performing a more complex task than simply discriminating smooth versus rough textures. In contrast, the haptic system that has been described herein proved to be able to use texture discrimination to afford minimal navigation capacities in a complex environment. Such capacities could be used to complement vision in daylight conditions or to replace it in the dark.

With this haptic system, texture recognition is possible in both fixed and mobile robot conditions. This tends to indicate that, despite the underlying simple algorithm and the various approximations on which it relies, the IMP feature is robust.

Conversely, we already mentioned that the whisker orientations in our system is not always well suited. Indeed, our whiskers are oriented toward the front (cf. Fig 5), which occasionally prevents all the whiskers from touching a texture. Within a corridor, for instance, about 10 whiskers only were touching the walls. Additionally, our implementation sometimes entails brusque return jumps of some whiskers when they are stuck on a given surface, rather than a gentle sweeping, which makes their signals totally unreliable. Fortunately, this problem only occurs in corridors and with a minority of whiskers

(usually the more dorsal and ventral ones) and thus the classifier can see it as mere noise. Obviously, a system in which the whisker orientation could be dynamically controlled - such as the one used in [14] - would be more adapted to alleviate this specific inconvenience and would be closer to the natural active whiskering system of rats.

Another remark concerns our feature extraction technique. We chose to design an algorithm that extracts an estimation of the amplitude-frequency product. This choice was based on a recent finding about how texture signals are encoded in a rat's brain [17]. Using such a feature, we were able to perform fine texture discrimination. This finding is an argument in favor of the so-called "kinetic signature" hypothesis which stands that each vibrissa encodes a specific signature of the touched surface in term of magnitude and temporal pattern.

Likewise, the fact that our results suggest that the texture discrimination capacities depend both on the length and number of the involved whiskers, seems to back up the "resonance hypothesis" [27][28] which stands that the self resonance property of a vibrissa plays a crucial role in vibration transduction and, in some way, helps to enhance texture perception. The exact manner in which this resonance property is exploited in rats is still unclear, but it seems quite reasonable to think that a kind of signal filtering is involved. Additional experiments with the current system might help clarify this issue.

Be that as it may, already acquired results strongly suggest that two hypotheses that are currently considered as mutually exclusive to explain texture recognition in rats - i.e., the "kinetic signature hypothesis" and the "resonance hypothesis" - may be, in fact, complementary.

## 5 Conclusions

We endowed a whiskered robot with a simple algorithm allowing to discriminate textures. Its efficiency has been demonstrated using both a fixed head and a mobile robot. Comparatively to previous similar approaches, this system affords greater behavioral capacities and may complement or supply vision in simple navigation tasks. Future work will be devoted to demonstrating its ability to perform shape recognition. On a fundamental level, it will also be used to investigate the influence of whiskers resonance properties on texture transduction.

## References

1. Carvell, G., Simons, D.: Biometric analyses of vibrissal tactile discrimination in the rat. *Journal of Neuroscience* 10, 2638–2648 (1990)
2. Guic-Robles, E., Valdivieso, C., Guarjardo, G.: Rats can learn a roughness discrimination using only their vibrissal system. *Behavioural Brain Research* 31, 285–289 (1989)
3. Brecht, M., Preilowski, B., Merzenich, M.: Functional architecture of the mystacial vibrissae. *Behavioural Brain Research* 84, 81–97 (1997)
4. Krupa, D.J., Matell, M.S., Brisben, A.J., Oliviera, L.M., Nicolelis, M.A.L.: Behavioural properties of the trigeminal somatosensory system in rats performing whisker-dependent tactile discrimination. *J. Neurosci.*, 5752–5763 (2001)
5. Petersen, R.S., Diamond, M.E.: Spatial-Temporal Distribution of Whisker-Evoked Activity in Rat Somatosensory Cortex and the Coding of Stimulus Location. *J. Neurosci.* 20, 6135–6143 (2000)

6. Hartmann, M.J.: Active sensing capabilities of the rat whisker system. *Autonomous Robots* 11, 249–254 (2001)
7. Brooks, R.A.: A robot that walks: Emergent behaviors from a carefully evolved network. Technical Report AI MEMO 1091, MIT (1989)
8. Russell, R.A.: Object recognition using articulated whisker probes. In: Proc. 15th Int. Symp. Industr. Robots, pp. 605–612 (1985)
9. Chapman, T., Hayes, A., Tilden, T.: Reactive maze solving with a biologically-inspired wind sensor. In: Meyer, J., Berthoz, A., Floreano, D., Roitblat, H., Wilson, S. (eds.) *From Animals to Animats 6*. Proc. of the 6th Int. Conf. on Simulation of Adaptive Behavior, pp. 81–87. MIT PRESS, A Bradford Book (2000)
10. Fend, M., Bovet, S., Yokoi, H., Pfeifer, R.: An active artificial whisker array for texture discrimination. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), vol. II, pp. 1044–1049 (2003)
11. Lungarella, M., Hafner, V., Pfeifer, R., Yokoi, H.: Artificial whisker sensors in robotics. In: IEEE/RSJ International Conference on Intelligent Robots and System, vol. 3, pp. 2931–2936 (2002)
12. Kim, D., Moller, R.: A biomimetic whisker for texture discrimination and distance estimation. *From Animals to Animats* 8, 140–149 (2004)
13. Seth, A.K., McKinsty, J.L., Edelman, G.M., Krichmar, J.L.: Spatiotemporal processing of whisker input supports texture discrimination by a brain-based device. In: Schall, S., Ijspeert, A., Billard, A., Vijayakumar, S., Hallam, J., Meyer, J. (eds.) *From Animals to Animats 8*. Proc. of the 8th Int. Conf. on Simulation of Adaptive Behavior. MIT Press, MA (2004)
14. Fox, C.W., Mitchinson, B., Pearson, M.J., Pipe, A.G., Prescott, T.J.: Contact type dependency of texture classification in a whiskered mobile robot. *Autonomous Robots* (2009) (in press)
15. Meyer, J.A., Guillot, A., Girard, B., Khamassi, M., Pirim, P., Berthoz, A.: The psikharpax project: Towards building an artificial rat. *Robotics and Autonomous Systems* 50, 211–223 (2005)
16. N'Guyen, S., Pirim, P., Meyer, J.A.: Elastomer-based tactile sensor array for the artificial rat psikharpax. In: ISEF 2009 - XIV International Symposium on Electromagnetic Fields in Mechatronics, Electrical and Electronic Engineering (2009) (in press)
17. Arabzadeh, E., Panzeri, S., Diamond, M.E.: Whisker Vibration Information Carried by Rat Barrel Cortex Neurons. *J. Neurosci.* 24, 6011–6020 (2004)
18. Arabzadeh, E., Zorzin, E., Diamond, M.E.: Neuronal encoding of texture in the whisker sensory pathway. *PLoS Biol.* 3, e17 (2005)
19. Ghitza, O.: Auditory models and human performance in tasks related to speech coding and speech recognition. *IEEE Transactions on Speech and Audio Processing* 2, 115–132 (1994)
20. Kim, D.S., Lee, S.Y., Kil, R.M.: Auditory processing of speech signals for robust speech recognition in real-world noisy environments. *IEEE Transactions on Speech and Audio Processing* 7, 55–69 (1999)
21. Sreenivas, T.V., Niederjohn, R.J.: Spectral analysis for formant frequency estimation in noise. *IEEE Transactions on Signal Processing* 40, 282–293 (1992)
22. Licklider, J.C.R., Pollack, I.: Effect of differentiation, integration, and infinite peak clipping upon the intelligibility of speech. *Journal of the Acoustical Society of America* 20, 42–52 (1948)
23. Hipp, J., Arabzadeh, E., Zorzin, E., Conradt, J., Kayser, C., Diamond, M.E., Konig, P.: Texture Signals in Whisker Vibrations. *J. Neurophysiol.* 95, 1792–1799 (2006)
24. Nissen, S.: Implementation of a Fast Artificial Neural Network Library (fann). Report, Department of Computer Science University of Copenhagen (DIKU) 31 (2003)
25. Igel, C., Hüsken, M.: Improving the rprop learning algorithm. In: Proceedings of the Second International Symposium on Neural Computation, NC 2000, pp. 115–121 (2000)

26. Fend, M.: Whisker-based texture discrimination on a mobile robot. In: Capcarrère, M.S., Freitas, A.A., Bentley, P.J., Johnson, C.G., Timmis, J. (eds.) ECAL 2005. LNCS (LNAI), vol. 3630, pp. 302–311. Springer, Heidelberg (2005)
27. Moore, C.I., Andermann, M.L.: 2. In: *The Vibrissa Resonance Hypothesis*, pp. 21–60. CRC Press, Boca Raton (2005)
28. Neimark, M.A., Andermann, M.L., Hopfield, J.J., Moore, C.I.: Vibrissa resonance as a transduction mechanism for tactile encoding. *The Journal of Neuroscience* (2003)



# Heart Rate Variability and the Acceleration Plethysmogram Signals Measured at Rest

Mohamed Elgendi<sup>1,2</sup>, Mirjam Jonkman<sup>1</sup>, and Friso DeBoer<sup>1</sup>

<sup>1</sup> School of Engineering and Information Technology, Charles Darwin University, Australia

<sup>2</sup> Nanyang Technological University

elgendi@ntu.edu.sg

{Mirjam.Jonkman, Friso.DeBoer}@cdu.edu.au

**Abstract.** It is well-known that the electrocardiogram (ECG) is a non-invasive method that can be used to measure heart rate variability (HRV). Photoplethysmogram (PPG) signals also reflect the cardiac rhythm since the mechanical activity of the heart is coupled to its electrical activity. Photoplethysmography is a non-invasive, safe, and easy-to-use technique that has been developed for experimental use in vascular disease. A useful algorithm for *a*-wave detection in the acceleration plethysmogram (APG, the second derivative of the PPG) is introduced to determine the interval between successive heartbeats and heart rate variability. In this study, finger-tip PPG signals were recorded for twenty seconds from 27 healthy subjects measured at rest. The use of the *aa* interval in APG signals showed very promising results in calculating the HRV statistical indices, SDNN and rMSSD.

**Keywords:** Heart rate, HRV, Photoplethysmogram, APG.

## 1 Introduction

The focus of this paper is the variation in the interval between consecutive heartbeats. Heart rate variability has become the conventionally accepted term to describe variations of both instantaneous heart rate and RR intervals in electrocardiogram (ECG) signals. A number of terms have been used in the literature to describe heart rate variability, for example, cycle length variability, heart period variability, RR variability, and RR interval tachogram. In this paper the acronym HRV will be used to describe the heart rate variability.

HRV measures the heart rate variations around the mean heart rate (HR). HRV provides information about the sympathetic-parasympathetic autonomic stability and consequently about the risk of unpredicted cardiac death. The traditional method to detect heart beats is detecting R peaks in ECG signals, a very promising tool to derive useful information about the hemodynamics as well as autonomic nerve system is the photoplethysmogram (PPG) signals. Accurate detection of inter-beat intervals from fingertip PPG signals is considered difficult. Ventricular pressure and other parameters of cardiac output however can influence the form and timing of the pulse waveform. In addition, peripheral effects, such as changes in vascular tone, may also influence distal pulse peak detection.


These possible weaknesses of using the fingertip PPG signals in measuring HRV are mentioned by Bernston *et al.* [1]. Hence, they recommend the usage of RR intervals from ECG signals to determine inter-beat intervals. They do believe, however, that with a sophisticated peak detection algorithm the use of intra-arterial pressure pulses may be acceptable. According to them indirect measures, such as photoplethysmographic signals need further validation.

Giardino *et al.*[2] demonstrated that under resting conditions the distal pulse pressure, as shown in Fig.1 (a), is sufficient for determining the heart rate. However, they recommended extra studies that include test–retest reliability evaluation of different data collection techniques.


These reasons may explain the lack in investigating the use of PPG signals instead of ECG signals to measure the heart rate (HR) and the heart rate variability (HRV).

The PPG contour itself can be used to detect the heart beat and consequently HRV can be measured [3], as shown in Fig.1 (a) the two circles represent two consecutive heart beats with the smallest positive values in the PPG signal. However, reliable detection of heart beats from the PPG contour is challenging due to PPG noise and its interference incorporated nature [4].

To overcome the PPG contour analysis, the second derivative of photoplethysmogram, also called the acceleration plethysmogram (APG), has been introduced, as shown in Fig.1 (b); the two circles represent two consecutive heart beats with the largest positive values in the APG signal. Because the peaks in the APG signal are more clearly defined than the peaks in the PPG contour the heart rate can be more accurately detected using the APG.



**Fig. 1.** Two successive beats in (a) fingertip photoplethysmogram (PPG) signal (b) second derivative wave of photoplethysmogram (APG) signal



**Fig. 2.** One beat measurement (a) PPG (b) APG

A typical APG tracing of the cardiac cycle (heartbeat) consists of four systolic waves and one diastolic wave [5], namely *a*-wave (early systolic positive wave), *b*-wave (early systolic negative wave), *c*-wave (late systolic reincreasing wave), *d*-wave (late systolic redecresing wave) and *e*-wave (early diastolic positive wave), as shown in Fig.2 (b). The height of each wave was measured from the baseline, with the values above the baseline being positive and those under it negative. The first systolic wave, the *a*-wave seems the most suitable for heart rate calculations because of its amplitude and steepness.

Taniguchi *et al* [6] used the *aa* interval in the APG signals to determine HR instead of using the ECG when assessing the stress experienced by surgeons.

In the present study, our goal was to determine if variations in the APG signal can be used instead of the ECG for measuring the HRV. In addition, we investigate the relationship between HR and HRV at rest. In order to measure HRV using the APG signals, the accurate detection of *a* wave is an essential step. To date, there is still a lack of studies focusing on the automatic detection of *a*-waves in APG signals. Therefore, this investigation also aimed to develop a fast and robust algorithm to detect *a*-waves in APG signals. The APG waveform was measured in a population-based sample of healthy males at rest.

## 2 Data

The PPG data were collected as a minor part of a joint project between Charles Darwin University, Defence Science and Technology Organisation (DSTO) and the Department of Defence. The background of the entire project can be found in [7].




Fig. 3. Flowchart for new  $a$  wave detection algorithm

Twenty seven healthy volunteers (males) with a mean $\pm$ SD age of 27 $\pm$ 6.9. PPG were measured by a photoplethysmography (Salus PPG), with the sensor located at the cuticle of the second digit of the left hand. Measurements were performed while the subject was at rest on a chair. PPG data were collected at a sampling rate of 200 Hz. The duration of each data segment is 20 seconds [7].

The test was conducted from 20<sup>th</sup> of April to 5<sup>th</sup> of May 2006 at Northern Territory Institution of Sport (NTIS).

All procedures were approved by the ethics committee of Charles Darwin University. Informed consent was obtained from all volunteers.

### 3 Methodology

Accurate HRV calculation depends on accurate *a*-waves detection in the APG signals. The calculation of HRV from APG signals consists of two steps: detection of *a*-waves and calculation of the heart rate variability.

#### 3.1 Detection of a-Waves

The proposed *a* wave detection algorithm consists of three main stages: pre-processing, feature extractions and classification. The flowchart of the algorithm is shown in Fig. 3.

1) *Pre-processing*: In this stage the measured PPG signal is filtered, differentiated twice (to provide clean APG signals) and then squared (to provide a clear position for each heart beat).


A) *Bandpass Filter*: the baseline wander and high frequencies, which do not contribute to *a*-wave detection, have been removed by a second order Butterworth filter with passband 0.5-10 Hz. typically a recommended bandpass filter is a bidirectional Butterworth implementation [8]. This bandpass filter performs zero-phase digital filtering by processing the input data in both the forward and reverse directions. Moreover, Butterworth filters offer good transition band chactersitics at low coefficient orders which make them efficient to implement [8].

A second order Butterworth filter with bandpass 0.5-10 Hz implemented by cascading a high pass and low-pass filter to remove the baseline wander and high frequencies which do not contribute to *a* and *b*-waves. Since one complete heart cycle takes approximately one second, the frequencies below 0.5 Hz can be considered as noise (baseline wander). The 10 Hz is chosen because most the energy of the PPG signal is below 10 Hz, as shown in Fig. 4(b).

The time domain representation of the filter is required because the filter is applied to PPG data samples, the filter can be defined [9] as

$$S[n] = \sum_{k=0}^N b_k PPG[n-k] - \sum_{k=1}^N a_k S[n-k] \quad (1)$$

where,  $a_k$  and  $b_k$  are the filter coefficients.  $S[n]$  is the result of applying the bandpass filter to the input signal  $PPG[n]$ . Fig. 5(b) is the result of applying Butterworth filter to the original APG signal measured at rest shown Fig. 5(a).




**Fig. 4.** Demonstrating the APG signals frequency bands (a) APG signal (b) the Fourier transform of the APG signal

B) *Second Derivative*: To increase accuracy and the detection rate of the inflection points and make interpretation easier, the second derivative of photoplethysmogram, also called the acceleration plethysmogram (APG), has been introduced. Because the peaks in the APG signal are more clearly defined than the peaks in the photoplethysmogram the heart rate can be more accurately detected using the APG, as shown in Fig. 5(c). In order to obtain the APG signal  $Z[n]$ , the second derivative will be applied to the filtered PPG signal  $S[n]$ .

$$\begin{aligned} S' [n] &= |S[n] - S[n-1]| \\ Z[n] &= |S' [n] - S' [n-1]| \end{aligned} \quad (2)$$

C) *Squaring*: At this stage the *a*-peak of the APG needs to be emphasized to distinguish it clearly for detection. This can be done by setting the negative parts of the signal equal to zero, followed by squaring the signal, resulting in  $y[n]$ .



**Fig. 5.** Algorithm structure. (a) original PPG signal (b) filtered PPG signal with Butterworth bandpass filter (c) second Derivative of PPG which is APG signal (d) squaring the APG signal (e) generating blocks of interest using two moving averages to detect *a*-waves

**IF**  $Z[n] < 0$  **THEN**

$Z[n] = 0$

**END**

$$y[n] = (Z[n])^2 \quad (3)$$

2) *Features Extraction*: In this stage the squared APG will be processed by two moving averages to generate blocks of interest and then detect *a*-waves. The moving averages are low pass filters with filter coefficients equal to the reciprocal of the window size.

A) *First Moving Average*: The first moving average is used to emphasize the *a* and *b* wave, shown as the dotted line in Fig. 5(e) for APG signal measured at rest.

$$MA_{Peak}[n] = \frac{1}{W_1}(y[n - (W_1 - 1)/2] + \dots + y[n] + \dots + y[n + (W_1 - 1)/2]) \tag{4}$$

where  $W_1 = 4I$  which is the window width of *ab* segment.

B) *Second Moving Average*: The purpose of the second moving average, shown as the solid line in Fig. 5(e) for APG signal measured at rest, is used as a threshold for the output of the first moving average.

$$MA_{OneBeat}[n] = \frac{1}{W_2}(y[n - (W_2 - 1)/2] + \dots + y[n] + \dots + y[n + (W_2 - 1)/2]) \tag{5}$$

where  $W_2 = 22I$  is the window width of approximately one beat.

When the amplitude of the first moving average filter ( $MA_{Peak}$ ) is greater than the amplitude of the second moving average filter ( $MA_{OneBeat}$ ), that part of the signal is selected as a block of interest, as follows:

```

IF  $MA_{Peak}[n] > MA_{OneBeat}[n]$  THEN
     $Blocks[n] = 1$ 
ELSE
     $Blocks[n] = 0$ 
END
    
```

Fig. 5(e) demonstrates the idea of using two filters to generate blocks of interest. Not all of the blocks are potential *a* and *b* waves. Some block are caused by noise and need to be eliminated.

The blocks associated with small width are considered as blocks caused by noise. Blocks which are smaller than half of the expected size for the *ab* interval are rejected. The expected size for the *ab* interval is based on the statistics for healthy adults, as described above. Blocks that is smaller than the expected width  $W_1$  that is expected for the *ab* interval will be rejected and the accepted blocks are considered to be containing *a*-wave.

The following statistical parameters were used to evaluate the algorithm:

$$\begin{aligned}
 Se &= \frac{TP}{TP + FN} \\
 +P &= \frac{TP}{TP + FP}
 \end{aligned}
 \tag{6}$$



**Table 1.** *a*-waves detection performance and HRV indices calculated using the APG

Record	No of beats	TP	FP	FN	SDNN	rMSSD
<b>A1</b>	26	26	0	0	0.024	0.013
<b>A2</b>	24	24	0	0	0.033	0.030
<b>B1</b>	17	17	0	0	0.080	0.108
<b>B2</b>	26	26	0	0	0.086	0.043
<b>C2</b>	20	20	0	0	0.105	0.118
<b>C3</b>	20	20	0	0	0.105	0.142
<b>D2</b>	22	22	0	0	0.059	0.051
<b>D3</b>	19	19	0	0	0.063	0.074
<b>E1</b>	22	22	0	0	0.045	0.056
<b>E2</b>	22	22	0	0	0.045	0.056
<b>E3</b>	19	19	0	0	0.095	0.144
<b>G2</b>	30	30	0	0	0.019	0.013
<b>G3</b>	19	19	0	0	0.056	0.033
<b>H3</b>	23	23	0	0	0.058	0.051
<b>I1</b>	22	22	0	0	0.034	0.037
<b>I2</b>	17	17	0	0	0.052	0.079
<b>J2</b>	23	23	0	0	0.0778	0.048
<b>L2</b>	24	24	0	0	0.033	0.030
<b>L3</b>	24	24	0	0	0.037	0.029
<b>N2</b>	18	18	0	0	0.032	0.039
<b>N3</b>	20	20	0	0	0.064	0.053
<b>O1</b>	24	24	0	0	0.043	0.037
<b>O2</b>	17	17	0	0	0.0459	0.050
<b>P1</b>	26	26	0	0	0.083	0.073
<b>P2</b>	20	20	0	0	0.129	0.137
<b>Q1</b>	22	22	0	0	0.036	0.043
<b>Q2</b>	18	18	0	0	0.140	0.161

True Positive (TP): *a* wave has been classified as *a* wave.

False Negative (FN): *a* wave has been missed.

False Positive (FP): Non- *a* wave classified as *a* wave.

The sensitivity (*Se*) is the percentage of true *a* waves that were correctly detected by the algorithm. The positive predictivity (*+P*) is the percentage of detected *a* waves which are real *a* waves.

Table 1 shows the result of  $a$  waves detection in 27 different records of collected APG at rest, containing a total of 584 heart beats. The number of false negatives (FN) and false positives (FP) were zero. The overall average sensitivity for  $a$  waves detection was 100% and the positive predictivity was 100%.

### 3.2 HRV Analysis

Once the location of each  $a$  wave is detected the duration of the consecutive  $aa$  intervals can be determined. Two time domain parameters are calculated and compared. These parameters are often used with ECG signals. The first parameter, the  $SDNN$ , is the standard deviation of the duration of heart beats; normal-to-normal R-to-R intervals in ECG signal. The second parameter, the  $rMSSD$ , is the root-mean square of the difference of successive heart beats or R-to-R intervals in ECG signals. Here, the RR interval is replaced by  $aa$  intervals. These were calculated in the previous step.

$$SDNN = \sqrt{\frac{1}{N} \sum_i (aa)^2 - \left( \frac{1}{N} \sum_i aa \right)^2} \quad (7)$$

$$rMSSD = \sqrt{\frac{1}{N} \sum_i (aa)^2} \quad (8)$$

where is  $N$  the total number of  $aa$  intervals,  $i$  is a counter.

## 4 Results

The  $SDNN$  and the  $rMSSD$  index were calculated for 27 subjects at rest, using recordings of the PPG of 20 seconds. The results are shown in table 1.


In order to calculate the correlation between the heart rate and the HRV indices, the correlation coefficient is used as follows:

$$r = \frac{cov(x, y)}{\sigma_x \sigma_y} \quad (9)$$

where is  $r$  the correlation coefficient,  $cov(x, y)$  is the covariance between data  $x$  and data  $y$ ,  $\sigma_x$  is the standard deviation of data  $x$  and  $\sigma_y$  is the standard deviation of data  $y$ .

As can be seen from Fig. 6 (a) and Fig. 6 (b) there is a negative correlation between the heart rate and the heart rate variability indices. The  $rMSSD$  index is more negative correlated with the HR ( $r = -0.565$ ) than the  $SDNN$  index ( $r = -0.39$ ).

Fig. 6 (c) shows the correlation between the two HRV indices. As can be expected, there is a strong positive correlation ( $r = 0.894$ ).



**Fig. 6.** Relations between (a) HR and *SDNN* (b) HR and *rMSSD* (c) *SDNN* and *rMSSD* calculated from APG signals for all subjects measured at rest

## 5 Conclusions

The acceleration photoplethysmogram can be used to calculate heart rate variability.

The accurate detection of *a*-waves leads to an accurate detection of *aa* interval and consequently reliable HRV measurement. Therefore, the proposed *a*-wave detection algorithm achieved an overall average sensitivity 100% and a positive predictivity was 100%. over 27 PPG records measured at rest, containing a total of 584 heart beats.

As discussed above, HRV indices which are normally used with ECG signals can also be applied to APG signals. The negative correlation between HR and both indices *SDNN* and *rMSSD* confirmed that the 20 seconds recordings of APG signal are suitable for short duration signal. As expected the *rMSSD* index decreased more than *SDNN*. The reason is *rMSSD* used to estimate the short-term components of HRV, related to the parasympathetic activity while *SDNN* is related to overall HRV. There is a strong positive correlation between the two HRV indices, indicating that the 20 second APG recordings are sufficient to reliably measure the HRV. The usage of APG can be useful for HRV analysis and identification of individuals at risk.

**Acknowledgements.** The authors would like to thank Dr. Aya Matsuyama for collecting the data.

## References

1. Berntson, G., et al.: Heart rate variability: origins, methods, and interpretive caveats. *Psychophysiology* 34(6), 623–648 (1997)
2. Giardino, N., Lehrer, P., Edelberg, R.: Comparison of finger plethysmograph to ECG in the measurement of heart rate variability. *Psychophysiology* 39(2), 246–253 (2002)
3. Lu, S., et al.: Can photoplethysmography variability serve as an alternative approach to obtain heart rate variability information? *J. Clin. Monit. Comput.* 22(1), 23–29 (2008)
4. Jianfeng, W., Zhiqian, Y., Jianling, W.: An Improved Pre-processing Approach for Photoplethysmographic Signal. *IEEE-EMBS in Engineering in Medicine and Biology Society* (2005)
5. Takazawa, K., et al.: Clinical usefulness of the second derivative of a plethysmogram (acceleration plethysmogram). *Cardiology* 23, 207–217 (1993)
6. Taniguchi, K., et al.: Evaluating the surgeon's stress when using surgical assistant robots. In: *The 16th IEEE International Symposium on Robot and Human Interactive Communication, ROMAN 2007*, pp. 888–893 (2007)
7. Matsuyama, A.: *ECG and APG Signal Analysis during Exercise in a Hot Environment*, in School of Engineering and Information Technology. Charles Darwin University, Darwin (2009)
8. Oppenheim, A., Shafer, R.: *Discrete-time Signal Processing*. Prentice-Hall, NJ
9. Rangayyan, R.M.: *Biomedical Signal Analysis A Case-Study Approach*. In: Akay, M. (ed.). *IEEE Press Series on Biomedical Engineering*, p. 516. John Wiley & Sons, New York (2002)

# Emotional Face Perception in Healthy Subjects and Parkinson's Disease: An Effective Connectivity Study

Elvis Lira da Silva<sup>1,4</sup>, Gabriela Castellano<sup>1,4</sup>, João Ricardo Sato<sup>2,3,4</sup>,  
Ellison Fernando Cardoso<sup>2,4</sup>, and Edson Amaro Jr.<sup>2,4</sup>

<sup>1</sup>Institute of Physics "Gleb Wataghin", University of Campinas - UNICAMP, Campinas, Brazil

<sup>2</sup>NIF, LIM-44, Department of Radiology, University of São Paulo, São Paulo, Brazil

<sup>3</sup>Center of Mathematics, Computation and Cognition

Universidade Federal do ABC, Santo André, Brazil

<sup>4</sup>CInAPCe Program (Cooperação Interinstitucional de Apoio a Pesquisas sobre o Cérebro)  
São Paulo State, Brazil

elvislir@ifi.unicamp.br

**Abstract.** We investigated the neural connectivity induced by face presentation with different emotional valences in Parkinson's disease (PD) patients and a control group of healthy, drug-free volunteers, using event-related fMRI in a parametric design. The focus of this work was applying Dynamic Causal Modelling (DCM), an approach that allows the assessment of effective connectivity within cortical networks [1], to the study of effective connectivity between maximally activated brain regions in response to passive viewing of facial stimuli. A connectivity model was built based on the literature and in our fMRI analyses, which included the fusiform gyrus (FG), the anterior cingulate cortex (ACG), the dorsolateral prefrontal cortex (DLPFC) and the dorso-medial prefrontal cortex (DMPFC) areas. The results showed differences in connectivity between the PD group and the control group. We found that the intrinsic connectivities among DLPFC/DMPFC and FG, DLPFC/DMPFC and ACG, were higher in PD patients than in healthy subjects, while the effective connectivity among FG and ACG was lower in PD patients.

**Keywords:** Dynamic causal modelling, fMRI, connectivity, Parkinson's disease, Face perception.

## 1 Introduction

The branch of Neuroscience that studies functional integration between cerebral areas has recently shown a significant growth. Functional integration refers to the interactions among specialized neuronal populations, where the integration is mediated by the so called effective connectivity. Effective connectivity is defined as the influence that regions, which encompass given neuronal populations, exert on each other. Effective connectivity studies are important to assess how different areas, involved in a particular brain processing task, are related.

Facial perception is one of the fundamental tasks in our daily life and plays a critical role in social interactions. It is a highly developed visual ability in humans and it

is mediated by activation in a distributed neural system that encompasses visual, limbic, and prefrontal regions [2,3]. Facial perception with different emotional valences involves the emotional recognition that is related to the activity of amygdala, insula, orbitofrontal cortex and ventral striatum. The areas linked to emotional regulation include the anterior cingulate, dorsolateral and medial prefrontal cortices [4, 5].

To investigate effective connectivity within the distributed cortical network for face perception, we combined conventional Statistical Parametric Mapping (SPM) and the technique of Dynamic Causal Modelling (DCM) [1]. DCM consists of a theoretical-experimental approach that treats the brain as a nonlinear deterministic dynamic system. DCM regards an experiment of fMRI as a designed perturbation of neuronal dynamics that is distributed throughout a system of coupled anatomical nodes to change region-specific neuronal activity [1]. In practical terms, a reasonably realistic neuronal model of interacting cortical regions is built, with neuro-physiologically meaningful parameters. This model is supplemented with a forward model of how neuronal or synaptic activity is transformed into a measured response, and the parameters of this model can be estimated by attempting to match the predicted BOLD signal to the observed BOLD signal. DCM has been previously used to investigate visual perception and visual imagery of faces [6] and face perception in healthy subjects [2, 7].

The etiology of depression in PD is multifactorial. Although neurological deficits might lead to depressive symptoms, increasing evidence suggests that PD depression is also related to the degenerative process that underlies the disease. Our previous studies in this field have shown a decreased activation of the dorso-medial pre-frontal cortex (DMPFC), which is restored with different treatments. Also, treatment has produced changes mainly in areas that are responsible for emotional regulation, but also in the fusiform gyrus [8,9].

In this study we investigated the effective connectivity induced by face presentation with different emotional valences in Parkinson's disease (PD) patients and in a control group of healthy, drug-free volunteers. Bearing in mind the physiological model of facial perception, and the changes induced by depression in PD, we conducted a ROI based analysis of the following areas: DMPFC (including the anterior cingulate cortex (ACG)), dorsolateral prefrontal cortex (DLPFC) and fusiform gyrus (FG). We hypothesized that the modulatory connectivity between the DLPFC/DMPFC and the FG, induced by faces, would change in both groups. In addition, PD patients compared to normal volunteers would present changes in the modulatory and intrinsic connectivity between the DLPFC/DMPFC and the FG.

## 2 Methods

### Study Population

We studied 17 male patients with idiopathic Parkinson's disease according to the UK Parkinson's Disease Brain Bank. We excluded patients with severe motor symptoms and severe motor fluctuation; ferromagnetic metallic implants; seizures; major head trauma; dementia; previous neurological surgery; depression with psychotic symptoms; or those under antidepressants treatment within one year of the beginning of the study. All patients were recruited from the Movement Disorders Clinics of

“Hospital das Clínicas” - University of São Paulo (São Paulo - Brazil) and all gave their written informed consent. The study was approved by the Ethics Committee of the University of São Paulo (Project Approval number: 414/03). Healthy volunteers were paired with patients in age and level of education. The demographic data is shown in Table 1.

**Table 1.** Demographic and clinical data from healthy subjects and patients with idiopathic Parkinson's disease


Clinical and Demographic Characteristics		
	Parkinson's Disease	Healthy Subjects
Age	62.5 (6.5)	59.5 (6.4)
Mini-Mental State Examination (MMSE)	28.1 (2.1)	27.5 (1.6)
Time since diagnosis	11.2 (6.9)	-
Daily dose of L-dopa (mg)	875 (500)	-
Hoehn and Yahr	2.1 (0.6)	-
Unified Parkinson's Disease Rating Scale (UPDRS)	30.1 (8.1)	-

### Image Acquisition

All images were acquired in a 1.5 T GE scanner, equipped with a 33 mT/m gradient. The images were oriented according to the AC–PC line; and 168 brain volumes were acquired, with 15 slices each (7 mm thickness, 0.7 gap), 64×64 pixels matrix, 20×20 cm FOV, 90° flip angle, 2.0 s TR, 40 ms TE, using a gradient echo EPI acquisition.

### fMRI Paradigm

Eckmann's faces were morphed to produce neutral, low and high intensities of sadness, as shown in Figure 1. An event-related fMRI paradigm, similar to Fu et al. [10] was used. Facial stimuli and baseline trials (crosshair fixation) were presented in random order. Each trial and control condition was presented for 2 s, and the inter-trial interval was randomly varied according to a Poisson distribution (2–12 s; mean 5 s).



**Fig. 1.** An example of stimuli set used. Eckmann's faces were morphed to produce neutral, low and high intensities of sadness. Facial stimuli and baseline trials (crosshair fixation) were presented in random order in the event related fMRI paradigm.

## Image Analysis

The fMRI statistical and DCM analyses were performed using the free software Statistical Parametric Mapping (SPM8, [www.fil.ion.ucl.ac.uk/spm/](http://www.fil.ion.ucl.ac.uk/spm/)). All volumes were slice time corrected, realigned to the middle volume, corrected for motion artifacts, mean-adjusted by proportional scaling, normalized into standard stereotactic space (template provided by SPM8), and smoothed using an 8 mm full-width at half-maximum (FWHM) Gaussian kernel. The time series were high-pass filtered to eliminate low-frequency components (filter width = 128 s) and adjusted for systematic differences across trials.

## Dynamic Causal Modelling

Dynamic Causal Modelling (DCM) is an approach that allows the assessment of effective connectivity within cortical networks [1]. The central idea behind DCM is to treat the brain as a deterministic nonlinear dynamic system that is subject to inputs and produces outputs. In this way, the effective connectivity is parameterized in terms of neuronal activity in different regions. DCM is constructed by a bilinear approximation given by equation 1:

$$\dot{z} \approx Az + \sum_j u_j B^j z + Cu, \quad (1)$$

where  $z$  is a vector that represents the change in activity with time of the neuronal regions involved,  $A$  is a matrix that represents the first order connectivity among the regions in the absence of input,  $B^j$  are matrices that account for the change in coupling among the regions induced by the  $j$ th input  $u_j$ , and the matrix  $C$  embodies the extrinsic influences of inputs  $u$  on the neuronal activity of the regions. This approximation allows the dynamics of the system to depend on these three groups of parameters: parameters that mediate intrinsic coupling among the areas (matrix  $A$ ), parameters that mediate the influence of extrinsic inputs on the areas (matrix  $C$ ) and (bilinear) parameters that allow the extrinsic inputs to modulate that coupling (matrix  $B$ ) [1]. The model depends on the experimental design, where the extrinsic inputs enter by two ways: directly influencing the areas (driving inputs) and/or influencing the coupling among the areas (contextual inputs).

### 2.1 Effective Connectivity Analysis

Initially, a standard fMRI (General Linear Model – GLM) analysis was used for locating brain responses to the experimental task (face perception with different emotional valences). Individual maps of activations were generated using voxel based analysis. Next, second-level analysis was used to generate maps of the group using a one sample t-test, with corrected (FWE) p-value < 0.05.

Based on our analysis of the healthy group data and on related literature data [4,5], we determined which areas should enter the DCM model. The model included 3 areas: left and right Fusiform Gyri (FG), Anterior Cingulate Gyrus (ACG), and Dorsolateral Prefrontal Cortex (DLPFC) / Dorsomedial Prefrontal Cortex (DMPFC). These regions of interest (ROIs) were defined using masks created with the WFU Pickatlas software [11, 12].



After delimiting the brain areas aforementioned in the individual brain activation map, averaged time series of voxels limited by a sphere of 8 mm were extracted. The centres of these spheres were located in the local maxima of the activation map for each anatomical area included in the model. This procedure was performed for each of the subjects. These three volumes of interest (VOIs) were identified for each individual subject. Mean localization and t-values of these areas are shown in Table 2 for healthy subjects and PD patients. All three VOIs were reliably delineated in 16 of the 19 healthy subjects and 10 of the 17 PD patients (Table 2) ( $p < 0.05$ , uncorrected).

**Table 2.** Regions of the DCM model for healthy subjects and PD patients. The x, y, z, columns give the average coordinates across all subjects for the location of each region (with the standard deviation – SD – in parentheses). The t column shows the average t-statistics across subjects in the first-level analysis (and the SD in parentheses). “L” and “R” mean left and right, respectively.

Healthy subjects				
Regions	x	y	z	T
FG (R)	31.3(7.1)	-59.5(6.9)	-18.4(3.1)	4.0(1.5)
ACG (R)	3.7(4.0)	14.2(9.9)	20(16,9)	2.53(0.8)
DMPFC/ DLPFC(R)	21.8(13.8)	19.3(14.7)	56(18)	4.6(1.5)
FG (L)	-31.3(7.9)	-58(14.4)	-20(4.9)	3.8(1.1)
ACG (L)	-7(2)	22.3(8.2)	8.3(14.9)	2.5(0.5)
DMPFC/ DLPFC(L)	-30.1(13.1)	-2.5(8.7)	63.2(8.7)	4.0(1.4)
PD patients				
Regions	x	y	z	T
FG (R)	30(8.6)	48.6(16.5)	-21.1(7.9)	4.1(1.2)
ACG (R)	5.2(3.2)	20.3(5.6)	28.1(12,8)	2.9(0.6)
DMPFC/ DLPFC(R)	17(15.2)	9.0(8.3)	59.7(13,3)	4.1(1.0)
FG(L)	-22(0)	-67.5(10.6)	-16.0(0)	4.9(1.2)
ACG (L)	-11(0)	26.0(0)	28.0(0)	2.9(0)
DMPFC/ DLPFC(L)	-32(2.8)	-3.5(10.6)	63.0(4.2)	4.3(0.7)

Initially the DCM model was estimated separately for each subject. In order to generalize our results to the population level, the estimated connection strengths from that analysis were then subjected to a second-level analysis (using Matlab functions), where the significance of inferred connections was tested using one-sample t-tests against the null hypothesis that the connection strength was equal to zero.

A model that connects the different areas was defined according to our hypotheses, taking into account the FG, ACG and prefrontal cortex regions. As a first model we admitted that all regions interacted with one another; and that the face stimuli entered


the model as driving inputs only on FG. The face stimuli also entered the model as contextual factors on all connections. From this model, the intrinsic connections between each of the components of the model were estimated using a DCM analysis.

### 3 Results

The fMRI analysis using the GLM approach (see above) showed significant BOLD effect in the FG, ACG, DLPFC and DMPFC brain areas ( $p < 0.05$ , FWE corrected).

The group connectivity maps for healthy subjects and PD patients are shown in Figures 2 and 3, respectively. Only connections that were significant (non-zero mean) after the one-sample t-test ( $p < 0.05$ ) are shown. Filled arrows among areas (and their values) are the results of intrinsic connections, which do not depend on external stimuli. Dashed arrows (and their values) are the results of modulatory connections given by contextual inputs; these values increase or decrease the influence of the intrinsic connections and depend on the external stimulus. Filled arrows among stimuli and the FG area are results of direct influence of stimuli on the areas. Values are shown when significant (one-sample t-test) for each of the experimental conditions (neutral, low and high sadness).

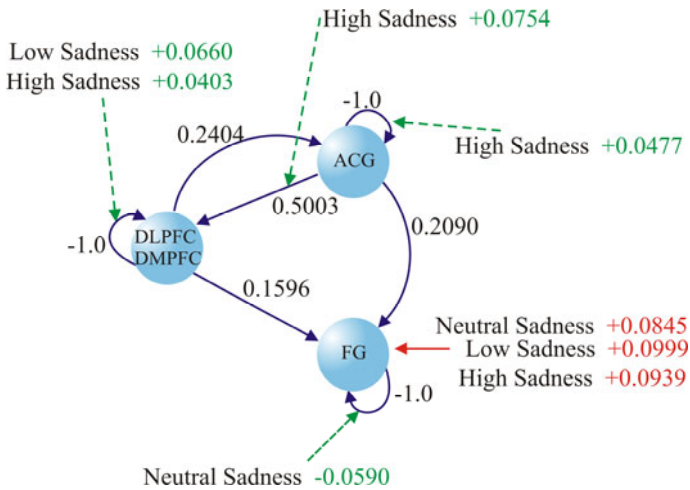
Control subjects (Figure 2) showed an increase in activity in the FG area induced by modulation of connectivity by neutral and high sadness faces (FG → FG connection) and in the DLPFC/DMPFC area induced by modulation of connectivity by all faces (DLPFC / DMPFC → DLPFC / DMPFC and ACG → DLPFC / DMPFC connections). Neutral faces increased activity in the ACG area (FG → ACG connection) and decreased activity in the FG (DLPFC / DMPFC → FG connection).



**Fig. 2.** DCM results for healthy subjects. Filled arrows among areas (and their values) are the results of intrinsic connections, filled arrows among stimuli and the FG area are the results of direct influence of the stimuli on this area, dashed arrows (and their values) are the results of modulatory connections.

For PD patients (Figure 3), we notice an increase in activity in the DLPFC and DMPFC areas induced by the modulation of connectivity by low sadness (DLPFC/DMPFC → DLPFC/DMPFC connection) and high sadness faces (DLPFC/DMPFC → DLPFC/DMPFC and ACG → DLPFC/DMPFC connections). We can also see an increase in activity in the ACG area induced by the modulation of connectivity by high sadness faces (ACG → ACG connection). In addition, we see a decrease in activity of the FG area induced by the modulation of connectivity by neutral faces (FG → FG).

Using a two-sample t-test (through Matlab functions) we compared the connections of the two groups (healthy and PD) and found a significant difference among the intrinsic connections (black lines in Figures 2 and 3) DLPFC/DMPFC → ACG (p-value of 0.0345), DLPFC/DMPFC → FG (p-value of 0.0303) and FG → ACG (p-value of 0.0487).



**Fig. 3.** DCM results for PD patients. Filled and dashed arrows represent the same as for Figure 2.

We found that the effective couplings DLPFC/DMPFC → FG and DLPFC/DMPFC → ACG were higher in PD patients than in healthy subjects, while the effective coupling FG → ACG was lower in PD patients. The results are shown in Table 3.

To compare the connectivity patterns between the stimuli of highly sad and neutral faces, we used a paired t-test among these conditions within each group (PD patients and controls), for every connection (intrinsic, modulatory and direct influence) in every region. We found a significant difference between the connectivity for these two conditions in healthy subjects and PD patients (Table 4). The connection FG → DLPFC (intrinsic and modulatory) is not shown in the table because it was not included in the model. For clarity only the results for the comparison between the neutral and high sadness conditions are shown.

**Table 3.** Comparison between healthy subjects and PD patients connections. Significant differences are marked with bold font. Connections marked with asterisk are higher in PD patients than in healthy subjects.

Comparison between groups			
Connections	Healthy Subjects	PD Patients	P-value
Intrinsic Connections			
FG→FG	-1.000	-1.000	-
<b>FG→ACG</b>	<b>0.4908</b>	<b>0.2620</b>	<b>0.0487</b>
ACG→FG	0.2080	0.2090	0.9850
ACG→ACG	-1.000	-1.000	-
ACG→DLPFC	0.5636	0.5003	0.6691
<b>DLPFC→FG*</b>	<b>0.0361</b>	<b>0.1596</b>	<b>0.0303</b>
<b>DLPFC→ACG*</b>	<b>0.0855</b>	<b>0.2404</b>	<b>0.0345</b>
DLPFC→DLPFC	-1.000	-1.000	-
FG→FG	0.0587	0.0590	0.9896
<b>FG→ACG</b>	<b>0.0557</b>	<b>-0.0820</b>	<b>0.0047</b>
ACG→FG	-0.0059	0.0017	0.6207
ACG→ACG	0.0014	0.0404	0.1135
ACG→DLPFC	0.0530	0.0390	0.6259
<b>DLPFC→FG*</b>	<b>-0.0373</b>	<b>0.0128</b>	<b>0.0345</b>
DLPFC→ACG	-0.0300	-0.0055	0.3046
DLPFC→DLPFC	0.0183	0.0261	0.6131
Modulatory Connections (Low Sadness)			
FG→FG	-0.0192	0.0180	0.2513
FG→ACG	0.0102	0.0236	0.7676
ACG→FG	-0.0245	-0.0010	0.3873
ACG→ACG	0.0042	0.0387	0.1190
ACG→DLPFC	0.0943	0.0671	0.5274
<b>DLPFC→FG*</b>	<b>-0.0246</b>	<b>0.0239</b>	<b>0.0324</b>
DLPFC→ACG	-0.0024	0.0011	0.8531
DLPFC→DLPFC	0.0621	0.0660	0.8996
Modulatory Connections (High Sadness)			
FG→FG	0.0337	0.0234	0.6424
FG→ACG	0.0047	0.0076	0.9513
ACG→FG	0.0226	0.0111	0.5818

**Table 3.** (Continued)

Comparison between groups			
Connections	Healthy Subjects	PD Patients	P-value
<b>ACG→ACG*</b>	<b>0.0087</b>	<b>0.0477</b>	<b>0.0344</b>
ACG→DLPFC	0.0649	0.0754	0.7691
DLPFC→FG	0.0179	0.0068	0.6267
DLPFC→ACG	0.0061	0.0095	0.8716
DLPFC→DLPFC	0.0404	0.0403	0.9974
Extrinsic Connections			
Stimulus → FG (All others connections not included in the model)			
<b>Neutral</b>	<b>0.1530</b>	<b>0.0845</b>	<b>0.0330</b>
Low Sadness	0.1469	0.0999	0.2165
High Sadness	0.1390	0.0939	0.2400

**Table 4.** Comparison between neutral and high sadness faces conditions, for healthy subjects and PD patients connections. Significant differences are marked with bold font. Connections marked with asterisk are higher in the high sadness than in the neutral condition. Healthy subjects show many more significant differences between high sadness and neutral faces than PD patients.

Healthy Subjects (Comparison between conditions)			
Modulatory Connections			
Connections	Neutral	High Sadness	P-value
FG→FG	0.0587	0.0337	0.2430
FG→ACG	0.0557	0.0047	0.1149
<b>ACG→FG*</b>	<b>-0.0059</b>	<b>0.0226</b>	<b>0.0459</b>
ACG→ACG	0.0014	0.0087	0.6311
ACG→DLPFC	0.0530	0.0649	0.5454
<b>DLPFC→FG*</b>	<b>-0.0373</b>	<b>0.0179</b>	<b>0.0100</b>
<b>DLPFC→ACG*</b>	<b>-0.0300</b>	<b>0.0061</b>	<b>0.0346</b>
<b>DLPFC→DLPFC*</b>	<b>0.0183</b>	<b>0.0404</b>	<b>0.0470</b>
Extrinsic Connections			
<b>Stimulus → FG</b> (All others extrinsic connections were not included in the model)	Neutral	High Sadness	P-value
	<b>0.1530</b>	<b>0.1390</b>	<b>0.0414</b>

**Table 4.** (Continued)

PD Patients (Comparison between conditions)			
Modulatory Connections			
	Neutral	High Sadness	P-value
FG→FG	0.0590	0.0234	0.1538
<b>FG→ACG*</b>	<b>-0.0820</b>	<b>0.0076</b>	<b>0.0427</b>
ACG→FG	0.0017	0.0111	0.5809
ACG→ACG	0.0404	0.0477	0.7515
ACG→DLPFC	0.0390	0.0754	0.2473
DLPFC→FG	0.0128	0.0068	0.7276
DLPFC→ACG	-0.0055	0.0095	0.6463
Healthy Subjects (Comparison between conditions)			
DLPFC→DLPFC	0.0261	0.0403	0.5526
Extrinsic Connections			
Stimulus →FG (All others extrinsic connections were not included in the model)	Neutral	High Sadness	P-value
	0.0845	0.0939	0.4041

## 4 Discussion

Dynamic Causal Modelling was able to detect significant connections between the known areas involved in facial perception and emotional regulation in both groups.

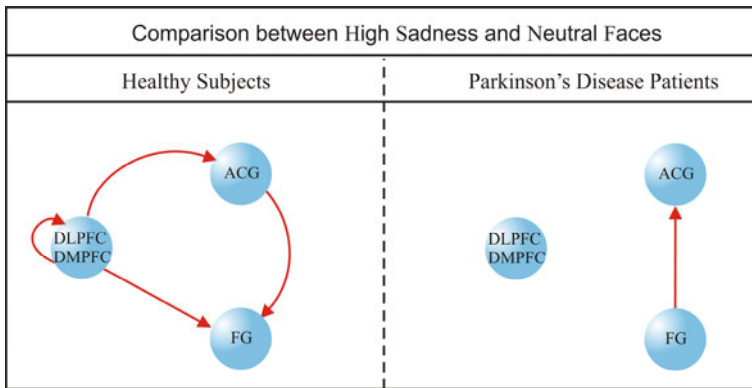
In healthy subjects, the increase of sad content in the stimuli induced a progressive increase in the connectivity between the prefrontal areas and ACG with FG. This supports the current model of emotional stimuli, showing a top-down modulation exercised by the pre-frontal areas over the fusiform gyrus, which increases as the emotional valence increases. The connectivity for the sad faces stimulus has greater modulation by prefrontal areas in the ACG and FG, which is in agreement to the article by Phillips et al. [4]. According to Phillips et al., the prefrontal areas are responsible for the regulation of the emotional state and the perception of emotion.

Therefore, we may conclude that in healthy subjects, the prefrontal area regulated the emotional state due to the presentation of the sad faces stimulus. On the other hand, we did not observe this difference in the group of Parkinson's disease patients: the t-test did not show significant differences in connectivity (DLPFC/DMPFC → FG and DLPFC/DMPFC → ACG) between the different face conditions (Table 4 and Figure 4). In fact, facial emotion recognition deficits have been reported in early Parkinson's disease [13-15]. The dysfunction in dopaminergic pathways (nigrostriatal and mesolimbic) leads to changes in the frontal-striatal circuitry, which are critical for emotional recognition [16]. Differently than what occurred for healthy volunteers, changes in emotional valence did not increase the connectivity between the DLPFC

and ACG with the FG in PD patients. In these patients, the only significant change observed when going from neutral to sad faces was an increased connectivity between the FG and the ACG (Figure 4). We might speculate that these changes might be a consequence of an ineffective top-down modulation of the DLPFC.

When comparing healthy volunteers with PD patients, we observed an increased intrinsic and modulatory connectivity between DLPFC with both the ACG and the FG in the latter. These areas are part of the system responsible for the regulation of emotional stimuli. As PD leads to a dysfunction in the DLPFC (as we observed in fMRI maps), we also see that these changes influence the behavior of areas that are connected to them, leading to a global dysfunction of the system that results in an abnormal perception.

Some limitations should be entertained. PD is a progressive disorder, which follows a known course of involvement of brain areas. We can not infer from our data if changes observed in the DLPFC in PD are due to changes in dopaminergic pathways, as the main areas of this system, such as the ventral tegmental area and the striatum, were not included in the model. We decided not to include these areas because they were not activated during the task in most of the individuals. Future studies should evaluate fronto-striatal circuits in PD.



**Fig. 4.** Comparison between high sadness and neutral faces, we used a paired t-test among these conditions. These connections are higher in the high sadness than in the neutral condition.

## 5 Conclusions

Using DCM, we explored the effective connectivity of the main cerebral regions involved in responses to facial stimuli with different intensities of sadness, for both PD patients and healthy subjects. The results showed differences in connectivity between the PD group and the control group, suggesting that these changes in connectivity can play an important role in Parkinson's disease and may thus provide insights on the underlying mechanisms of PD impairment in processing human facial emotions.

However, as the success of DCM is dependent on the experimental design and on the specified interacting regions model, other models involving those regions should be tested for a more definitive conclusion.

## References

1. Friston, K., Harrison, L., Penny, W.: Dynamic causal modeling. *NeuroImage* 19, 1273–1302 (2003)
2. Fairhall, S., Ishai, A.: Effective Connectivity within the Distributed Cortical Network for Face Perception. *Cerebral Cortex* 17(10), 2400–2406 (2007)
3. Haxby, J., Hoffman, E., Gobbini, M.: The distributed human neural system for face perception. *Trends in Cognitive Sciences* 4(6), 223–233 (2000)
4. Phillips, M., Drevets, W., Rauch, S., Lane, R.: Neurobiology of emotion perception I: The neural basis of normal emotion perception. *Biological Psychiatry* 54, 504–514 (2003)
5. Phillips, M., Drevets, W., Rauch, S., Lane, R.: Neurobiology of emotion perception II: Implications for major psychiatric disorders. *Biological Psychiatry* 54, 515–528 (2003)
6. Mechelli, A., Price, C., Friston, K., Ishai, A.: Where bottom-up meets top-down: neuronal interactions during perception and imagery. *Cerebral Cortex* 14, 1256–1265 (2004)
7. Rotshtein, P., Vuilleumier, P., Winston, J., Driver, J., Dolan, R.: Distinct and convergent visual processing of high and low spatial frequency information in faces. *Cerebral Cortex* 17, 2713–2724 (2007)
8. Cardoso, E., Fregni, F., Maia, F., Boggio, P., Myczkowski, M., Coracini, K., Vieira, A., Melo, L., Sato, J., Marcolin, M., Rigonatti, S., Cruz, A., Barbosa, E., Amaro, E.: rTMS treatment for depression in Parkinson's disease increases BOLD responses in the left prefrontal cortex. *International Journal of Neuropsychopharmacology* 11(2), 173–183 (2007)
9. Cardoso, E.F., Maia, F.M., Fregni, F., Myczkowski, M.L., Melo, L.M., Sato, J.R., Marcolin, M.A., Rigonatti, S.P., Cruz Jr., A.C., Barbosa, E.R., Amaro Jr., E.: Depression in Parkinson's disease: convergence from voxel-based morphometry and functional magnetic resonance imaging in the limbic thalamus. *NeuroImage* 47, 467–472 (2009)
10. Fu, C., Williams, S., Cleare, A., Brammer, M., Walsh, N., Kim, J., Andrew, C., Pich, E., Williams, P., Reed, L., et al.: Attenuation of the neural response to sad faces in major depression by antidepressant treatment: a prospective, event-related functional magnetic resonance imaging study. *Archives of General Psychiatry* 61, 877–889 (2004)
11. Maldjian, J., Laurienti, P., Burdette, J., Kraft, R.: An Automated Method for Neuroanatomic and Cytoarchitectonic Atlas-based Interrogation of fMRI Data Sets. *NeuroImage* 19, 1233–1239 (2003)
12. Maldjian, J., Laurienti, P., Burdette, J.: Precentral Gyrus Discrepancy in Electronic Versions of the Talairach Atlas. *Neuroimage* 21(1), 450–455 (2004)
13. Assogna, F., Pontieri, F., Caltagirone, C., Spalletta, G.: The recognition of facial emotion expressions in Parkinson's disease. *European Neuropsychopharmacology* 18, 835–848 (2008)
14. Dujardin, K., Blairy, S., Defebvre, L., Duhemb, S., Noël, Y., Hess, U., Destée, A.: Deficits in decoding emotional facial expressions in Parkinson's disease. *Neuropsychologia* 42, 239–250 (2004)
15. Sprengelmeyer, R., Young, A., Mahna, K., Schroeder, U., Woitalla, D., Büttner, T., Kuhn, W., Przuntek, H.: Facial expression recognition in people with medicated and unmedicated Parkinson's disease. *Neuropsychologia* 41, 1047–1057 (2003)
16. Yip, J.T.H., Lee, T.M.C., Ho, S.L., Tsang, K., Li, L.S.W.: Emotion Recognition in Patients With Idiopathic Parkinson's Disease. *Movement Disorders* 18, 1115–1122 (2003)



# Biometric Authentication with Electroencephalograms: Evaluation of Its Suitability Using Visual Evoked Potentials

André Zúquete<sup>1</sup>, Bruno Quintela<sup>2</sup>, and João Paulo Silva Cunha<sup>1</sup>

<sup>1</sup>IEETA / University of Aveiro, Aveiro, Portugal

<sup>2</sup>University of Aveiro, Aveiro, Portugal

**Abstract.** This paper studies the suitability of brain activity, namely electroencephalogram signals, as raw material for conducting biometric authentication of individuals. Brain responses were extracted in particular scenarios, namely with visual stimulation leading to biological brain responses known as visual evoked potentials. In our study, we evaluated a novel method, using only 8 occipital electrodes and the energy of differential EEG signals, to extract information about the subjects for further use as their biometric features. To classify the features obtained from each individual we used a one-class classifier per subject. These classifiers are trained only with target class features, which is the correct procedure to apply in biometric authentication scenarios. Two types of one-class classifiers were tested, K-Nearest Neighbor and Support Vector Data Description. Two other classifier architectures were also studied, both resulting from the combination of the two previously mentioned classifiers. After testing these classifiers with the features extracted from 70 subjects, the results showed that brain responses to visual stimuli are suitable for an accurate biometric authentication.

## 1 Introduction

This document presents a study on the suitability of induced electroencephalograms (EEGs) for implementing high-quality, practical biometric authentication systems. EEGs are impossible to forge because they reflect the inner self of a person, and they are likely to be different from person to person when performing similar mental activities. However, EEGs are complex and noisy signals, being affected by different brain activities and other body activities as well. Thus, we conducted our study with EEG signals measured in particular scenarios, namely with visual stimulations leading to very focused brain activities known as Visual Evoked Potentials (VEP). To the best of our knowledge, this is the first work on EEG-based authentication using VEPs, though some works exist on EEG-based identification using VEPs and EEG-based authentication using other brain activity stimuli (e.g. specific imagining tasks).

A biometric authentication system has four fundamental requirements [8]:

- *Universality*: it should be possible to use the system with all persons.
- *Uniqueness*: the system should be able to separate different persons with a reasonably low failure probability.

- *Constancy*: biometric characteristics of the persons should remain fairly constant for a reasonable time (months, years).
- *Collectability*: biometric values should be easy to obtain, easy to quantify and cause no discomfort.

Considering the first requirement (universality), we believe that only a small percentage of people could not use the presented EEG authentication procedure. As we used the perception of simple drawings for triggering EEG signals, which are the raw input for the biometric authentication system, people with severe visual impairments or blindness cannot be authenticated; other evoked potentials must be studied for them (e.g. sound-evoked potentials).

Considering the second requirement (uniqueness), we did an empirical observation of the separation of individuals among a limited population of 70 people for which we had several EEG samples. Therefore, we have no proof that it will work on other populations, but we cannot as well anticipate any reason for not working. Furthermore, the proposed system is configurable, as most biometric authentication system are, and allows administrators to tune several parameters for achieving the required levels of quality in the separation of individuals.

Considering the third requirement (constancy), our study is still limited. Our authentication system uses images to trigger brain cognitive activities, which are then measured and classified. Cognitive activities may be affected by several factors, such as stress, fatigue, medication, alcohol ingestion, etc., some of them with natural daily variations. However, the raw EEG data used was collected from a set of people at a particular measurement session, thus not reflecting daily variations or even variations along the required time spans (months or years). Nevertheless, in our study we concluded that EEGs collected in a row upon many similar visual stimuli are constant enough for implementing an authentication system based upon them, which is a good starting point.

Considering the fourth requirement (collectability), the current EEG measurement technology raises many problems. As EEG signals are very low-power, EEGs measurement must be done with special care to increase signal-to-noise ratios. This means that electrodes must have a good contact with the scalp, that all relevant induced electromagnetic signals (e.g. caused by power lines) must be filtered out and that interferences from other body activities (e.g. eye blinking) should be taken into consideration when collecting EEG signals or in a posteriori analysis. Finally, electrodes must be placed always in the same scalp location, an issue usually solved by using EEG helmets. We anticipate that the actual technological problems for EEG measurement may disappear in a near future, for instance, by using sensors under the scalp, thus we do not see it as a definitive barrier to the use of EEGs for biometric authentication. Nevertheless, in our study we made an effort to facilitate its deployment, both with the current technology or with future solutions. More specifically, we tried to get the best authentication results with the minimum possible set of electrodes (or EEG channels), all of them located in the occipital area of the brain, were the relevant EEG signals are to be measured.

For this study we did not obtain our own EEG samples from people. Instead, we used a public data set<sup>1</sup> containing EEG signals of 70 individuals, acquired with 64

<sup>1</sup> Hosted in <http://kdd.ics.uci.edu/>. It contains data for 122 subjects, but we could one find 70 subjects with more than 40 single object trials.

electrodes after their visual stimulation. After an initial period of evaluation, we learned that 8 channels were enough to achieve acceptable authentication results, namely reduced false positive and false negative ratios. These 8 channels are all located in the occipital area, where cognitive workload is more relevant.

For authenticating people using VEP features we used personal one-class classifiers (OCCs). These classifiers get as input the VEP features of the person being classified and produce a TRUE/FALSE output value. We used two different OCCs in order to study which one would produce better authentication results: K-Nearest Neighbor (KNN) and Support Vector Data Description (SVDD) [29]. After testing both classifiers, we also tested a two classification architectures combining both KNN and SVDD. These combined classifiers, that we nicknamed OR and AND, produce outputs after computing a logic function of the outputs of each individual classifier.

The results, obtained with personal OCCs of the four types, showed that VEPs can be used as a biometric data for authentication systems, producing results with high correctness, namely low false positive and false negative ratios. The results also showed that correctness is fairly stable for all evaluated subjects, an important requirement of biometric authentication systems.

## 2 Electroencephalograms

EEG signals are electric signals gathered in the scalp of an individual and result from the combination of signals from two different sources: (i) close-by cerebral activity (ii) and non-cerebral origins, such as eye motion, eye blinking and electrocardiac activity, called *artifacts*.

EEG signals generated by cerebral activity are usually decomposed in several frequency bands. Each band contains signals associated with particular brain activities [3]. The standard EEG frequency bands are: 0.5–3.5 Hz ( $\delta$ ), 4–7 Hz ( $\theta$ ), 8–15 Hz ( $\alpha$ ), 15–30 Hz ( $\beta$ ), 30–70 Hz or around 40 Hz ( $\gamma$ ). This last one,  $\gamma$  band, has been related both to gestalt perception [10] and to cognitive functions such as attention, learning, visual perception and memory.

For each particular brain activity there is one particular area that produces stronger electrical activity in one of the previously referred frequency bands; similarly, artifact manifestations are more relevant in some parts of the scalp than in others. Consequently, EEG signals are multi-channel signals, where each channel corresponds to a specific scalp electrode location. In this study we will consider only the occipital area of the scalp, which is known to provide stronger electrical signals in the  $\gamma$  band in response to visual stimulation and perception of pictures [30,28,7].

### 2.1 Visual Evoked Potentials (VEPs)

Visual evoked potentials (VEPs) are brain activity responses to visual stimuli, which may comprise different components, such as color, texture, motion, objects, readability (text vs. non-text), etc. Each of these components has impact in the spacial dispersion of the VEP through the scalp, being observed differently in each EEG channel and in different frequency bands. Therefore, for focusing the VEP production and analysis,



**Fig. 1.** Pictures of the Snodgrass & Vanderwart standard 260 picture set

the set of visual stimuli must be coherent, i.e., it should stimulate always the same brain areas.

Several research works (see Section 3) were previously conducted for achieving individual identification using VEPs produced upon the presentation of images from the Snodgrass and Vanderwart picture set [26]. This standard set of 260 pictures was conceived for experiments investigating differences and similarities in the cognitive processing of pictures. The pictures are black-and-white line drawings executed according to a set of rules that provide consistency of pictorial representation (see example in Fig. 1).

Various studies [12,6] showed that VEPs recorded from the human scalp contain a train of short latency wavelets in the  $\gamma$  band, precisely time locked to the stimulus and lasting approximately 100 ms. Furthermore, a more recent study showed that perception of pictures from the Snodgrass & Vanderwart picture set induced highly synchronized neural activity, in the  $\gamma$  band, between posterior electrodes [7].

### 3 Related Work

Poulos *et al.* [17,18] proposed a method to distinguish an individual from the rest using EEG signals. They performed a parametric spectral analysis of  $\alpha$  band EEG signals by fitting to them a linear all-pole autoregressive model. The coefficients of the fitted model were then used as features for the identification component. In [17] the identification component was built with computational geometric algorithms; in [18] they changed it to a neural network, namely for a Kohonen's Linear Vector Quantizer [11]. The cerebral activity was recorded from subjects at rest, with closed eyes, using only one channel and during three minutes.

Although the goal of Poulos *et al.* was person identification using his brain activity, in [18] they experimented classification of a person as one of a finite set of known persons. In the tests they recorded 45 EEG features from each of 4 individuals (the X set) and one EEG feature from each of 75 individuals (the non-X set). The neural network was trained using 20 features from each X member and 30 features from non-X members. Then the system was used to classify the remaining 25 features of each X member and the 45 features from the remaining non-X members. This process was repeated for all the 4 X members, attaining a correct classification score between 72% and 84%.

Using VEPs and signals in the  $\gamma$  band to perform subject identification was proposed by Palaniappan [13] and followed on his posterior studies [15,20,19,21,16]. In all these

works is used the same dataset of VEPs, recorded from 40 individuals and comprising a 61-channel EEG for 30 VEPs triggered by pictures from the Snodgrass & Vanderwart set.

These six subject identification studies are all similar; they mainly differ in filtering and classification components. First VEP signals are processed to remove artifacts and noise caused by other background brain activities not related with the VEP. Next they are filtered with a pass-band, digital filter in to isolate signals from the  $\gamma$  band. Then, for each of the 61 channels is computed its spectral power and normalized with the energy values from all the 61 channels; the 61 resulting values form a feature array. These features are then used to perform subject identification using a classifier with as many output categories as the number of individuals used to train it; in this case there were 40 individuals, thus the classifier has 40 different outputs. In the experiments, half of the features from each individual were used to train the classifier and the other half for testing the correctness of its output. The tested correctness of all these approaches is somewhat similar, ranging from 85.59% up to 99.62%.

For filtering, in [13][19][21] a Butterworth filter was used, while in [20][15][16] an elliptic FIR filter was used (in the latter the lower pass-band threshold was lower, 20 Hz). For classifying, in [20][15][21][16] was used an Elman back-propagation neural network [5], in [13] a back-propagation multi-layer perceptron, in [19][21] a simplified fuzzy ARTMAP [9] and in [19] a KNN.

Some attempts were made to reduce the number of channels used in these VEP-based approaches. In [16] Davies-Bouldin Indexes [4] were used to order the channels according to their relevance. Correct identification results using the most relevant DBI-oriented channels gave 13.63% with 1 channel, about 50% with 6 channels and 99.0% with 50 channels.

There are already several studies on authentication with EEG signals, but all them use different approaches [12][27][14]. Table 1 resumes some of their differences.

**Table 1.** Some differences among EEG-based authentication proposals

	[12] Marcel and de R. Milln, 2007	[26] Sun, 2008	[14] Palaniappan, 2008
EEG channels	$\alpha, \beta$	$\alpha, \beta, \gamma$	$\alpha, \beta, \gamma$
Electrodes	8	15	8
Feature array elements	96 (12 freqs./channel)	8 (CSP reduction)	11 (PCA reduction)
Tested subjects	9	9	5

In [12], authors collected EEGs from subjects performing 3 mental activities: imagination of movements with the right or left hand and imagination of words beginning with the random letter. Features' classification uses Gaussian Mixture Models and Maximum A Posteriori model adaptation. The correctness results were satisfactory but not very conclusive, because the number of evaluated subjects was too small (we used 70). A drawback of the classification approach is that it relies on a generic EEG model, which may not exist or requires training with EEGs from many people.

In [27], authors used 15 signals from the same dataset used in [12], raw feature reduction with common spatial patterns (CSP) and using multi-task learning to evaluate the advantage regarding single-task learning.

In [14], authors collected EEGs from subjects performing 5 imagined activities: nothing in particular (baseline activity), mathematical multiplication, geometric figure rotation, letter composition and visual counting. Feature arrays are initially composed by 18 channel spectral powers, 27 inter-hemispheric channel spectral power differences and 18 entropy values (yielding the non-linearity of channel signals). Features are then reduced to 11 elements using Principal Component Analysis (PCA). Features' classification uses a two-stage authentication process using maximum and minimum threshold values stored in personal profiles. Like in the previous article, correctness results were satisfactory but even less conclusive, due to the extremely small number of evaluated subjects (only 5).

All these three works used imagined activities to focus EEG-signals; we used VEPs instead. The advantages of VEPs is that they do not require any effort from the subjects being authenticated, as VEPs occur without any sort of human control. Furthermore, we did an evaluation with a larger population (roughly an order of magnitude more) than all these works, therefore our results yield a more trustworthy evaluation of the universality and uniqueness requirements. Finally, we did not use more electrodes than any of them, thus we do not require a more complex EEG acquisition setup.

Finally, some studies have been done with multi-modal biometrics involving EEG signals [24]. The main advantage of this approach regarding the simple EEG authentication is that one can reduce the number of electrodes (only 2 were used).

## 4 An EEG-Based Authentication System Using VEP

As previously stated, our goal was to build an authentication system based only in occipital VEP EEG signals gathered by a small number of electrodes. Note that, authentication is different from identification: an identification system gives the identity of the subject being evaluated, while an authentication system gives a yes/no answer whether or not the subject being evaluated is who he claims to be.

The VEP-based identification systems developed by Palaniappan *et al.* are also not directly usable as authentication systems. These systems were designed for identifying members of a set  $X$  of  $N$  subjects, having  $N$  possible output classifications. When these systems are used by other non- $X$  subjects, these will be identified as someone belonging to  $X$ . Thus, a non- $X$  person being authenticated only has to guess the erroneous identity the system gives to him, in order to get an authentication match.

Therefore, a new architecture is required to use EEG patterns for authenticating individuals. We propose a new one where we merge part of the contributions of the previously referred systems with some new ideas introduced by us.

### 4.1 Personal Classifiers

Our key design principle is to analyse EEG patterns in the  $\gamma$  band, namely VEPs in occipital area of the brain, with *one classifier per individual*, and not a classifier for all individuals. Furthermore, we used an OCC for each personal classifier, which is the correct type of classifier for an authentication scenario. Thus, when a subject claims to be  $X$ , we use  $X$ s OCC classifier to evaluate the correctness of the claim.




Fig. 2. Authentication components and inner activities of each component

OCCs may have many inputs to handle the features obtained from subjects, but always two possible output responses: TRUE or FALSE. Each personal OCC is trained only with inputs provided by its owner. When the individual being evaluated is the owner of the classifier, the output should be TRUE; otherwise, the output should be FALSE. Other outputs are errors, either false negatives or false positives, respectively (see table below).

Feature provider	OCC output	
	TRUE	FALSE
OCC owner	Correct result	Error (false negative)
other individual	Error (false positive)	Correct result

As previously referred, the goal for this new architecture was to use a reduced number of EEG channels. In the limit we would like to use only one channel, just like in the work of Poulos *et al.*. However, unlike the approach described in [16], we have not tried to detect the “best” channels (the ones with less correlation) from a set of measured features. Instead, we chose specific channel locations in the occipital area of the scalp and we ran authentication tests with them to find out the set of channels providing the highest authentication quality.

### 4.2 Authentication Process


Our authentication process is formed by three main components (see Fig. 2): (i) EEG signal acquisition, (ii) feature array extraction and (iii) feature array classification.

First VEP EEG signals are acquired from electrodes placed in the subject’s scalp. Then the feature array extractor processes raw EEG samples from  $C$  channels in order to extract a biometric measure of the subject: a feature array with  $C'$  energy values. Finally, this feature is processed by the OCC of the subject being authenticated, either to train the OCC or to get a TRUE or FALSE authentication outcome.


### 4.3 Description of the Data Set

As previously explained, we did not collect EEG signals for this study. Instead, we used a public data set registered for conducting other EEG studies, namely the genetic predisposition of people to alcoholism. Thus, it was not in any way specially gathered for authenticating people.


The data set is composed by EEG signals recorded from 70 individuals, both alcoholic and non-alcoholic, while exposed to short latency (300ms) visual stimuli. Each



**Fig. 3.** Location of the 64 electrodes used to collect the used data set and the 8 occipital electrodes used in our authentication system (grey)



**Fig. 4.** Stimuli visualization procedure



**Fig. 5.** Feature extraction: (a) VEP acquired by 8 occipital channels, (b)  $\gamma$ -band filtered VEP (30–50Hz), (c) feature signals formed by 8  $\gamma$ -band filtered VEPs and 28 differential VEPs and (d) normalized energy feature from original and differential VEP signals

individual completed a total number of 45 trials corresponding to the visualization of 45 pictures from the Snodgrass & Vanderwart picture set. EEG signals were acquired by 64 electrodes (61 actives + 3 reference), placed in the individuals' scalp, hardware filtered with a 0.1–50Hz passband and measured at a sampling rate of 256 samples per second. For building our authentication system we considered only 8 occipital channels from the 64 available in the dataset — channels PO3, PO4, POZ, PO7, PO8, O1, O2 and OZ (see Fig. 3).

Individuals were asked to recognise the pictures as soon as they were presented in a CRT screen, located 1 meter away from individuals' eyes. Each picture was presented only for 300 ms, separated by blank screen intervals of 5.1 seconds (see Fig. 4). After each picture presentation, only 1 second of EEG signal was recorded, corresponding to the VEP occurrence interval.



#### 4.4 Feature Array Extraction

VEP signals, which are raw EEG signals with 1 second measured after the presentation of the stimuli images, are the source data for the biometric authentication process of each individual. The feature extraction procedure from these signals is detailed below (cf. Fig. 5).

**Detection of Artifacts.** First, VEP signals containing artifacts are discarded. We considered only artifacts produced by eye blinking, which are the most common and intrusive ones. Detection of eye blinking artifacts is achieved with an amplitude threshold method: VEP signals with magnitude above  $50 \mu V$  are assumed to be contaminated [25] so they are discarded.

**EEG  $\gamma$ -band (30-50Hz) Frequency Filtering.** The resulting artifact-free VEP signals are filtered with a 30-50 Hz pass-band, using a  $10^{th}$  order Butterworth digital filter. The non-linearity of this filter was cancelled by using forward and reverse filtering. The resulting signal has zero phase distortion and an amplitude multiplied by the square of the amplitude response of the filter. After filtering, the 20 first and 20 last output samples are discarded, because they do not represent a properly filtered signal.

**Signal Composition.** For computing feature arrays we use  $C$  original EEG signals plus differential signals resulting from the subtraction of pairs of the  $C$  EEG signals. Thus, features include  $C' = C + \binom{C}{2}$  signals, which in our case, for  $C = 8$ , means that  $C' = 36$ .

By computing differential signals from the subtraction of pairs of EEG signals we expect to provide to classifiers information about the phase of the EEG signals and not just information about their amplitudes (energies). Phase shifts between subtracted sinusoidal signals with equal frequency and amplitude produce non-null signals with an energy that is a function of the phase shift. Therefore, we included the energy of differential signals in the features because it could denote phase shifts between EEG channels, thus more information about the subjects.

These differential signals are somewhat similar to the ones used in [14] but with two main differences: (i) we compute the energy of differential signals, while they compute differences between powers of different signals and (ii) we produce a differential signal from all pairs of signals, while they only compute differential powers between signal on different hemispheres. Thus, we are able to evaluate phase shifts on differential signals and we produce more information that may help to differentiate subjects.

**Energy Calculation and Normalization.** The energy of original and differential signals is computed with the Parseval's spectral power ratio theorem:

$$E(s) = \frac{1}{N} \sum_{n=1}^N s_n^2$$

where  $s_n$  is the  $n$ -th sample of signal  $s$  and  $N$  is the total number of samples in the signal. In our case  $N = 216$ , because we discard 40 samples of the 256 measured in 1 second of VEP after the  $\gamma$  filtering stage.

Finally, feature values are computed by normalizing the energy feature array. For this normalization we divide all array values by the maximum among them. This way, we get features with elements in the  $[0, 1]$  interval.

$$F [1 \dots C'] = \frac{E [1 \dots C']}{\max (E [1 \dots C'])}$$

#### 4.5 The Feature Classifier

The feature classifier is formed by independent, personal classifiers; so, for authenticating someone claiming to be X, we use the personal classifier of X, or the classifier *owned* by X. Each personal classifier is formed by an OCC, providing two different outputs (or classifications): TRUE and FALSE.

One-Class Classification is a type of classification where we deal with a two-class classification problem (target and outlier) but we only need to provide information to the classifier about the target. During an OCC train, the boundary between the target class and all other possible outlier classes is estimated from the target class data only. In our authentication goal, the target class is the classifier owner while the outlier class represents all other individuals.

In our study we used two types of OCCs in order to check which one would produce better authentication results: KNN with  $k=1$  and SVDD with a Radial Basis Function kernel [29]. We also tested two other OCC architectures, combining the outputs the KNN with SVDD. The goal of the combinations was to evaluate if there was any advantage in combining them in order to complement their individual correctness. The OR combination uses arithmetical averages, and the AND geometrical averages. For simplicity, we will call the first a *OR KNN-SVDD* and the second a *AND KNN-SVDD*.

We also found out that each classifier should be trained with single features from its owner, but should be used for authentication with average tests  $F$  features, obtained from the visualization of  $F$  images of the subject. A possible explanation of this fact is the following. Perception activities performed by individual's brain are not exactly the same for all visual stimuli, resulting in different VEP features. By training the classifier with as different as possible VEP features from its owner, we improve its ability to recognize them in the future, disregarding possible noise occurrences. On the other hand, by averaging VEP features during authentication processes, we reduce the probability of presenting to the classifier features from its owner too different from the ones it was trained with.

During the training of each classifier, we have to provide a *rejection fraction threshold* that will be used to establish acceptance or rejection ratios. Low rejection fraction values lead classifiers to produce low false negative and high false positive ratios, while high rejection fraction values, on the contrary, lead classifiers to produce high false negative and low false positive ratios. The choice of the best rejection fraction threshold implies a balance between security (low false positive ratio) and comfort for the individuals engaged in a correct authentication process (low false negative ratio).

## 5 Experimentation

The number of feature arrays used per subject was variable, both from start (in the data set) and furthermore after eye blink detection. Therefore, we decided to test classifiers

**Table 2.** Features evaluated per classifier

Composition of features	Features evaluated per classifier	
	From the owner	Total
Single features	15	$70 \times 15 = 1050$
Pairs of features	$\binom{15}{2} = 105$	$70 \times \binom{15}{2} = 7350$
Trios of features	$\binom{15}{3} = 455$	$70 \times \binom{15}{3} = 31850$

**Table 3.** Mean and standard deviation (inside parenthesis) of correctness results for the four OCC classifiers obtained in 10 independent classification tests. Columns labeled 1, 2 and 3 represent tests using singular features and average combinations of pairs and trios of features, respectively

	Rejection fraction	Owners Correctness(%)			Outliers acceptance (%)		
		1	2	3	1	2	3
KNN	0.2	78.7 (13.1)	90.6 ( 7.5)	95.1 ( 5.3)	5.2 (1.3)	5.6 (1.6)	6.4 (1.9)
	0.5	50.1 (15.9)	65.3 ( 9.3)	74.1 ( 8.6)	1.9 (0.7)	2.2 (0.9)	2.3 (1.3)
	0.7	31.3 (18.5)	46.1 (11.3)	66.3 ( 9.8)	0.9 (0.4)	1.1 (0.5)	1.3 (0.7)
SVDD	0.2	76.1 (12.8)	95.2 ( 4.9)	98.5 ( 3.5)	5.7 (1.8)	8.5 (2.1)	10.1 (3.2)
	0.5	58.5 (17.4)	88.3 ( 8.7)	93.7 ( 5.7)	2.8 (1.6)	4.4 (1.8)	5.1 (2.7)
	0.7	44.2 (20.5)	77.7 (12.3)	85.3 ( 9.8)	1.7 (1.6)	2.6 (1.7)	3.6 (1.9)
AND	0.2	83.3 (12.1)	96.4 ( 6.1)	99.0 ( 3.0)	4.5 (5.6)	6.0 (7.4)	6.8 (8.2)
	0.5	60.4 (16.8)	85.6 (12.4)	92.8 (11.3)	1.3 (2.2)	1.8 (2.9)	1.9 (3.0)
	0.7	37.8 (18.7)	68.4 (18.1)	79.4 (19.2)	0.4 (0.9)	0.6 (1.1)	0.6 (1.4)
	0.2	83.8 (11.0)	96.5 ( 6.0)	99.1 ( 3.5)	4.7 (5.7)	6.2 (7.6)	6.8 (8.4)
	0.5	59.7 (17.2)	85.7 (13.1)	92.8 (11.5)	1.2 (1.9)	1.7 (2.8)	2.0 (3.2)
	0.7	38.7 (18.2)	69.8 (18.1)	80.5 (18.1)	0.4 (0.8)	0.6 (1.2)	0.6 (1.3)


with fixed numbers of features and train classifiers with the maximum possible number of features until a given maximum. This is a conservative approach, since some classifiers may not have enough features to be properly trained. Nevertheless, we did not observe abnormal errors in such classifiers.

Thus, to train each personal classifier we used no more than 30 features of its owner. For testing each personal classifier, we used 15 features of its owner and 15 features from each of the other 69 subjects, which makes a total of 1050 test features. Note that each classifier had never “seen” the test feature before. The test features of each individual were used alone or averaged in pairs or trios. The number of features evaluated per classifier is presented in Table 5.

We run authentication tests with all the proposed four classifiers, in order to verify which one of them is more suitable for our authentication system. In the tests we tried also to assess the impact of two configuration parameters for the overall correctness of the authentication system: OCC rejection fraction threshold and classification of multiple, averaged features.

### 5.1 Overall Evaluation Results

The overall biometric authentication results of the 70 classifiers for three rejection fraction thresholds and different combinations of features are summarized in Table 5. The values presented are the mean and standard deviation obtained from 10 independent tests with the 70 OCCs, each one of them using different features from the owner (to




**Fig. 6.** Average individual classification results of the combined OR KNN-SVDD classifier, obtained after 10 independent tests. The upper (green) curve in each graphic shows the average correct owner classifications per classifier, while the lower (red) curve shows the average false positives per classifier. The vertical line under each average value shows the maximum and minimum values observed in the 10 tests.

train and test his classifier) and from outliers (to test). The graphics of Fig. 6 show the values of these 10 independent tests per personal classifier, but only for the combined OR KNN-SVDD classifier, using single and trios of features.

The results show that the rejection fraction threshold used while training classifiers had the expected impact on authentication results: for low rejection values the classifier provides a correct classification of its owner (low false negative ratio) but can be misled by many other individuals (high false positive ratio), while for higher rejection values the correct classification of owners decreases but the same happens to the wrong acceptance of other individuals.

Comparing KNN and SVDD, we can conclude that both have advantages and disadvantages: KNN gives lower outliers acceptance ratios (false positives), while SVDD gives higher owners' correctness ratios (true positives).

The combined OR and AND KNN-SVDD classifiers also have advantages and disadvantages when compared with the isolated OCCs. In general, they decrease the false positive ratio and most times (67%) they increase the owners' correctness ratio. However, they have a noticeable tendency to increase the standard deviation of the results,



**Fig. 7.** ROC of the 70 individual classifiers and their average AUC for each type of OCC considered and using feature trios

being thus less assertive than the isolated OCCs. Comparing the two combined KNN-SVDD classifiers, the results show that they are quite similar, but the OR combination is slightly better.

Finally, these results clearly show that the quality of the authentication increases when we use combinations of features instead of singular features. In absolute value, the owners' correctness gain is much higher than the loss in the false positive ratio.

## 5.2 Evaluation of Individual Classifiers

From the graphics of Fig. 6 we can conclude that average classification results are fairly stable for all the considered subjects. Therefore, with these tests we have reasons to believe that a biometric authentication system using EEGs may be suitable for a large majority of the population. Note that the evaluated subjects already include a group of people (alcoholics) that may have visual cognition problems and that was not noticeable in the authentication results.

The impact of false negatives and false positives can be reduced by making several independent authentication attempts in a row and accepting the subject as the legitimate classifier owner after a certain percentage of successful classifications. This percentage can be defined globally for all classifiers or specifically for each personal classifier.

A good indicator about an OCC performance is the plot of its receiver operating characteristic (ROC) curve. A ROC curve is calculated with several tests of the classifier with different rejection fraction thresholds applied to target objects and shows the

percentage of true positives in order to the percentage of the false positives during each test. Thus, ROC curves are useful to assert the effect of the rejection fraction threshold in tuning the correctness of the OCC.

The OCC with the best performance is the one that simultaneously maximizes true positive ratios and minimizes false negative ratios. This performance can be measure by calculating the area under curve (AUC). This way, the OCC with the higher AUC is assumed to be the OCC with best performance.

Figure 7 shows the ROCs of the 70 individual classifiers and their average AUC values for each of the four OCC types considered; these ROC curves where obtained with feature trios. These results clearly show that for evaluating trios of features the best OCC is the combined OR KNN-SVDD, while the worse is KNN.

## 6 Open Issue: Constancy

As referred in the introduction, a fundamental requirement of biometric authentication systems is constancy. If raw biometric data does not show enough constancy, then authentication may fail unexpectedly.

In our work we have not evaluated constancy because the data set we used was collected in a single session per subject. Therefore, at this point we can only extrapolate constancy expectations from other studies using VEPs.

The evaluation of constancy within EEG data is a complex task, mainly because high-quality EEG signal acquisition is still a non-trivial operation. For getting trustworthy data for evaluating constancy within EEGs, electrodes must be placed always in the same scalp location, electrodes must have a good contact with the scalp, all relevant induced electromagnetic signals must be filtered out and all relevant interferences from other body activities must be taken into consideration. Not surprisingly, all the works regarding identification and authentication with EEGs signals referred in Section 3 did not evaluate constancy and many of them used exactly the same data set we did.

Nevertheless, there are some publications on the constancy of VEPs, although using visual stimuli other than the images from the Snodgrass and Vanderwart set. Regan & Neima [22] reported an amplitude attenuation on VEPs caused by visual fatigue (loss of visual sensitivity). Rhodes *et al* [23] studied the effect of alcohol on VEP, both with and without a visual discrimination tasks. Results showed that a mean blood alcohol content of 90 mg% affected VEPs recorded from the central scalp by attenuating the overall amplitude of the later VEP components (60-200 msec) and by significantly reducing hemispheric asymmetry in the amplitude of these VEP components.

Both this studies evaluated amplitude changes in VEPs, and not energy. Therefore, the results cannot be easily extrapolated to our authentication system. Nevertheless, we assume that it is natural that VEPs could be influenced by many factors, such as diseases, stress, fatigue (either physical, visual or mental), or by ingestion of several substances (alcohol, stimulants, drugs, etc.).

In conclusion, we imagine that two different scenarios are possible, regarding the constancy of the VEPs we used:

- Non-constancy: VEPs change substantially during short time periods, such as daily or hourly, without any specific cause. In this case, our authentication system could not simply be used.

- Quasi-constancy: VEPs are fairly stable for a subject at the occasions when the VEPs are recorded, but can be influenced by several factors, such as stress, fatigue and ingestion of alcohol, stimulants or drugs. In this case, our system is able to detect deviations from the normal, which may be an interesting area to explore for authenticating people performing highly demanding jobs (air traffic controllers, truck drivers, medical doctors, etc.).

## 7 Conclusions

We presented in this paper a novel method for authenticating individuals using their brain activity. The EEG signals used were VEPs, i.e., brain responses associated to visual stimuli. In the described system we used EEG signals acquired with only 8 electrodes placed on occipital area of the brain, which is associated to visual and cognitive perception.

The authentication system presents several improvements over other previous works in the area of subject identification using VEPs:

- We used a reduced number of electrodes (8) and we placed them in the scalp area where EEG signals have more correlation with the stimuli.
- We used the differences between pairs of the 8 EEG signals to create other signals (differential signals) that provide extra information to classifiers. Namely, the energy of differential signals may reflect phase shifts between the original EEG signals. As far as we know, this technique was never used to increase the information provided by VEPs.
- Feature arrays with the energy of original and differential EEG signals are classified using personal classifiers. Personal classifiers are not general classifiers that use personal biometric templates, but classifiers trained to recognise their owner and to reject all other subjects.
- We used OCC personal classifiers instead of other types of classifiers (e.g. neural networks) or generic classifiers. OCCs are more appropriate to authentication because they only have to be trained with the owner biometric data. By using personal classifiers instead of generic ones we avoid the definition of EEG template models.
- OCCs are trained with single owner features but provide better results when tested with average of features instead of single ones.

Regarding other systems performing EEG-based authentication, we used VEPs, which are effortless for subjects, while others used more complex and annoying brain stimuli, such as activity imagining, we obtained satisfactory results with a population one order of magnitude larger than the other proposals (70 vs. 5 or 9 subjects), and we did not use more electrodes (only 8). Therefore, our system has clear advantages regarding the collectability requirements.

Average results obtained with authentication tests with 70 individuals, using a public VEP data set, showed that authentication with EEGs is possible and may be used in future applications. The ratios of owner's correctness and false positives are fairly stable for the tested population, which is a positive indication for the universality and uniqueness of the process.

A fundamental requirement of biometric authentication was not evaluated in this document: constancy. In fact, we assumed that the public data set, per subject, was collected in a very short time, thus it is not possible to take any conclusions about the constancy of VEPs in other scenarios. Since several factors may affect VEPs significantly, such as stress and fatigue, the constancy of VEP-based biometric authentication must be addressed by future research.

Since we are using cognitive brain activity as raw material for conducting authentication, authentication failures may reveal substantial deviations from the normal. Therefore, the lack of constancy would not necessarily be a problem for our authentication system. If we could experience such authentication errors in some specific scenarios, such as subjects under the influence of drugs or alcohol, or with mental or physical fatigue, this system could be used to detect dangerous cognitive deviations in subjects performing highly demanding jobs (air traffic controllers, truck drivers, medical doctors, etc.).

**Acknowledgements.** We would like to acknowledge H. Begleiter, from the Neurodynamics Laboratory at the State University of New York Health Center at Brooklyn, for the EEG data set, and D. Tax, for the OCC Matlab toolboxes.

## References

1. Başar, E.: Relation between EEG and Brain Evoked Potentials. In: EEG-Brain Dynamics. Elsevier, North-Holland biomedical Press, Amsterdam (1980)
2. Başar, E., Rosen, B., Başar-Eroglu, C., Greitschus, F.: The associations between 40 Hz-EEG the Middle Latency Response of the Auditory and Evoked Potential. *The Int. J. of Neuroscience* 33(1-2), 103–117 (1987)
3. Basar, E., Basar-Eroglu, C., Demiralp, T., Schrmann, M.: Time and Frequency Analysis of the Brain's Distributed Gamma-Band System. *IEEE Eng. in Medicine and Biology* 14, 400–410 (1995)
4. Davies, D., Bouldin, D.: A Cluster Separation Measure. *IEEE Trans. on Pattern Analysis and Mach. Intelligence* 1(2), 224–227 (1979)
5. Elman, J.L.: Finding structure in time. *Cognitive Science* 14(2), 179–211 (1990)
6. Galambos, R.: A comparison of certain gamma band (40-Hz) brain rhythms in cat and man. In: *Induced Rhythms in the Brain*, pp. 201–216. Birkhuser, Boston (1992)
7. Gruber, T., Mller, M.M., Keil, A.: Modulation of Induced Gamma Band Responses in a Perceptual Learning Task in the Human EEG. *J. of Cognitive Neuroscience* 14(5), 732–744 (2002)
8. Jain, A., Hong, L., Pankanti, S.: Biometric Identification. *Communications of the ACM* 43(2), 90–98 (2000)
9. Kasuba, T.: Simplified Fuzzy ARTMAP, pp. 18–25. *AI Expert, USA* (1993)
10. Keil, A., Mller, M.M., Ray, W.J., Gruber, T., Elbert, T.: Human Gamma Band Activity and Perception of a Gestalt. *The J. of Neuroscience* 19(16), 7152–7161 (1999)
11. Kohonen, T.: *Self-organization and associative memory*, 3rd edn. Springer, New York (1989)
12. Marcel, S., del Millán, J.R.: Person Authentication Using Brainwaves (EEG) and Maximum A Posteriori Model Adaptation. *IEEE Trans. on Pattern Analysis and Mach. Intelligence* 29(4), 743–748 (2007)
13. Palaniappan, R.: Method of identifying individuals using VEP signals and neural network. *IEE Proc. - Science Measurement and Technology* 151(1), 16–20 (2004)



14. Palaniappan, R.: Two-stage biometric authentication method using thought activity brain waves. *Int. J. of Neural Systems* 18(1) (2008)
15. Palaniappan, R., Mandic, D.P.: Energy of Brain Potentials Evoked During Visual Stimulus: A New Biometric? In: Duch, W., Kacprzyk, J., Oja, E., Zadrożny, S. (eds.) *ICANN 2005*. LNCS, vol. 3697, pp. 735–740. Springer, Heidelberg (2005)
16. Palaniappan, R., Mandic, D.P.: EEG Based Biometric Framework for Automatic Identity Verification. *J. of VLSI Signal Processing* 49, 243–250 (2007)
17. Poulos, M., Rangousi, M., Kafetzopoulos, E.: Person identification via the EEG using computational geometry algorithms. In: *Proc. of the 9th European Signal Processing (EUSIPCO 1998)*, Rhodes, Greece, pp. 2125–2128 (1998)
18. Poulos, M., Rangoussi, M., Chrissikopoulos, V., Evangelou, A.: Person Identification Based on Parametric Processing of the EEG. In: *Proc. of the 6th IEEE Int. Conf. on Electronics, Circuits and Systems (ICECS)*, pp. 283–286 (1999)
19. Ravi, K.V.R., Palaniappan, R.: Leave-one-out Authentication of Persons Using 40 Hz EEG Oscillations. In: *Proc. of the Int. Conf. on "Computer as a tool" (EUROCON 2005)*, Belgrade, Serbia & Montenegro (November 2005)
20. Ravi, K.V.R., Palaniappan, R.: Recognising Individuals Using Their Brain Patterns. In: *Proc. of the 3th Int. Conf. on Information Tech. and Applications (ICITA 2005)*. IEEE Computer Society, Los Alamitos (2005)
21. Ravi, K.V.R., Palaniappan, R.: Neural network classification of late gamma band electroencephalogram features. *Soft Computing* 10, 163–169 (2006)
22. Regan, D., Neima, D.: Visual fatigue and visual evoked potentials in multiple sclerosis, glaucoma, ocular hypertension and Parkinson's disease. *Journal of Neurology, Neurosurgery, and Psychiatry* 47(7) (July 1984)
23. Rhodes, L.E., Obitz, F.W., Creel, D.: Effect of alcohol and task on hemispheric asymmetry of visually evoked potentials in man. *Electroencephalography Clinical Neurophysiology* 38(6), 561–568 (1975)
24. Riera, A., Soria-Frisch, A., Caparrini, M., Grau, C., Ruffini, G.: Unobtrusive Biometric System Based on Electroencephalogram Analysis. In: *EURASIP J. on Advances in Signal Processing 2008* (2008)
25. Sivakumar, R., Ravindran, G.: Identification of Intermediate Latencies in Transient Visual Evoked Potentials. *Academic Open Internet J.* 17 (2006)
26. Snodgrass, J.G., Vanderwart, M.: A Standardized Set of 260 Pictures: Norms for Name Agreement, Image Agreement, Familiarity and Visual Complexity. *J. of Experimental Psychology: Human Learning and Memory* 6(2), 174–215 (1980)
27. Sun, S.: Multitask learning for eeg-based biometrics. In: *19th Int. Conf. on Pattern Recognition (ICPR 2008)*, Tampa, Florida, USA (December 2008)
28. Tallon-Baudry, C., Bertrand, O., Peronnet, F., Pernier, J.: Induced  $\gamma$ -Band Activity during the Delay of a Visual Short-Term Memory Task in Humans. *The J. of Neuroscience* 18(11), 4244–4254 (1998)
29. Tax, D.M.J.: One-class classification; Concept-learning in the absence of counter-examples. Ph.D. thesis, Delft University of Technology, Delft, Netherlands (June 2001)
30. Lutzenberger, W., Pulvermiller, F., Elbert, T., Birbaumer, N.: Visual stimulation alters local 40-Hz responses in humans: an EEG-study. *Neuroscience Letters* 183(1-2), 39–42 (1995)

**Part IV**  
**Bioinformatics**

# Reconstructing the Phylogenetic Relationships of the Cyclophyllidean Cestodes: A Case Study Using ITS2 rDNA and Sequence-Structure Alignment

Veena Tandon<sup>1,2</sup>, Devendra Kumar Biswal<sup>2</sup>, Pramod Kumar Prasad<sup>1</sup>,  
and Chenkual Malsawmtluangi<sup>1</sup>

<sup>1</sup> Department of Zoology, North-Eastern Hill University, Shillong 793022, Meghalaya, India  
{tandonveena, pramod812, mschenkual}@gmail.com

<sup>2</sup> Bioinformatics Centre, North-Eastern Hill University, Shillong 793022, Meghalaya, India  
{tandonveena, devbioinfo}@gmail.com

**Abstract.** Different authors have subscribed to varied opinions on the phylogenetic relationships and systematic position of cyclophyllidean cestodes owing to high degree of non-significant morphological variations among these taxa. Molecular phylogenetic analysis based on ITS2 rDNA of 16 representatives spanning 6 different families (Mesocestoididae, Davaineidae, Anoplocephalidae, Taeniidae, Dipylidiidae and Hymenolepididae) of the Order Cyclophyllidea and one out group from the family Diphylobothriidae of the Order Pseudophyllidea revealed the order Order Cyclophyllidea to be a monophyletic clade. The results were further validated by bayesian analysis, primary sequence-structure alignment and subsequent molecular morphometrics analysis. A consistent support for all the major nodes was seen across all the phylogenetic trees. Interestingly, the genus *Mesocestoides*, in which the genital pores are ventral in position, a feature of the Order Pseudophyllidea, was accommodated within Cyclophyllidea and served as a sister clade close to the families Taeniidae, Anoplocephalidae, Hymenolepididae and Dipylidiidae.

**Keywords:** *Mesocestoides*, Cyclophyllidea cestodes, ITS2 rDNA, Molecular morphometrics, Sequence-structure alignment.

## 1 Introduction

Cyclophyllidean cestodes are parasites occurring as adult or larval stage in a wide variety of avian and mammalian hosts and are cosmopolitan in nature commonly called tapeworms that live in the digestive tract of vertebrates as adults and often in the bodies of various animals as juveniles. Taxonomists face inconsistent or even contradictory information when they examine the systematic relationships between cestodes at higher taxonomic groupings [1]. The phylogeny at higher levels is of little significance as the taxonomy is primarily based on morphological characters and in case of cestodes it is often difficult to discern between secondary loss and convergence of morphological characters as several authors have weighed the characters of taxa differently [2]. This applies well to the taxonomic positions of *Mesocestoides*,

which is quite complicated by a high degree of non-significant morphological variation. Cyclophyllidean cestodes of the Family Mesocestoididae differ from other taxa in the Order Cyclophyllidea in some important characteristics. The life cycle of *Mesocestoides* spp. requires three hosts and not two; the median ventral position of the genital atrium and the presence of bipartite vitelline gland in *Mesocestoides* spp. appear to be unique among the Cyclophyllidea [3].

Although, nowadays, most authors agree that there are 15 families included in the monophyletic Order Cyclophyllidea, till date no study has pinpointed the complete taxonomic linkage of all the 15 families persuasively with genetic evidence. Besides, the taxonomic position of Mesocestoididae remains unclear as some of the cestodes at the generic level of the family bear similarity to those of the Order Pseudophyllidea; one such is the case of *M. lineatus*, with a wide range of hosts thus lacking host specificity, which is atypical of Cyclophyllidea but resembles the cestodes (viz. *Diphyllobothrium dendriticum*) of the Order Pseudophyllidea [4].

The phylogeny of cyclophyllidean cestodes has been reported with aid of 12S rRNA markers of mitochondrial genomic regions from 21 cestode species spanning eight families [5]. In the present study, we address the overall taxonomic resolution of cyclophyllidean cestodes with a different phylogenetic marker using a combinatorial approach of sequence analysis and molecular morphometrics. The internal transcribed spacer 2 (ITS2), the region of ribosomal RNA between 5.8S rRNA gene and the large subunit (28S rRNA) has proven to be appropriate marker for analysis of microscale phylogenies of close relatives [6]. Moreover, the ITS2 sequence data can be subjected to secondary structure predictions and as the secondary structure seems to be well conserved, it can provide clues for higher taxonomic studies [7]. This is quite obvious that phylogenetic analyses are improvable by inclusion of molecular morphometrics information in common sequence analysis [8]. Here, we combine sequence with structural information and apart from the biological problem, address the different in-silico practices in vogue for phylogeny studies using ITS2 r-DNA.

## 2 Materials and Methods

### 2.1 Sequence Alignments and Molecular Phylogenetic Analysis

The Cyclophyllidean Cestoda sequences from several geographical locations spanning six different families and one from the Pseudophyllidea order (taken as out group) was retrieved from the NCBI GenBank databases in the present study. Nucleotide sequences were aligned and edited using ClustalW [9]. A phylogenetic tree was constructed using the Neighbor-Joining as well as Maximum Parsimony methods in MEGA 4.0 [10]. Branch support was given using 1000 bootstrap replicates. Maximum parsimony was accomplished with gaps treated as missing data and all characters coded as “unordered” and equally weighted.

### 2.2 Bayesian Phylogenetic Analysis

A Bayesian analysis using MrBayes V 3.12 [11] was carried out for tree construction using a general time reversible substitution model (GTR) with substitution rates estimated by MrBayes. Metropolis-Coupled Markov Chain Monte Carlo (MCMCMC)

sampling was performed with two incrementally heated chains that were combinatorially run for 20,000 generations. The convergence of MCMCMC was then monitored by examining the value of the marginal likelihood through generations. Coalescence of substitution rate and rate model parameters were also examined. Average standard deviation of split frequencies was checked and the generations were kept on adding until the standard deviation value was below 0.01. The values slightly differed because of stochastic effects. The sample of substitution model parameters and samples of trees and branch lengths were summarized by the “sump burnin” and “sumt burnin” commands, respectively. The values in the following commands were adjusted as per the 25% of our samples. A cladogram with the posterior probabilities for each split and a phylogram with mean branch lengths were generated and subsequently read by the tree drawing program Tree view V1.6.6 [12].

### 2.3 ITS2 Secondary Structure Prediction and Sequence Structure Alignment Footnotes

ITS2 secondary structures of the cestodes were folded with the help of MFold [13] by screening for thermodynamically optimal and suboptimal secondary structures (default settings, with  $T=25^{\circ}\text{C}$ ). The secondary structures in Vienna (dot-bracket-dot) format was used as an input for MARNAs [14] to calculate sequence-structure multiple alignment. However, there was a limitation with the online server that the maximum length of one RNA sequence is restricted to 500 bases; hence some of the ITS sequences whose exact boundary information was available from GenBank graphics view was trimmed for facilitating MARNAs to run. Some more cestode sequences were also taken whose 5.8S, 28S and ITS2 regions were clearly defined so as to include in the sequence-structure multiple alignment dataset. A phylogenetic tree was created using ProfDistS [15] that takes the multiple aligned sequence-structure as an input and a consensus tree was built using RNA/DNA structure profile neighborhood-joining method with 100 bootstraps.

Besides, the GC content of the ITS 2 regions was calculated using Oligo Calculator available at <http://www.pitt.edu/~rsup/OligoCalc.html>.

## 3 Results

### 3.1 Neighbour-Joining (NJ) and Maximum Parsimony (MP) Trees

GenBank accession numbers of ITS2 sequences for the cestodes spanning 6 families of the Order Cyclophyllidea and one from the Order Pseudophyllidea (as out group) are given in Table 1. The evolutionary history was inferred using the NJ method [16] and the bootstrap consensus tree (Fig. 1), inferred from 1000 replicates, depicted an overall robust topology of the cyclophyllidean cestodes' phylogeny. Branches corresponding to partitions reproduced in less than 50% bootstrap replicates are collapsed. The evolutionary distances were computed using the Maximum Composite Likelihood method [17] and are in the units of the number of base substitutions per site. All positions containing gaps and missing data were eliminated from the dataset (complete deletion option). There was a total of 168 positions in the final dataset.

**Table 1.** Percentage of GC content in the ITS2 region of various cyclophyllidean cestodes

Species	GenBank Accession No.	Order & Family
<i>Hymenolepis diminuta</i> India_Meghalaya	FJ939132.1	Cyclophyllidea: Hymenolepididae
<i>Hymenolepis diminuta</i> Japan	<u>AB494475.1</u>	Cyclophyllidea: Hymenolepididae
<i>Hymenolepis microstoma</i> Japan	<u>AB494478.1</u>	Cyclophyllidea: Hymenolepididae
<i>Hymenolepis nana</i> Japan	AB494477.1	Cyclophyllidea: Hymenolepididae
<i>Dipylidium caninum</i> China	AM491339.1	Cyclophyllidea: Dipylidiidae
<i>Taenia saginata</i> China	AY825540.1	Cyclophyllidea: Taeniidae
<i>Echinococcus granulosus</i> China	AY389985.1	Cyclophyllidea: Taeniidae
<i>Taenia serialis</i> Canada	DQ099571.1	Cyclophyllidea: Taeniidae
<i>Taenia crassiceps</i> California	DQ099564.1	Cyclophyllidea: Taeniidae
<i>Taenia taeniaeformis</i> India_Andhra Pradesh	EU051352.1	Cyclophyllidea: Taeniidae
<i>Taenia taeniaeformis</i> India_Meghalaya	FJ939133.1	Cyclophyllidea: Taeniidae
<i>Moniezia expansa</i> Japan	AB367793.1	Cyclophyllidea: Anoplocephalidae
<i>Anoplocephala perfoliata</i> Germany	AJ578153.1	Cyclophyllidea: Anoplocephalidae
<i>Mesocestoides</i> spp. USA	AF119697.1	Cyclophyllidea: Mesocestoididae
<i>Raillietina beveridgei</i> Australia	AY382318.1	Cyclophyllidea: Davaineidae
<i>Raillietina dromaius</i> Australia	AY382320.1	Cyclophyllidea: Davaineidae
* <i>Diphyllobothrium latum</i> Japan	AB302387.1	Pseudophyllidea: Diphyllo- bothriidae

With MP method [18] the most parsimonious tree drawn had the length=727. The percentage of replicate trees in which the associated taxa clustered together in the bootstrap test (1000 replicates) are shown next to the branches (Fig. 2) and the MP tree was obtained using the Close-Neighbor-Interchange algorithm [19] with search level 3 in which initial trees were obtained with the random addition of sequences. After deletion of the positions containing the gaps, there were a total of 168 positions in the final dataset out of which 139 were parsimony informative.

The phylogenetic analysis using the distance and character state methods showed very good bootstrap values (Figs. 1 & 2) and all the six cyclophyllidean families depicted reliable monophyletic groupings. Bootstrap values for the six monophyletic groups ranged from 70-100%.

### 3.2 Bayesian Phylogeny

Bayesian analysis of the alignment retained the same topology and supported the branches with good bootstrap values (Fig. 3), though there were slight variations in the placing of some species of the Family Taeniidae (*Taenia saginata* from China and *Taenia taeniaeformis* from India) that were grouped in another node from the rest of

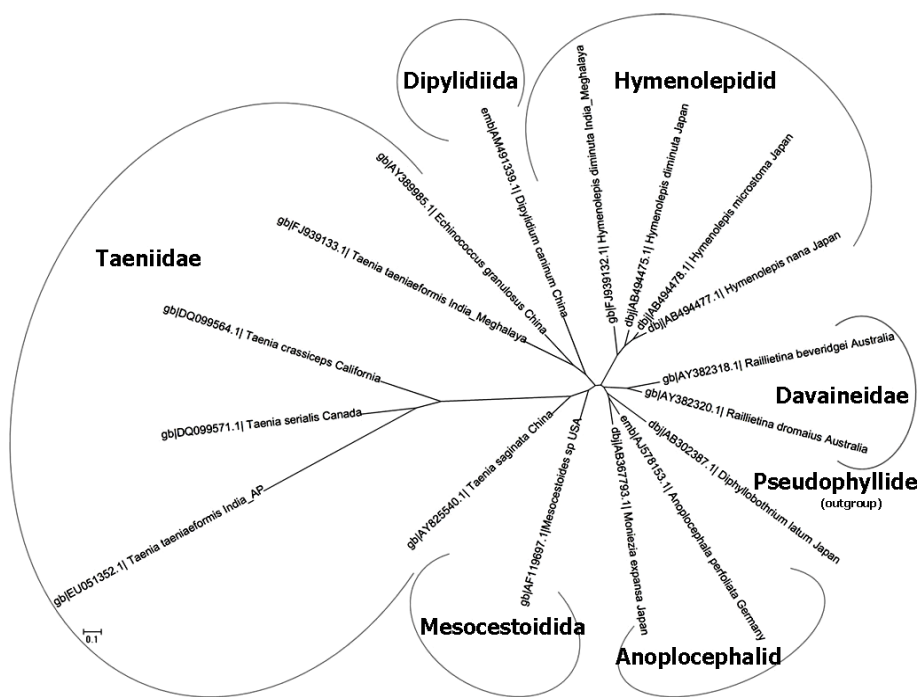


Fig. 1. Neighbor-Joining Tree




Fig. 2. Maximum Parsimony Tree

the *Taenia* species. *Diphyllobothrium latum* of the Order Pseudophyllidea was rooted as an out group.




Fig. 3. Bayesian Analysis Phylogeny

### 3.3 Secondary Structure Analysis and GC Content


#### ITS2 Secondary Structures

The ITS 2 secondary structures (Figs. 4, 5A-J) were analyzed for conserved stem and loop. The *Hymenolepis* species showed characteristic hallmark of ITS 2 secondary structure, i.e., four helices were clearly visible in secondary structures with third one as the longest. However, the third helix contains a side branch (Fig. 4). Species of *Mesocestoides* and *Raillietina* also maintained common secondary ITS2 core structure. *Taenia* and *Echinococcus* species showed a lot of variation in the secondary structure with many extra helices, loops and side branches. UGGU motif (Fig. 4) in the secondary structure was present in almost all the species of *Hymenolepis* genus and the U-U mismatch motif was completely absent in them. Secondary structures of species, belonging to the same genus showed high overall structural similarity except *Taenia* species in which considerable differences were noticed. The grouping of the families Taeniidae, Mesocestoididae and Hymenolepididae together in phylogenetic trees (Figs. 1, 2 & 3) forming a monophyletic group was supported by ITS2 secondary structure similarity.

#### GC Content

The GC content in the ITS2 region was calculated (Table. 2) and it was found that for *Taenia* species the GC content varied from 54.6% to 62.6%. For species of *Hymenolepis*, it ranged from 43.7% to 54%. *Taenia* species showed a higher GC content compared to others. The GC content also somewhat reflected grouping pattern of the organisms in the phylogenetic tree. Among many other factors, GC content is one of the factors related with stability of the secondary structure.







**Fig. 4.** ITS2 Secondary Structure showing UGGU motif: *Hymenolepis nana*

**Table 2.** Percentage of GC content in the ITS2 region of various cyclophyllidean cestodes


<b>Organism name</b>	<b>GC Content</b>
<i>Taenia saginata</i>	62.4%
<i>Taenia serialis</i>	62.6%
<i>Taenia crassiceps</i>	57%
<i>Taenia taeniaeformis</i>	61.6%
<i>Taenia pisiformis</i>	54.6%
<i>Hymenolepis nana</i>	50.5%
<i>Hymenolepis diminuta</i>	43.7%
<i>Raillietina Beveridgei</i>	50.5%
<i>Raillietina australis</i>	49.4%
<i>Mesocestoides spp.</i>	58.8%
<i>Anoplocephala perfoliata</i>	46%
<i>Echinococcus granulosus</i>	59%




5A



5B




5C




5D

**Fig. 5. A-J. ITS2 Secondary Structures of Cyclophyllidean cestodes**




**dG = -235.57 [initially -245.20] *E. granulosis\_China***

5E



**dG = -190.71 [initially -219.20] *T. taeniformis\_India\_Andhra Pradesh***

5F



**dG = -279.49 [initially -301.30] *T. taeniformis\_India\_Meghalaya***

5G



**dG = -253.01 [initially -265.30] *A. perfoliata\_Germany***

5H

Fig. 5. (Continued)




Fig. 5. (Continued)

### Primary Sequence-secondary Structure Alignment

Apart from the ITS 2 sequences initially used for the primary sequence analysis, some more sequences were included for secondary structure alignment with the primary sequence. Multiple sequence-structure alignment from MARNAs was used in ProfDistS program to build phylogenetic tree (Fig. 6); though secondary structure information helped to improve alignment, the proposed phylogeny showed slight differences. However, the monophyletic character of different groups was retained with few exceptions. The *Echinococcus* species were placed close to species of *Hymenolepis*. *Taenia saginata* and *Taenia crassiceps* were placed in the upper branch of the phylogenetic tree away from the basal group of *Taenia* species. These slight differences in tree topology may be due to specific ITS2 rate matrix used in the analysis. As most of the studies related to ITS2 have been carried out pertaining to plants and fungi, the specific rate matrix developed largely depends on those data. The ITS2 region of cestodes may follow different rates of evolution and thus ITS2 rate matrix specific to cestode may provide better results. Overall there was considerable similarity between the ITS analysis and the consensus of previous phylogenetic reconstruction using other DNA loci.

## 4 Discussion

The ITS2 region is a requisite in ribosome biogenesis [20] and its gradual removal from mature rRNA is driven by its specific secondary structure. Using the secondary structure of ITS2 sequences from various cyclophyllidean cestodes covering six important families in this study, we aimed to pursue three consecutive queries concerning their systematic relationships (i) whether the cyclophyllidean cestodes retain the

same taxonomic positions at higher levels and can be regarded as monophyletic considering ITS2 as molecular markers for drawing the phylogeny; (ii) how does the secondary structure of ITS2 sequences contribute to the Cyclophyllidea phylogeny using molecular morphometrics approach, and whether the latter would corroborate the monophyletic characteristics at the family level?




Fig. 6. Multiple sequence-structure alignment

The order Cyclophyllidea comprises 15 families. Of these 6 families viz. Hymenolepididae, Dipylidiidae, Taeniidae, Anoplocephalidae, Mesocestoididae and Davaineidae were considered for analysis in the present study. Using the ITS2 sequence data of *Diphyllobothrium latum* that represents the Order Pseudophyllidea as an out group, we constructed phylogenetic trees using distance-based, character-based and Bayesian methods. Besides, molecular morphometrics approach was employed taking sequence-structure alignment into consideration. Our study shows that all the taxa were clearly monophyletic within their families and principally correspond to earlier classifications based upon morphology and biology [2].

The genus *Mesocestoides* (Family: Mesocestoididae) has a complicated taxonomy owing to its high degree of nonsignificant morphological variations and uncertain relationship with Cyclophyllidea [2]. The two known genera (*Mesocestoides* and *Mesogyna*) differ from the other taxa in Cyclophyllidea in several important features. The midventral (and not lateral) position of the genital atrium, the vitelline gland consisting of two compact masses, and the requirement of three hosts in the life cycle of *Mesocestoides* are the unique features that make the group stand apart from other Cyclophyllidea. The families Mesocestoididae and Taeniidae share a common origin as tentatively suggested on the basis of tegumental hairs of their metacestodes; in contrast the cysticercoids of other cyclophyllideans have series of fibrous layers instead [21]. Moreover, the entire lifecycle of *Mesocestoides* is quite aberrant and the number of intermediate hosts remains enigmatic. Our dendrograms, predicted through

several in-silico approaches, demonstrate that *Mesocestoides* spp from USA are closer to the families Anoplocephalidae and Taeniidae. Due to the lack of ITS 2 sequences of other Mesocestoidae taxa in the public domain, we could not build a better dataset for accurate resolution of the family with high precision. Nevertheless our data supported Khalil et al's (1994) arrangement of Mesocestoides into Cyclophyllidea and a narrow relationship between Mesocestoididae, Taeniidae, Hymenolepididae and Anoplocephalidae.

Taeniids are the best-known cestodes. The various phylogenetic methods applied to *Taenia* and *Echinococcus* corroborates the monophyletic grouping of the family Taeniidae. The present analysis also agrees to the monophyly of other families under Cyclophyllidea; further analysis can be done once more and more molecular markers are deposited in public gene bank databases.

## 5 Conclusions

A molecular morphometrics approach that uses combined features both from anatomical and quantitative morphometrics and molecular primary sequence comparison was the basis of our study. The approach differentiates significant features between anatomical and molecular characters that make molecular morphometrics a strong predictive tool for phylogenetic resolution. There is always more than one gene involved in anatomical variations and most importantly the genetic sites responsible for morphological characters are usually not known. On the contrary, molecular structural variations are attributable to identifiable mutations that can be characterized at the single mutational level. The observed anatomical characters are the outcome of both the genetic characters as well as epigenetic effects (environmental influences), whereas the molecular morphometrics method takes advantage of the fact that molecular characters are independent of their somatic expression (Smith, 1992).

The analysis corroborated strong results for phylogenetic relationships of cyclophyllidean cestodes. The use of ITS2 data as phylogenetic molecular markers and the inclusion of secondary structure information in this analysis offered a higher resolution power for relationships from the level of sub species up to the order level.

**Acknowledgements.** This study was carried out under the DIT (Ministry of Communication & Information Technology (Government of India) - sponsored project, "Northeast Parasite Information and Analysis Centre" sanctioned to VT & DKB. Partial support from UPE-Biosciences Programme of UGC in the School of Life Sciences at NEHU is gratefully acknowledged.

## References

1. Mariaux, J.: Cestode systematics: Any progress? *Int. J. Parasitol.* 26(3), 231–243 (1996)
2. Rausch, R.L.: Family Mesocestoididae Fuhrmann, 1907. In: Khalil, L.F., Jones, A., Bray, R.A. (eds.) *Keys to the Cestode Parasites of Vertebrates*, pp. 309–314. CAB International, Wallingford (1994)

3. Chandler, A.C.: Observations on the anatomy of Mesocestoides. *J. Parasitol.* 32, 242–246 (1946)
4. Kamegai, S., Ichihara, A., Nonobe, H., Machida, M.: The 6th and 7th records of human infection with *Mesocestoides lineatus* (Cestoda) in Japan. *Research Bulletin of the Meguro Parasitological Museum* 1, 1–7 (1967)
5. von Nickisch-Roseneck, M., Lucius, R., Loos-Frank, B.: Contributions to the phylogeny of the Cyclophyllidea (Cestoda) inferred from mitochondrial 12S rDNA. *J. Mol. Evol.* 48(5), 586–596 (1999)
6. Coleman, A.W.: ITS2 is a double-edged tool for eukaryote evolutionary comparisons. *Trends Genet.* 19(7), 370–375 (2003)
7. Schultz, J., Maisel, S., Gerlach, D., Müller, T., Wolf, M.: A common core of secondary structure of the internal transcribed spacer 2 (ITS2) throughout the Eukaryota. *RNA* 11(4), 361–364 (2005)
8. Billoud, B., Guerrucci, M.A., Masselot, M., Deutsch, J.S.: Cirripede phylogeny using a novel approach: molecular morphometrics. *Mol. Biol. Evol.* 17(10), 1435–1445 (2000)
9. Thompson, J.D., Higgins, D.G., Gibson, T.J.: CLUSTAL W: improving the sensitivity of progressive multiple sequence alignment through sequence weighting, position-specific gap penalties and weight matrix choice. *Nucleic Acids Res.* 22, 4673–4680 (1994)
10. Tamura, K., Dudley, J., Nei, M., Kumar, S.: MEGA4: Molecular Evolutionary Genetics Analysis (MEGA) software version 4.0. *Mol. Biol. Evol.* 24(8), 1596–1599 (2007)
11. Ronquist, F., Huelsenbeck, J.F.: MRBAYES 3: Bayesian phylogenetic inference under mixed models. *Bioinformatics* 19, 1572–1574 (2003)
12. Page, R.D.: TreeView: an application to display phylogenetic trees on personal computers. *Comput. Appl. Biosci.* 12, 357–358 (1996)
13. Zuker, M.: Mfold web server for nucleic folding and hybridization prediction. *Nucleic Acids Res.* 31, 3406–3415 (2003)
14. Siebert, S., Backofen, R.: MARNA: multiple alignment and consensus structure prediction of RNAs based on sequence structure comparisons. *Bioinformatics* 21, 3352–3359 (2005)
15. Wolf, M., Ruderisch, B., Dandekar, T., Schultz, J., Müller, T.: ProfDistS (profile-) distance based phylogeny on sequence-structure alignments. *Bioinformatics* 24(20), 2401–2402 (2008)
16. Saitou, N., Nei, M.: The neighbor-joining method: A new method for reconstructing phylogenetic trees. *Mol. Biol. Evol.* 4, 406–425 (1987)
17. Tamura, K., Dudley, J., Nei, M., Kumar, S.: Prospects for inferring very large phylogenies by using the neighbor-joining method. *Proc. Natl. Acad. Sci. USA* 101, 11030–11035 (2004)
18. Dayhoff, M.O., Eck, R.V., Chang, M.A., Sochard, M.R.: Atlas of Protein Sequence and Structure. National Biomedical Research Foundation, Silver Spring, Maryland (1965)
19. Nei, M., Kumar, S.: *Molecular Evolution and Phylogenetics*. Oxford University Press, New York (2000)
20. Cote, C.A., Peculis, B.A.: Role of the ITS 2-proximal stem and evidence for indirect recognition of processing sites in pre-rRNA processing in yeast. *Nucleic Acids Res.* 29, 2106–2116 (2001)
21. Brooks, D.R., Hoberg, E.P., Weekes, P.J.: Preliminary phylogenetic systematic analysis of the major lineages of the Eucestoda (Platyhelminthes: Cercomeria). *Proceedings of the Biological Society of Washington* 104, 651–688 (1991)
22. Smith, A.B., Lafay, B., Christen, R.: Comparative variation of morphological and molecular evolution through geologic time: 28S ribosomal RNA versus morphology in echinoids. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 338, 365–382 (1992)

# ***ReHap*: A Framework for Single Individual Haplotyping from Next-Generation Sequencing Data**

F. Geraci and M. Pellegrini

Istituto di Informatica e Telematica, CNR – Consiglio Nazionale delle Ricerche  
56100 Pisa, Italy

{filippo.geraci,marco.pellegrini}@iit.cnr.it

**Abstract.** Next-Generation Sequencing technologies (NGS) are transforming today's biology by making it economically feasible to read the complete genome of individuals. Single nucleotide polymorphism (SNP) is the most common form of individual DNA variation; and the set of SNPs present in a chromosome (called the *haplotype*) is of interest in a wide area of applications in molecular biology and biomedicine. Personalized haplotyping of (portions of/all) the chromosomes of individuals through NGS is one of the most promising basic ingredients leading to effective personalized medicine (including diagnosis, and eventually therapy).

## **1 Introduction**

The Sanger automated sequencing technology [1], which has dominated the scene for the past 20 years, has been key to the determination of the genome of several higher species, most notably *homo sapiens* [2]. However its limitations due to high costs in labour and materials became ever more apparent over time, as biologists required higher throughput and more economical sequencing equipment [3,4]. In recent years Sanger sequencing has also evolved, but the most promising technological advances are due to the so called Next Generation Sequencing technologies [5]. The steady decrease in the cost of sequencing equipment will thus soon render cost-effective the effort of producing complete genomic profiles of individuals versus collecting data on predefined genetic markers, for which other technologies (such as those based on micro-arrays) may be more suitable. In a second, converging, trend the attention of researchers has shifted more and more from what is common among members of a species (the genome) to what is different among members of the same species, thus to the individual variations in the chromosomal content. As a consequence, we witness the rising importance of profiling individual genetic features for the purposed of individual genomic medicine [6].

The first publication of a complete individual diploid human genome sequence has been announced [7] in 2007, and two individual diploid human sequences were published in 2008 [8,9]. An ambitious project to determine the genetic profile of about 1200 individuals all over the world has been launched in 2008 (see. “1000 Genomes Project - A Deep Catalog of Human Genetic Variation” <http://www.1000genomes.org>). As for the cost-effectiveness of the sequencing technology, the current aim of the research community is to attain within a few years a cost of 1000 USD per individual



genome sequenced [10]. The cost of state-of-the-art technology is decreasing [11], so to make the 1000 USD target a realistic one.

The role of individual genomic variations and their impact in the emergence of diseases is also at the core of the so called *Common disease common variation* hypothesis (CDCV) [12], [13]. The compilation of complete individual genomes including rare (or private) variants might help in predicting traits and diseases in that particular person. Thus it is important to produce accurate individual haplotype maps in a timely manner and with reduced costs.

A key component of the reconstruction pipeline is the so called “haplotype assembly” problem, also known as the “single individual haplotyping problem” (SIH), ([14], [15], [16]). During the fragment generation phase, depending on the precise technology used [17], a large set of fragments of length in the range of 40/900 bases are read which cover with multiple overlaps the (portion of) chromosome being analyzed. By the use of reference maps it is relatively easy to locate the position and orientation of such fragments in the chromosome, but not to determine the association of the fragments to the two homologous copies (paternal and maternal) of a chromosome. The haplotype assembly problem is the problem of determining the association of each fragment to one of the two homologous chromosomes and to determine the haplotype (i.e. the value of the bases in the SNP positions of the two homologous chromosomes). In absence of errors or gaps, with sufficient coverage, this problem is easily solved, however more realistic models of the problem take into account several types and sources of error/noise in the data, and in this setting the problem is much more challenging (intractable in some cases). Several algorithms and heuristics have been proposed in the literature to solve the haplotype assembly problem (see e.g. [18,19,20,21,22,23,24,25,26,27]).

A common feature of all these researches is that, due to lack of a common benchmark and of publicly available implementations in a unified framework, experimental comparisons among different methods are based on a narrow choice of data, parameters and algorithms. Thus, in the view of large scale haplotyping projects, when the need comes to test the available algorithms onto the set of data and parameters of a specific experiment and/or technology, the question as to the best algorithm to employ is currently not explored to the its full extent. To remedy this situation, and to provide a service to the bioinformatics community, we have developed the web-based tool (*ReHap*) that provides an easy to use environment for testing several algorithms, on a variety of data sets, and with a variety of parameter settings. We envision three modalities of use for our tool.

- (I) In modality (I) the user supplies as input the two testing haplotype strings and sets a series of parameters relative to the simulation of the fragment generation. The system simulates the fragmentation process, applies any number of algorithms (among the 5 methods currently supported) and is able to measure the reconstruction rate and the reconstruction time of those methods.
- (II) In modality (II) the user provides to the system a matrix of fragments and the system returns the reconstructed haplotypes. Note that in this case the reconstruction error cannot be computed by *ReHap*. However the quality of the reconstruction can be assessed by the user that has generated the fragments.

- (III) In modality (III) the user has no data yet but an hypothesis to be tested on the basic parameters of the experiment/technology to be employed. In this case *ReHap* will use data from Hapmap ([28]) as haplotype test data. The reconstruction rate and the reconstruction time is then measured.

Modality III is interesting since in this modality *ReHap* can give guidance in an early critical stage of experimental design. In particular the length of the haplotypes chunks to be reconstructed, the (minimum, average) coverage of each SNP by fragments, and the error rate that can be handled in the final reconstruction phase are critical parameters to be assessed in early stage of design. Modality I and II can be useful by researchers wishing to develop new assembly algorithms (even exploiting additional private information) since *ReHap* can provide a common testing framework and a reference benchmark for the assessment of performance.

*ReHap* is complementary to other tools such as *Haploview* [29], where the emphasis is on visualization and computation of statistics for haplotype data from population studies (mainly genotype data).

The paper is organized as follows. Section 2 describes the web interface and the practical use of the tool. Section 3 gives a brief self-contained description of the five reconstruction algorithms currently supported. Section 4 describes the main output quality index (i.e. the reconstruction rate). Section 5 gives an example of the results that can be obtained with this tool.

## 2 Functionalities

*ReHap* allows the user to generate a SNP matrix starting from a pair of haplotypes provided by the HapMap project [28] (or submits his/her own data) according to some user-defined parameters. Haplotypes in [28] came from four different populations. For each population, a number ranging from 44 to 60 of different individuals, equally divided in males and females, are available (in total 209 unrelated individuals). For each individual the two haplotypes of chromosomes 1-22 are available. For females also the haplotypes of X chromosome are available. The SNP matrix is generated according to a model similar to that proposed in [30], which is considered a good approximation of the actual physical process of shotgun sequencing technology. The generator allows users to select a specific chromosome, individual and population. Once selected the desired haplotype, it is possible to control many other parameters. Among them: the haplotype length, the average coverage of each haplotype, the error probability, a range specifying the length of each fragment, a range specifying the expected hamming distance of the two haplotypes, and the gaps rate. By default *ReHap*'s parameters are set with technologically reasonable values. *ReHap* allows also the user to upload his/her own SNP matrix. See Figure 1 for a snapshot of the control pannel for the fragment generation process.

The generated SNP matrix (with and without errors) can be downloaded or inspected inside *ReHap*. Errors are highlighted using color codes and the correct value showed just clicking on an error basis. Besides the SNP matrix, *ReHap* also computes as a reference baseline the output of an "omniscient algorithm" (called the *baseline*) that

**Fig. 1.** A snapshot of *ReHap*'s interface in which it is possible to select the parameters of the fragmentation process

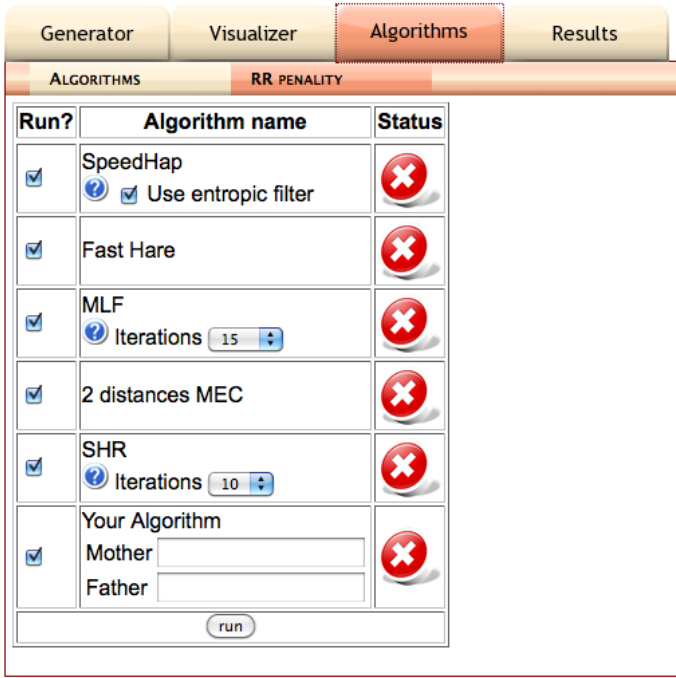
can access the true subdivision of the fragments but still decides the imputation of SNPs based on majority.

*ReHap* allows the user to test and evaluate the effectiveness of five recent reconstruction algorithms: SpeedHap [25] (a robust algorithm for individual haplotype reconstruction developed by the authors), Fast Hare [20] (a well known algorithm in the bio-informatics community), MLF [21], a two-distances clustering algorithm for the MEC model [23], and SHR [26] (a probabilistic approach to the SIH problem). Moreover *ReHap* allows the user to upload her/his own solution to evaluate.

All methods have been re-implemented by the authors. SpeedHap has a flag to enable the use of an entropic filter to break ties in the reconstruction of the haplotypes. MLF and SHR accept as a parameter the number of trials to run before selecting the “best” haplotypes. Besides running time, for each selected algorithm *ReHap* displays the returned haplotype and highlights the reconstruction errors. Moreover a synthetic numerical evaluation of the performance of each algorithm is computed by a weighted variant of the standard *reconstruction rate*. All these outputs are available when the data is generated by *ReHap*. When the SNP matrix is uploaded from the user, only the reconstructed haplotypes are returned. *ReHap* provides a contextual help for each feature of the interface, and a general help page.

### 3 Algorithms Supported

*ReHap* implements five recently proposed algorithms for the haplotype assembly problem: Speedhap [25], FastHare [20], MLF [21], 2D [23] and SHR [26]. These five algorithms have been selected because (i) they represent different approaches to the assembly problem and (ii) they are described in the literature with sufficient detail so that accurate re-implementation is possible.



**Fig. 2.** A snapshot of *ReHap*'s interface in which it is possible to select the algorithms to run and the corresponding parameters, or to upload a candidate solution

A common feature of the above cited algorithms is the use of a notion of distance similar to the standard Hamming distance to compare two fragments. When considering a certain position of the compared strings, if both character in that position are not gaps they contribute to the distance as in the case of standard Hamming distance, otherwise they do not contribute.

As shown in figure 2, *ReHap* also allows the user to upload a pair of candidate consensus haplotypes generated by the user with a private algorithm. This gives to the user the opportunity of easily comparing her/his solution with all the other algorithms with respect to the same scoring function.

In the following, for sake of completeness, we provide synthetic description of the implemented algorithms.

SpeedHap [25] follows a multi-phase greedy approach: it has four phases and makes choices that are optimal in a local sense. In each phase SpeedHap performs three tasks: 1) detect likely positions of errors 2) allocate fragments to the two partially built haplotype strings, and 3) build partial haplotype strings deciding via majority on ambiguous SNPs calls. The difference among the four phases is twofold: on one hand the algorithm uses knowledge built up in the previous phases, and on the other hand in passing from one phase to the next it relaxes the conditions for the decisions to be taken regarding tasks 1), 2) and 3). The aim is to be always very conservatives in the choices made so to

introduce as little error as possible in the early phases. Moreover SpeedHap [25] solves ties in the allele determination using an entropy-based biased filter so to give priority to choices that reduce the entropy of the final solution.

Fash Hare [20] is the oldest simple and practical heuristic algorithm for the haplotype reconstruction problem. It is widely used as a benchmark for comparisons. In a nutshell, Fast Hare works as follows: as first step the fragments of the SNP matrix are sorted so that a fragment precedes another fragment if the position of the first non-gap character of the former is lower or equal to the first non-gap position of the latter. The main loop of Fast Hare consists in processing the fragments in the SNP matrix one at time according to their ordering and dividing them in two sets. The first set is initialized with the first fragment while the second set is left empty. At each iteration, from each set Fast Hare derives the consensus haplotypes. A fragment is compared to the consensus haplotypes and assigned to the set with lower distance.

The MLF algorithm [21] is based on the well-known one-pass  $k$ -means clustering algorithm due to McQueen [31]. The procedure initialization consists in randomly splitting the SNP matrix in two partitions and, for each one, compute its consensus string. Then, in the main loop fragments are partitioned again in two sets according to the following procedure. Each fragment is compared with the two consensus strings and assigned to the set associated to the closest haplotype. Once all fragments are processed two new consensus strings can be computed and returned. Since this procedure performance strongly depends from the initial random partition, it is repeated a certain (large) number of times and as final result is returned the pair of consensus haplotypes that minimize the MEC (Minimum Error Correction) score.

The 2D algorithm [23] main contribution is the introduction of a notion of distance that overcome the structural drawbacks of Hamming distance which assign distance 0 to both: two equal strings and two strings which cover two disjoint set of SNPs. The distance introduced in 2D not only gives a penalty to different characters in a certain position, but also gives a bonus to equal (not gaps) characters in a certain position. The procedure goal is to partition the rows of the SNP matrix in two partitions and works as follows: using the Hamming distance compute the two furthest fragments and initialize each partition with one of them. Each fragment is compared using the Hamming distance with the consensus strings derived from each partition. In case of tie the second distance function is used to break the tie and select the partition to which assign the fragment. The corresponding consensus haplotype is updated.

The SHR algorithm [26] uses a probabilistic framework to approach the SIH problem. Fragments in the SNP matrix are divided in two partitions according with the following procedure. At the beginning two fragments  $f_1$  and  $f_2$  randomly extracted from the SNP matrix are used to initialize two empty sets. Then, each fragment is compared to  $f_1$  and  $f_2$  and inserted in the set corresponding to the lower distance. Once all fragments are assigned to a set, the procedure computes the MEC score induced from  $f_1$ ,  $f_2$  and the computed bipartition. Due to its probabilistic nature, the above procedure is repeated a certain number of times and as final bipartition it is selected the one with lower MEC score. At the end the consensus haplotypes are computed by majority from the final bipartition. The main contribution of this algorithm stands in the theoretical framework developed by its authors.




**Fig. 3.** A snapshot of *ReHap*'s interface in which are shown the results of two algorithms. For each one *ReHap* gives the consensus sequence (in which errors are highlighted), the reconstruction rate and the running time.

## 4 Algorithms Evaluation

*ReHap* is endowed with an interface in which details about the results of each algorithm are reported. As shown in figure 3, for each algorithm, *ReHap* reports the consensus haplotypes (in which errors are highlighted), the running time and the reconstruction rate. As mentioned in Section 2, we slightly modified the reconstruction rate formula. We observed that, using the standard formula, the insertion of a gap in the consensus string is considered as good as the insertion of the correct character. This can introduce a positive bias for those algorithms that output a gap when they are unable to reconstruct the haplotype in a certain position. In order to control such a situation, *ReHap* introduces a parameter in the reconstruction rate that allows to penalize gaps in the consensus string. Our modified version of the reconstruction rate is:

$$RR_{\hat{H}, H} = 1 - \frac{\min(D(h_1, \hat{h}_1) + D(h_2, \hat{h}_2), D(h_1, \hat{h}_2) + D(h_2, \hat{h}_1))}{2m}$$

where  $H = \{h_1, h_2\}$  is the correct haplotype pair,  $\hat{H} = \{\hat{h}_1, \hat{h}_2\}$  is the consensus haplotype pair,  $m$  is the length of each haplotype and  $D()$  is defined as follows:



**Fig. 4.** Reconstruction rate of 5 reconstruction algorithms and the baseline algorithm for covering in the range [3,...,10]. The other parameters are fixed: error rate 0.15 and haplotype length 250 bases. Each point is the average of 20 experiments with randomly chosen haplotype strings of Hapmap.

$$D(h_i, \hat{h}_j) = \sum_{k=1}^m d(h_i[k], \hat{h}_j[k])$$


where

$$d(h_i[k], \hat{h}_j[k]) = \begin{cases} 0 & \text{if } h_i[k] = \hat{h}_j[k] \\ \alpha & \text{elseif } \hat{h}_j[k] = - \\ 1 & \text{otherwise} \end{cases}$$


Note that when  $\alpha = 0$  the above formula reduce to the standard reconstruction rate. When  $\alpha = 1$  the insertion of a gap in the consensus string is penalized as a reconstruction error. The user can select any value for  $\alpha$  in the range  $[0, 1]$ .

## 5 Experiments

In this section we report on experiments in modality (III), using Hapmap data to explore relative performance in a range of possible parameter choices. We concentrate on three parameters (covering, error rate, length) and we measure the reconstruction rate, with  $\alpha = 0$ , of the reconstructed haplotype returned by five algorithms. The reconstruction rate is a real number in the range  $[0.0, \dots, 1.0]$  and has value 1.0 for perfect reconstruction. We also include the outcome of the *baseline* algorithm that has access to the true association of the fragments to the two haplotype strings and decides the base call by majority vote. This algorithm is clearly useful only as a reference in simulations, as an ideal for the proper algorithms that have no access to this information. To the best of our knowledge this is the first published report of a comparison among five different haplotype assembly algorithms on a common data and algorithmic framework.



**Fig. 5.** Reconstruction rate of 5 reconstruction algorithms and the baseline algorithm for error rates in the range [0.0,...,0.30]. The other parameters are fixed: covering 8 and haplotype length 250 bases. Each point is the average of 20 experiments with randomly chosen haplotype strings of Hapmap.



**Fig. 6.** Reconstruction rate of 5 reconstruction algorithms and the baseline algorithm for length in the range [50,...,500]. The other parameters are fixed: error rate 0.15 and covering 8. Each point is the average of 20 experiments with randomly chosen haplotype strings of Hapmap.

As fragment coverage increases we expect the performance of all the algorithms also to improve. This is expectation confirmed for all algorithms, except SHR that seems unaffected. From figure 4 one can observe that for high coverage (say 10 and more)



SpeedHap, Fast Hare, MLF and 2D match the baseline and are above 95% reconstruction rate, while at low coverage (3-7) Speedhap is notably more accurate and very close to the baseline.

From figure 5 one observes that increasing the error rate in the range [0.0,...,0.3] the baseline reference algorithm is also affected. This is to be expected, as the accumulation of many errors in the same SNP position makes it impossible even for the baseline method to decide correctly by majority. Speedhap although performing worse than the baseline is not far from it, and it is much better than the other 4 algorithms at high error rate, attaining a respectable reconstruction rate only 0.1 below the baseline, even with 30% of reading errors. Fast Hare MLF, and 2D perform reasonably well, better than SHR, but all four converge to the similar reconstruction rate (about 0.70-0.75) for the higher reading error situation.

From figure 6 one can observe that for the set of parameters tested the length of the haplotype has little influence on the relative performance of the 5 algorithms w.r.t. the baseline. Speedhap, Fast hare and 2D have reconstruction rate above 90% for the range of length considered, while MLF and SHR have reconstruction rate in a lower region (0.75% - 0.85%).

## 6 Conclusions

The integrated framework *ReHap* for testing and comparing five different haplotype assembly algorithms is described. Our hope is that *ReHap* will help the bioinformatics community in selecting the most suitable algorithm for each specific haplotype assembly task, as large scale individual haplotyping programmes using NGS techniques are getting momentum. This system can be accessed and used by researchers freely at the url: <http://bioalgo.iit.cnr.it/rehap/> without any login requirement.

**Acknowledgements.** This work was supported by the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement N 223920 (Virtual Physiological Human Network of Excellence).

## References

1. Pop, M.: Shotgun sequence assembly. *Advances in Computers* 60, 193–248 (2004)
2. Istrail, S., Sutton, G.G., Florea, L., Halpern, A.L., Mobarry, C.M., Lippert, R., Walenz, B., Shatkay, H., Dew, I., Miller, J.R., Flanigan, M.J., Edwards, N.J., Bolanos, R., Fasulo, D., Halldorsson, B.V., Hannenhalli, S., Turner, R., Yooseph, S., Lu, F., Nusskern, D.R., Shue, B.C., Zheng, X.H., Zhong, F., Delcher, A.L., Huson, D.H., Kravitz, S.A., Mouchard, L., Reinert, K., Remington, K.A., Clark, A.G., Waterman, M.S., Eichler, E.E., Adams, M.D., Hunkapiller, M.W., Myers, E.W., Venter, J.C.: Whole-genome shotgun assembly and comparison of human genome assemblies. *Proceedings of the National Academy of Sciences of the United States of America* 101(7), 1916–1921 (2004)
3. Metzker, M.: Sequencing technologies - the next generation. *Nature reviews* 11 (2010)
4. Schuster, S.: Next-generation sequencing transforms today's biology. *Nature Methods* 5(1) (January 2008)

5. Metzker, M.: Emerging technologies in dna sequencing. *Genome Research* 15, 1767–1776 (2005)
6. Crawford, D., Nickerson, D.: Definition and clinical importance of haplotypes. *Annu. Rev. Med.* 56, 303–320 (2005)
7. Levy, S., Sutton, G., Ng, P., Feuk, L., Halpern, A., et al.: The diploid genome sequence of an individual human. *PLoS Biology* 5(10) (2007)
8. Wang, J., et al.: The diploid genome sequence of an asian individual. *Nature* 456, 60–65 (2008)
9. Wheeler, D., et al.: The complete genome of an individual by massively parallel dna sequencing. *Nature* 452, 872–876 (2008)
10. Mardis, E.R.: Anticipating the \$1,000 genome. *Genome Biology* 7 (July 2006)
11. von Bubnoff, A.: Next-generation sequencing: The race is on. *Cell* 132(5), 721–723 (2008)
12. Iles, M.M.: What can genome-wide association studies tell us about the genetics of common disease? *PLoS Genet.* 4(2) (February 2008)
13. Schork, N.J., Murray, S.S., Frazer, K.A., Topol, E.J.: Common vs. rare allele hypotheses for complex diseases. *Current Opinion in Genetics & Development* 19(3), 212–219 (2009)
14. Halldórsson, B.V., Bafna, V., Edwards, N., Lippert, R., Yooseph, S., Istrail, S.: A survey of computational methods for determining haplotypes. In: Istrail, S., Waterman, M.S., Clark, A. (eds.) DIMACS/RECOMB Satellite Workshop 2002. LNCS (LNBI), vol. 2983, pp. 26–47. Springer, Heidelberg (2004)
15. Zhao, Y., Xu, Y., Zhang, Q., Chen, G.: An overview of the haplotype problems and algorithms. *Frontiers of Computer Science in China* 1(3), 272–282 (2007)
16. Bonizzoni, P., Vedova, G.D., Dondi, R., Li, J.: The haplotyping problem: an overview of computational models and solutions. *J. Comput. Sci. Technol.* 18(6), 675–688 (2003)
17. Morozova, O., Marra, M.A.: Applications of next-generation sequencing technologies in functional genomics. *Genomics* 92(5), 255–264 (2008)
18. Rizzi, R., Bafna, V., Istrail, S., Lancia, G.: Practical algorithms and fixed-parameter tractability for the single individual SNP haplotyping problem. In: Guigó, R., Gusfield, D. (eds.) WABI 2002. LNCS, vol. 2452, pp. 29–43. Springer, Heidelberg (2002)
19. Li, L., Kim, J.H., Waterman, M.S.: Haplotype reconstruction from SNP alignment. In: Proceedings of the Seventh Annual International Conference on Computational Molecular Biology, pp. 207–216. ACM Press, New York (2003)
20. Panconesi, A., Sozio, M.: Fast hare: A fast heuristic for single individual SNP haplotype reconstruction. In: Jonassen, I., Kim, J. (eds.) WABI 2004. LNCS (LNBI), vol. 3240, pp. 266–277. Springer, Heidelberg (2004)
21. Zhao, Y.Y., Wu, L.Y., Zhang, J.H., Wang, R.S., Zhang, X.S.: Haplotype assembly from aligned weighted SNP fragments. *Computational Biology and Chemistry* 29(4), 281–287 (2005)
22. Lindsay, S.J., Bonfield, J.K., Hurler, M.E.: Shotgun haplotyping: a novel method for surveying allelic sequence variation. *Nucl. Acids Res.* 33(18) (2005)
23. Wang, Y., Feng, E., Wang, R.: A clustering algorithm based on two distance functions for mec model. *Computational Biology and Chemistry* 31(2), 148–150 (2007)
24. Cilibrasi, R., van Iersel, L., Kelk, S., Tromp, J.: On the complexity of the single individual SNP haplotyping problem. *Algorithmica* (2007)
25. Genovese, L., Geraci, F., Pellegrini, M.: Speedhap: An accurate heuristic for the single individual snp haplotyping problem with many gaps, high reading error rate and low coverage. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 5(4), 492–502 (2008)

26. Chen, Z., Fu, B., Schweller, R.T., Yang, B., Zhao, Z., Zhu, B.: Linear time probabilistic algorithms for the singular haplotype reconstruction problem from snp fragments. In: Brazma, A., Miyano, S., Akutsu, T. (eds.) APBC, vol. 6, pp. 333–342. Imperial College Press, London (2008)
27. Bansal, V., Bafna, V.: Hapcut: an efficient and accurate algorithm for the haplotype assembly problem. *Bioinformatics* 24(16), i153–i159 (2008)
28. Consortium, T.I.H.: A haplotype map of the human genome. *Nature* 437, 1299–1320 (2005)
29. Barrett, J.C., Fry, B., Maller, J., Daly, M.J.: Haploview: analysis and visualization of LD and haplotype maps. *Bioinformatics* 21(2), 263–265 (2005)
30. Myers, G.: A dataset generator for whole genome shotgun sequencing. In: Proceedings of the Seventh International Conference on Intelligent Systems for Molecular Biology, pp. 202–210. AAAI Press, Menlo Park (1999)
31. McQueen, J.: Some methods for classification and analysis of multivariate observations. In: Proc. of the 5<sup>th</sup> Berkeley Symposium on Mathematical Statistics and Probability, vol. 1, pp. 281–297. University of California Press, Berkeley (1967)

# Inference of Large Phylogenies Using Neighbour-Joining

Martin Simonsen, Thomas Mailund, and Christian N.S. Pedersen

Bioinformatics Research Centre (BIRC), Aarhus University  
C. F. Møllers Allé 8, DK-8000 Århus C, Denmark  
{zxr, mailund, cstorm}@birc.au.dk

**Abstract.** The neighbour-joining method is a widely used method for phylogenetic reconstruction which scales to thousands of taxa. However, advances in sequencing technology have made data sets with more than 10,000 related taxa widely available. Inference of such large phylogenies takes hours or days using the Neighbour-Joining method on a normal desktop computer because of the  $O(n^3)$  running time. RapidNJ is a search heuristic which reduce the running time of the Neighbour-Joining method significantly but at the cost of an increased memory consumption making inference of large phylogenies infeasible. We present two extensions for RapidNJ which reduce the memory requirements and allows phylogenies with more than 50,000 taxa to be inferred efficiently on a desktop computer. Furthermore, an improved version of the search heuristic is presented which reduces the running time of RapidNJ on many data sets.

## 1 Introduction

The neighbour-joining (NJ) method [10] is a widely used method for phylogenetic inference, made popular by reasonable accuracy combined with a cubic running time algorithm by Studier and Kepler [14]. The NJ method scales to hundreds of species, and while it is usually possible to infer phylogenies with thousands of species, tens or hundreds of thousands of species is computationally infeasible.

Implementations like QuickTree [5] and QuickJoin [6,7] use various approaches to reduce the running time of NJ considerably, and recently we presented a new heuristic, RapidNJ [12], which uses a simple branch and bound technique to reduce the running time even further. Though RapidNJ is able to build NJ trees very efficiently it requires, like the canonical NJ method,  $O(n^2)$  space to build a tree with  $n$  taxa. The space consumption of RapidNJ, and the NJ method in general, is thus a practical problem when building large trees, and since RapidNJ uses some additional data structures of size  $O(n^2)$ , this method has limited application to data sets with more than 10,000 taxa which is of interest when building phylogenetic trees from e.g. Pfam [4] sequence data.

In this paper we present two extensions for RapidNJ that reduce the memory requirements of RapidNJ. The first extension uses a simple heuristic which takes advantage of RapidNJ's memory access pattern to reduce the internal memory (RAM) consumption. The second extension is based on the first extension and makes use of external memory, i.e. a hard disk drive (HDD), to reduce internal memory consumption. We also present an improved version of the search heuristic used in RapidNJ which increases performance on data sets that RapidNJ has difficulties handling.

The two extensions combined with the improved search heuristic allow large NJ trees to be reconstructed efficiently which is important as sequence family data with more than 50,000 taxa are becoming widely available [42]. Also, the NJ method is used as a clustering method in both micro array data analysis and metagenomics where data sets can become very large. Using the methods proposed in this paper, clustering of large data sets can be handled efficiently on normal desktop computers.

We evaluate the performance of the extended RapidNJ method (ERapidNJ), by comparing running times of an implementation of the ERapidNJ method with other fast implementations for building canonical NJ trees.

## 2 Methods

### 2.1 The Neighbour-Joining Method

NJ is a hierarchical clustering algorithm. It takes a distance matrix  $D$  as input, where  $D(i, j)$  is the distance between clusters,  $i$  and  $j$ . Clusters are then iteratively joined using a greedy algorithm, which minimises the total sum of branch lengths in the tree. Basically the algorithm uses  $n$  iterations, where two clusters,  $i$  and  $j$ , are selected and joined into a new cluster in each iteration. The pair  $i$  and  $j$  is selected by minimising

$$Q(i, j) = D(i, j) - u(i) - u(j), \tag{1}$$

where

$$u(l) = \sum_{k=0}^{r-1} D(l, k) / (r - 2), \tag{2}$$

and  $r$  is the number of remaining clusters. When a minimum  $q$ -value  $q_{\min} = \min_{0 \leq i, j < r} Q(i, j)$  is found,  $D$  is updated, by removing the  $i$ 'th and  $j$ 'th row and column. A new row and a new column are inserted with the distances to the new cluster. The distance between the new cluster  $a = i \cup j$  and one of the remaining clusters  $k$ , is calculated as

$$D(a, k) = \frac{D(i, k) + D(j, k) - D(i, j)}{2}. \tag{3}$$

The result of the algorithm is an unrooted bifurcating tree where the initial clusters correspond to leafs and each join corresponds to inserting an internal node in the tree.

### 2.2 RapidNJ

RapidNJ [12] computes an upper bound for the distance between clusters which is used to exclude a large portion of  $D$  when searching for a minimum  $q$ -value. To utilise the upper bound two new data structures,  $S$  and  $I$ , are needed. Matrix  $S$  contains the distances from  $D$  but with each row sorted in increasing order and matrix  $I$  maps the ordering in  $S$  back to positions in  $D$ . Let  $o_1, o_2, \dots, o_n$  be a permutation of  $1, 2, \dots, n$  such that  $D(i, o_1) \leq D(i, o_2) \leq \dots \leq D(i, o_n)$ , then

$$S(i, j) = D(i, o_j), \tag{4}$$

and

$$I(i, o_j) = j. \quad (5)$$

The upper bound is computed and used to speed up the search for a minimum  $q$ -value as follows.

1. Set  $q_{\min} = \infty$ ,  $i = -1$ ,  $j = -1$ ,  $u_{\max} = \max(u(l))$
2. for each row  $r$  in  $S$  and column  $c$  in  $r$ :
  - (a) if  $S(r, c) - u(r) - u_{\max} > q_{\min}$  then move to the next row.
  - (b) if  $Q(r, I(r, c)) < q_{\min}$  then set  $q_{\min} = Q(r, I(r, c))$ ,  $i = r$  and  $j = I(r, c)$ .

The algorithm searches  $S$  row-wise and stops searching within a row when the condition

$$S(r, c) - u(r) - u_{\max} > q_{\min} \quad (6)$$

is true, or the end of a row is reached. If we reached an entry in  $S$  where (6) is true, we are looking at a pair  $(i, j)$ , where  $D(i, j)$  is too large for  $(i, j)$  to be a candidate for  $q_{\min}$ , and because  $S$  is sorted in increasing order, all the following entries in  $S(i)$  can now be disregarded in the search.


When the cluster-pair  $(i', j')$  with the minimum  $q$ -value is found,  $D$  is updated as described in Sec. 2.1. The  $S$  and  $I$  matrices are then updated to reflect the changes made in the  $D$  as follows. Row and column  $i'$  and  $j'$  are marked as deleted and entries in  $S$  belonging to these rows/columns are then identified using  $I$  and ignored in the following iterations of the NJ method. A new row containing the distances of the new cluster is sorted and inserted into  $S$ .

### 2.3 Reducing the Memory Consumption of RapidNJ

RapidNJ consumes about four times more memory than a straightforward implementation of canonical neighbour-joining, which makes it impractical to use on large data sets. We propose an extension to RapidNJ which reduces the memory consumption significantly while only causing a small reduction in performance.

First we reduce the size of the  $D$ -matrix. RapidNJ stores the complete  $D$ -matrix in memory, even though only the upper or lower triangular matrix is needed, because it allows for a more efficient memory access scheme. By only storing the lower triangular matrix, the size of  $D$  is halved without affecting performance seriously.

Secondly, the size of  $S$  and, consequently,  $I$  are reduced. As seen in Fig. 1, RapidNJ rarely needs to search more than a few percent of each row in  $S$ . Hence it is not necessary to store the full  $S$ -matrix in memory to receive a speed up similar to the original RapidNJ method. An increase in both maximum and average search depth is observed when the last quarter of the clusters remains, but as the number of remaining clusters is low at this point, the increase only causes a relatively small increase in the total number of entries searched. The size of  $S$  is reduced by only storing as many columns of  $S$  as can fit in the available internal memory after  $D$  has been loaded. Of course we might not store enough columns of  $S$  to complete the search for  $q_{\min}$  in all rows of  $S$ , i.e. we might not reach an entry where (6) becomes true. If this happens we simply search the corresponding row in  $D$ .



**Fig. 1.** The maximum and average number of entries of each row in  $S$  that RapidNJ searched during each iteration of the NJ method when building a typical tree containing 10,403 taxa

There is a lower limit on the number of columns of  $S$  we must store before the performance is severely affected, but there is no exact number as it depends on the data set. Our experiments imply that at least 5% of the columns in  $S$  are needed to receive a significant speed up in general.

#### 2.4 An I/O Algorithm for Building Very Large Trees

Even when using the extension described in Sec. 2.3, RapidNJ will run out of memory at some point and begin to swap out memory pages to the HDD. This will seriously reduce the performance because the data structures used by RapidNJ are not designed to be I/O efficient. I/O efficiency is achieved by accessing data in the external memory in blocks of typical 4-8 KB corresponding to the block size  $B$  of the HDD used [1], and it is often better to access data in blocks larger than  $B$  to take full advantage of hardware and software caching. However, even when using an I/O efficient algorithm, accessing data in the external memory has very high latency compared to accessing data in the internal memory, thus external memory data access should be kept at a minimum.

RapidDiskNJ is an extension to RapidNJ which employs both internal and external memory storage efficiently. Because RapidNJ only uses  $S$  (and  $I$ ) to search for  $q_{\min}$ ,  $D$  can be stored in the external memory without affecting performance significantly. Furthermore, as explained in Sec. 2.3, RapidNJ usually only needs to search a small fraction of  $S$  in each iteration, so the total internal memory consumption can be reduced by only representing a sufficient part of  $S$  in the internal memory. Using external memory to store  $D$  affects the running time by a large but constant factor, thus RapidDiskNJ has

the same  $O(n^3)$  asymptotic running time as RapidNJ.  $q_{\min}$  is found as described in Sec. 2.3 the only difference being that searching  $D$  is done using the external memory.

**Data Structures.**  $D$  is stored row-wise in the external memory, so all access to  $D$  must be done row-wise as accessing a column of  $D$  would result in  $r$  I/O operations (read/write operations) assuming that an entry in  $D$  has size  $\leq B$ . A row in  $D$  can be accessed using  $\frac{r \cdot \alpha}{B}$  I/O operations where  $\alpha$  is the size of an entry in  $D$ , which is much more efficient.

As explained in Sec. 2.3 storing half of  $D$  is sufficient, but by storing the whole  $D$ -matrix in the external memory, all distances from one cluster to all other clusters can be accessed by reading one row of  $D$ . After each iteration of the NJ method, at least one column of  $D$  needs to be updated with new distances after a join of two clusters. This would trigger column-wise external memory access but by using an internal memory cache this can be avoided as described below. Deletion of columns in  $D$  is done in  $O(1)$  time by simply marking columns as deleted and then ignoring entries in  $D$  belonging to deleted columns. This gives rise to a lot of “garbage” in  $D$ , i.e. deleted columns, which needs to be removed to avoid a significant overhead. An efficient garbage collection strategy to handle this problem is proposed below.

RapidDiskNJ builds the  $S$ -matrix by sorting  $D$  row by row and for each sorted row the first  $\frac{1}{\gamma}$  entries are stored in the internal memory where the size of  $\frac{n}{\gamma}$  is  $\frac{M}{2}$  and  $M$  is the size of the internal memory. If enough columns of  $S$  can be stored in the internal memory, RapidDiskNJ can usually find  $q_{\min}$  using only  $S$  which means that RapidDiskNJ rarely needs to access the external memory.

The other half of the internal memory is used for caching columns of  $D$ . After each iteration a new column for  $D$  is created but instead of inserting this in  $D$ , The column is stored in an internal memory cache  $C$ . By keeping track of which columns have been updated and in which order, updated entries in  $D$  can quickly be identified and read from  $C$ . When  $C$  is full (i.e. the size has reached  $\frac{M}{2}$ ), all updated values in  $C$  are flushed to  $D$ , by updating  $D$  row by row which is more efficient than writing columns to  $D$  when  $C$  is large.

**Garbage Collection.** Entries belonging to deleted columns are left in both  $D$  and  $S$  after clusters are joined. We just skip these entries when we meet them. This is not a problem for small data sets but in larger data sets they need to be removed to keep  $S$  and  $D$  as small as possible. Garbage collection in both  $D$  and  $S$  is expensive so RapidDiskNJ only performs garbage collection when  $C$  is flushed. During a flush of  $C$ , all rows in  $D$  are loaded into the internal memory where deleted entries can be removed at an insignificant extra cost. By removing entries belonging to both deleted rows and columns the size of  $D$  is reduced to  $r$  which makes both searching  $D$  and future flushes of  $C$  more efficient.

Garbage collection in  $S$  is performed by completely rebuilding  $S$  during a flush of  $C$ . Our experiments showed that rebuilding  $S$  each time we flush  $C$  actually decreases performance because of the time it takes to sort  $D$ . We found that the best average performance was achieved if  $S$  was rebuild only when more than half of  $S$  consisted of garbage. During garbage collection of  $S$  the number of rows in  $S$  decreases to  $r$ , which allows more columns to be added to  $S$  so that  $S$  attains size  $\frac{M}{2}$  again.



## 2.5 Improving the Search Heuristic

RapidNJ uses the maximum average row sum,  $u_{\max}$ , to compute an upper bound on  $q$ -values. However, because  $D$  is a symmetric matrix row  $i$  in  $S$  only needs to contain  $i$  columns which makes it possible to compute tighter bounds as follows. Initially, we compute  $u_{\max}$  for each row in  $S$  as  $u(i)_{\max} = \max_{0 \leq l \leq i} (u(l))$ , instead of the global  $u_{\max}$  value. For each new row,  $i'$ , created after a join, we compute  $u(i')_{\max} = \max_{0 \leq l \leq r} (u(l))$ . Updating the existing  $u(i)_{\max}$  values can be done in time  $O(r)$  by updating  $u$ -values in the same order as the rows of  $S$  were created, assuming that the initial rows of  $S$  were created in the order, shortest to longest. Now  $u(i)_{\max} = u'_{\max}$  where  $u'_{\max}$  is the largest  $u$ -value seen when  $u(i)$  is updated. These new bounds are particularly effective on data sets containing cluttered taxa (where a large group of taxa has almost identical distances to each other and a small or zero mutual distance), which gave rise to poor performance in RapidNJ (see Fig. 2 and Fig. 3).

Redundant taxa (taxa with equal distances to all other taxa and a mutual distance of 0) are quite common in Pfam data sets. Redundant data often causes a significant loss of performance in RapidNJ because a lot of  $q$ -values fall under the upper bound at the same time forcing RapidNJ to search all pairs of redundant taxa in each iteration until they are joined. To address this problem we initially treat redundant taxa as a single taxon. When a cluster representing such a taxon is selected for a join, we only delete the cluster if the number of redundant taxa it represents drops to 0. Identifying and processing redundant taxa can be done in  $O(n^2)$  time in a preprocessing phase and reduces the problem of redundant taxa considerably (see Fig. 3).


## 3 Experiments

The methods described in Sec. 2 were used to extend the original RapidNJ tool and create the ERapidNJ tool (Source code available at <http://birc.au.dk/research/software/rapidnj/>). To assess the performance of the ERapidNJ tool, we compared running times on Pfam data sets with running times of three other tools which to our knowledge are the fastest tools available for computing canonical NJ trees at the moment.


- **QuickTree** [5]. An efficient implementation of the NJ method with a heuristic for handling redundant data.
- **QuickJoin** [7]. Reduces the running time of the NJ method by using information from previous iterations of the NJ method to reduce the search space significantly.
- **NINJA** [15]. Uses an upper bound on  $q$ -values like RapidNJ but computes a tighter bound. NINJA also uses the same technique as QuickJoin to further reduce the search space and can also utilise external memory efficiently.

QuickTree is implemented in *C*, QuickJoin and ERapidNJ in *C++* while NINJA is implemented in Java.

Tools such as Fast Neighbor-Joining [3], Clearcut [11] and FastTree [9] which modify the NJ optimisation criteria are not included in the experiments. FastTree is able to construct large trees efficiently, but as this tool operates on multiple alignments and not on distance matrices a direct comparison of the performance of ERapidNJ and FastTree is difficult. See [12] and [9] for a comparison of the original RapidNJ tool and



**Fig. 2.** A NJ tree reconstructed from a data set with 1111 taxa where the upper bound used in RapidNJ is ineffective. All taxa in the rightmost part of the tree fall under the upper bound, used in RapidNJ, simultaneously resulting in a large search space.



**Fig. 3.** The difference in running time between the original RapidNJ search heuristic and the improved search heuristic. We observe that many of the outliers have been removed when using the improved bounds on  $q$ -values and efficient redundant data handling.




Fig. 4. Running times on data sets with 100 to 8,000 taxa

some of these tools. The data used in our experiments is distance matrices computed by QuickTree using multiple alignments from the Pfam database.

The ERapidNJ tool automatically chooses one of three methods for building trees, based on the memory requirements of a given data set and the available amount of memory in the system. For small data sets the original RapidNJ method is used, for medium sized data sets the method described in Sec. 2.3 is used to reduce the memory consumption and for large inputs RapidDiskNJ described in Sec. 2.4 is used. The improved search heuristic described in Sec. 2.5 is used in all three methods to limit the impact of redundant data and reduce the search space.


NINJA, like the ERapidNJ tool, is capable of using both internal and external memory. In the experiments NINJA was configured to use only internal memory for data sets which could fit in the 2 GB memory. For larger data sets NINJA used both internal and external memory.

### 3.1 Experimental Setup

All experiments were performed on machines with an Intel Core 2 6600 2.4 GHz CPU, 2 GB 667 MHz RAM and a 7200 RPM 160 GB, Western Digital HDD. The operating system was Red Hat Enterprise 5.2 32 bit with Java 1.6 installed.

### 3.2 Results and Discussion

As seen in Fig. 4, ERapidNJ is faster than the three other tools on data sets up to 3,000 taxa. NINJA seems to suffer from an overhead on small data sets which is probably



**Fig. 5.** Running times on data sets with 5,000 to 55,000 taxa

caused by the time required to initialise Java and the data structures NINJA needs, which are larger and more complex than those used by ERapidNJ. Except for a few outliers from NINJA, ERapidNJ and NINJA have roughly the same running time on data sets with 3,000 to 7,000 taxa. On data sets with more than 7,000 taxa NINJA runs out of internal memory and starts using external memory. Both QuickJoin and QuickTree are consistently slower than ERapidNJ and QuickJoin runs out of memory on data sets with more than 7,000 taxa like NINJA.

Figure 5 shows running times on data sets with 5,000 to 55,000 taxa. Because ERapidNJ is able to scale its memory consumption to the size of the data set, we observe that ERapidNJ is significantly faster than NINJA on data sets containing less than 28,000 taxa. On larger data sets ERapidNJ is still much faster than NINJA on most data sets, and we only found two large data sets where NINJA outperformed ERapidNJ. One of these results (a data set with 49,376 taxa) is not shown in Fig. 5 for ERapidNJ because ERapidNJ did not finish within 48 hours due to cluttering of data. NINJA was able to finish this data set in 16 hours because NINJA computes much tighter bounds on  $q$ -values than ERapidNJ. NINJA also uses a technique called  $q$ -filtering [15] to further reduce the search space when searching for  $q_{min}$ . This is computationally expensive but on a few large data sets with cluttering the tighter bounds give NINJA an advantage because ERapidNJ cannot store enough columns of  $S$  in the internal memory. More memory improves ERapidNJ's performance on these data sets significantly.

The performance of QuickTree was inferior to both NINJA and ERapidNJ on all data sets. When trying to build trees with more than 32,000 taxa using QuickTree the running time exceeded 48 hours because more than 2GB of internal memory is needed to build such trees which results in memory page swapping. Since QuickTree is not I/O

efficient, page swapping causes a large time penalty which prevents QuickTree from finishing within a reasonable amount of time.

**Improving Performance by Parallelisation.** Parallelsation of the original NJ method can be done by dividing the rows of  $D$  into  $t$  sets of approximately the same size and then searching each set for  $q_{\min}$  in parallel. Similarly, ERapidNJ can be parallelised by searching rows of  $S$  in parallel. The performance of the canonical NJ method can easily be improved in this way, as searching for  $q_{\min}$  is the most time consuming step of the canonical NJ method. This is not always the case with ERapidNJ where the time spent on operations such as accessing external memory, sorting and updating data structures is similar to the total time used on searching for  $q_{\min}$  when building relatively small trees. As an example, ERapidNJ uses only 24% of the total running time on searching for  $q_{\min}$  when building a tree containing 10,403 taxa. The remaining time is spent on loading  $D$  from external memory and updating data structures after each iteration of the NJ method, which are operations that does not benefit significantly from parallelisation. Consequently, the performance gain from parallelisation of ERapidNJ is expected to be small on small data sets. On larger data sets the total time used to search for  $q_{\min}$  takes up a substantial part of the total time consumption and here parallelisation is expected to be more effective.

Experiments with a parallelised version of ERapidNJ showed a reduction of the total running time by a factor 2.2 on a 2.66 GHz quad core Intel Core 2 Duo processor compared to the unparallelised ERapidNJ on the same processor when building a tree with 25,803 taxa. When building a tree with 10,403 taxa the total running time was only reduced by a factor 1.12. Both these data sets were computed in internal memory and parallelisation of RapidDiskNJ will not increase performance significantly as RapidDiskNJ is I/O bound, i.e. most of the running time is spent on waiting for the external memory.

## 4 Conclusions

We have presented two extensions and an improved search heuristic for the RapidNJ method which both increases the performance of RapidNJ and decreases internal memory requirements significantly. Using the methods described in this paper, we were able to overcome RapidNJ's limitations regarding the memory consumption and performance on data sets containing redundant and cluttered taxa. We have presented experiments with the extended RapidNJ tool showing that canonical NJ trees containing more than 50,000 taxa can be built in a few hours on a desktop computer with only 2GB of RAM. Our experiments also showed that in comparison with the fastest tools available for building canonical NJ trees, the ERapidNJ tool is significantly faster on any size of input.

We are aware that statistical phylogenetic inference methods with better precision than distance based methods are available. However, the time complexity of these methods are high compared to the NJ method and currently they do not scale well to large data sets [13,8], which justify the existence of distance based methods as presented in this paper.

## References

1. Aggerwal, A., Vitter, T.S.: The input output complexity of sorting and related problems. *Communications of the ACM* 31(9), 1116–1127 (1988)
2. Alm, E.J., Huang, K.H., Price, M.N., Koche, R.P., Keller, K., Dubchak, I.L., Arkin, A.P.: The microbesonline web site for comparative genomics. *Genome Research* 15(7), 1015–1022 (2005)
3. Elias, I., Lagergren, J.: Fast neighbor joining. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) *ICALP 2005. LNCS*, vol. 3580, pp. 1263–1274. Springer, Heidelberg (2005)
4. Finn, R.D., Mistry, J., Schuster-Böckler, B., Griffiths-Jones, S., Hollich, V., Lassmann, T., Moxon, S., Marshall, M., Khanna, A., Durbin, R., Eddy, S.R., Sonnhammer, E.L.L., Bateman, A.: Pfam: clans, web tools and services. *Nucleic Acids Research Database Issue* 34, D247–D251 (2006)
5. Howe, K., Bateman, A., Durbin, R.: QuickTree: Building huge neighbour-joining trees of protein sequences. *Bioinformatics* 18(11), 1546–1547 (2002)
6. Mailund, T., Brodal, G.S., Fagerberg, R., Pedersen, C.N.S., Philips, D.: Recrafting the neighbor-joining method. *BMC Bioinformatics* 7(29) (2006)
7. Mailund, T., Pedersen, C.N.S.: QuickJoin – fast neighbour-joining tree reconstruction. *Bioinformatics* 20, 3261–3262 (2004)
8. Ott, M., Zola, J., Stamatakis, A., Aluru, S.: Large-scale maximum likelihood-based phylogenetic analysis on the ibm bluegene/l. In: *Proceedings of the 2007 ACM/IEEE Conference on Supercomputing*, pp. 1–11 (2007)
9. Price, M.N., Dehal, P.S., Arkin, A.P.: Fasttree: Computing large minimum-evolution trees with profiles instead of a distance matrix. *Mol. Biol. Evol.* 26(7), 1641–1650 (2009)
10. Saitou, N., Nei, M.: The neighbor-joining method: A new method for reconstructing phylogenetic trees. *Molecular Biology and Evolution* 4, 406–425 (1987)
11. Sheneman, L., Evans, J., Foster, J.A.: Clearcut: A fast implementation of relaxed neighbor-joining. *Bioinformatics* 22(22), 2823–2824 (2006)
12. Simonsen, M., Mailund, T., Pedersen, C.N.S.: Rapid neighbour-joining. In: Crandall, K.A., Lagergren, J. (eds.) *WABI 2008. LNCS (LNBI)*, vol. 5251, pp. 113–122. Springer, Heidelberg (2008)
13. Stamatakis, A.: Raxml-vi-hpc: maximum likelihood-based phylogenetic analyses with thousands of taxa and mixed models. *Oxford Journals* 22(21), 2688–2690 (2006)
14. Studier, J.A., Kepler, K.J.: A note on the neighbor-joining method of Saitou and Nei. *Molecular Biology and Evolution* 5, 729–731 (1988)
15. Wheeler, T.J.: Large-scale neighbor-joining with ninja. In: Salzberg, S.L., Warnow, T. (eds.) *WABI 2009. LNCS*, vol. 5724, pp. 375–389. Springer, Heidelberg (2009)

# Prediction of Minimum Free Energy Structure for Simple Non-standard Pseudoknot

Thomas K.F. Wong and S.M. Yiu

Department of Computer Science, The University of Hong Kong, Hong Kong  
{kfwong, smyiu}@cs.hku.hk

**Abstract.** Predicting the secondary structure with minimum free energy of an RNA molecule is an important problem in computational biology. Unfortunately, the problem is in general NP-hard if there are pseudoknots in the structure. Existing algorithms usually target at some restricted classes of pseudoknots. In this paper, we extend the current classification of pseudoknots to capture more complicated pseudoknots, namely the simple non-standard pseudoknots of degree  $k$ . We provide an algorithm to compute the structure with minimum free energy for this type of pseudoknots of degree 4 which covers all known secondary structures of RNAs in this class. Our algorithm runs in  $O(m^4)$  time where  $m$  is the length of the input RNA sequence.

**Keywords:** RNA, Secondary structure prediction, Simple non-standard pseudoknot, Complex pseudoknot.

## 1 Introduction

RNA molecules are known to play important roles in living cells and are involved in many biological processes [5] [8] [13]. The structure of an RNA molecule provides important information about the functions of the molecule. Thus, finding the structure of an RNA molecule is an important problem. Unfortunately, finding or predicting the 3D (or tertiary) structure of an RNA molecule is a complicated and time consuming task. A more promising direction is to predict the secondary structure (that is, which pair of bases of the molecule forms a hydrogen bond) of the molecule with minimum free energy based on its primary sequence. The predicted secondary structure can already help researchers to deduce the functions of the molecule. However, predicting secondary structure of an RNA molecule is not an easy problem and is computationally difficult in the presence of pseudoknots (base pairs crossing each other, secondary structure without pseudoknots are referred as regular structures). In general, the problem is NP-hard [7]. Most of the existing algorithms aim at restricted classes of pseudoknots.

Pseudoknot structures can be classified as follows in increasing order of complexity: H pseudoknots [7], simple pseudoknots [1], standard pseudoknots (see the definition in Section 2), recursive pseudoknots (i.e., pseudoknot/regular structure inside a pseudoknot) [1]. For the definitions of these pseudoknot structures, please refer to the given references.

Rivas and Eddy were among the first who tackled the problem of RNA secondary structure prediction with pseudoknots [10]. They described a dynamic programming

algorithm to solve the problem for simple pseudoknots, certain restricted cases for standard pseudoknots and recursive pseudoknots, as well as some restricted cases in a more complicated class of pseudoknots, simple non-standard pseudoknots, to be defined in Section 2. Their algorithm runs in  $O(m^6)$  time where  $m$  is the length of the input RNA sequence. Their algorithm is still the most powerful algorithm that can handle the largest set of pseudoknot structures including some complicated ones for which none of the existing polynomial-time algorithms can handle.

Lyngso and Pedersen provided a faster algorithm that runs in  $O(m^5)$  time [7], but can only handle H pseudoknots. Later on, Uemura *et al.* gave an improved prediction algorithm for H pseudoknots that runs in  $O(m^4)$  time [12]. The algorithm can, in fact, handle simple pseudoknots and some limited cases for standard pseudoknots and recursive pseudoknots. Their algorithm is based on a tree adjoining grammar (TAG) to model RNA secondary structures that include pseudoknots. However, tree adjoining grammar is not easy to understand.


Akutsu provided another dynamic programming algorithm that runs in  $O(m^4)$  time for predicting simple pseudoknots [1]. This algorithm is much simpler than the algorithm given in [12]. He also extended the algorithm to support recursive simple pseudoknot structures (i.e., there can only be simple pseudoknots/regular structures inside another simple pseudoknot). The algorithm in [1] was implemented and evaluated by Deogun [3]. Recently, Chen *et al.* provided a faster algorithm that runs in  $O(m^5)$  time [2] that handles almost all pseudoknot structures that can be handled by the algorithm in [10]. The ones that cannot be handled by the algorithm in [2] are those in which there are three base pairs that cross one another. Other related work includes [4] [9] in which the pseudoknots considered are also more restricted than those in [10] [2].

**Our Contributions.** In this paper, we consider a more complicated class of pseudoknot structures, called simple non-standard pseudoknots, which extend the current classification of pseudoknots to capture more complicated pseudoknots and include some cases for which three base pairs cross one another. For example, it can include a complex structure (shown in Figure 1) with three base pairs crossing one another that is known to be a topology present in *Escherichia coli*  $\alpha$  mRNA [6] and some other complex structures as shown in Figure 2 listed in [4]. We provide an  $O(m^4)$  time algorithm for predicting simple non-standard pseudoknots of degree 4 which already include all known structures in this class. Note that our algorithm can handle all structures defined in the class of simple non-standard pseudoknots with degree 4 while algorithms in [10] [1] cannot. Our algorithm can be extended for general degree  $k$  with running time  $O(m^k)$ .


## 2 Simple Non-standard Pseudoknots

Let  $A = a_1a_2a_m$  be a length- $m$  RNA sequence with alphabet  $\{A, C, G, U\}$  and  $M$  be a secondary structure of  $A$ .  $M$  can be represented as a set of base pair positions, i.e.,  $M = \{(i, j) | 1 \leq i \leq j \leq m, (a_i, a_j) \text{ is a base pair}\}$ . Let  $M_{x,y} \subseteq M$  be the set of base pairs within the subsequence  $a_xa_{x+1}\dots a_y$ ,  $1 \leq x < y \leq m$ , i.e.,  $M_{x,y} = \{(i, j) \in M | x \leq i < j \leq y\}$ . Note that  $M = M_{1,m}$ . We assume that there is no two base pairs sharing the same position, i.e., for any  $(i_1, j_1), (i_2, j_2) \in M$ ,  $i_1 \neq j_2$ ,  $i_2 \neq j_1$ , and  $i_1 = i_2$  if and only if  $j_1 = j_2$ .

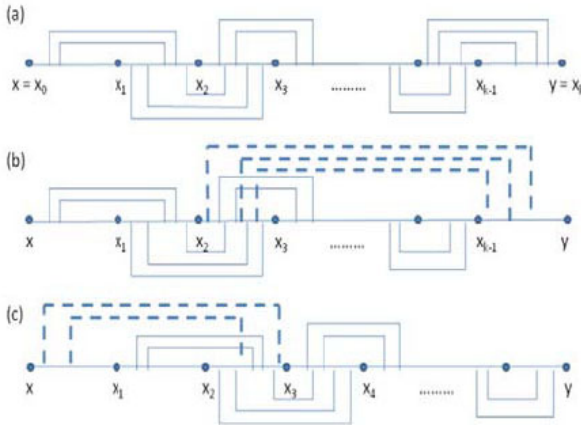




**Fig. 1.** A complex pseudoknot structure present in Esherichia coli  $\alpha$  mRNA [6]. It consist three base pairs that cross one another.



**Fig. 2.** The two complex pseudoknot structures (a) and (b) listed in [11]. The gray box represents a set of base pairs in the regions.



**Fig. 3.** (a) A standard pseudoknot of degree  $k$ . (b) A simple non-standard pseudoknot of degree  $k$  (Type I). (c) A simple non-standard pseudoknot of degree  $k$  (Type II).

Pseudoknots are base pairs that cross each other. For example, let  $(i, j)$  and  $(i', j')$ , where  $i < j$  and  $i' < j'$ , be two base pairs. They form a pseudoknot if  $i < i' < j < j'$  or  $i < i' < j' < j$ .  $M_{x,y}$  is a regular structure if there does not exist pseudoknots. Note that an empty set is also considered as a regular structure.

We now define a standard pseudoknot of degree  $k$  as follows. A structure is a standard pseudoknot of degree  $k$  if the RNA sequence can be divided into  $k$  consecutive regions (see Figure 3(a)) such that base pairs must have end points in adjacent regions and base

pairs that are in the same adjacent regions cannot cross each other. The formal definition is as follows.

$M_{x,y}$  is a standard pseudoknot of degree  $k \geq 3$  if there exists a set of pivot points  $x_1, x_2, \dots, x_{k-1}$  ( $x = x_0 < x_1 < x_2 < \dots < x_{k-1} < x_k = y$ ) that satisfy the following. Let  $M_w(1 \leq w \leq k - 1) = \{(i, j) \in M_{x,y} | x_{w-1} \leq i < x_w \leq j < x_{w+1}\}$ . Note that we allow  $j = x_k$  for  $M_{k-1}$  to resolve the boundary case.

- For each  $(i, j) \in M_{x,y}$ ,  $(i, j) \in M_w$  for some  $1 \leq w \leq k - 1$ .
- $M_w(1 \leq w \leq k - 1)$  is a regular structure.

Note that a standard pseudoknot of degree 3 is simply referred as a simple pseudoknot. Now, we define a simple non-standard pseudoknot that extends the standard pseudoknot to include some structures with three base pairs crossing each other. For a simple non-standard pseudoknot of degree  $k$ , similar to a standard pseudoknot, the RNA sequence can be divided into  $k$  regions with the region at one of the ends (say, the right end) designated as the special region. Base pairs with both end points in the first  $k - 1$  regions have the same requirements as in a standard pseudoknot. And there is an extra group of base pairs that can start in one of the first  $k - 2$  regions and end at the last special region and again these pairs do not cross each other (see Figure 3(b)). See the formal definition below.

$M_{x,y}$  is a simple non-standard pseudoknot of degree  $k \geq 4$  (Type I) if there exist  $x_1, \dots, x_{k-1}$  and  $t$  where  $x = x_0 < x_1 < \dots < x_{k-1} < x_k = y$  and  $1 \leq t \leq k - 2$  that satisfy the following. Let  $M_w(1 \leq w \leq k - 2) = \{(i, j) \in M_{x,y} | x_{w-1} \leq i < x_w \leq j < x_{w+1}\}$ . Let  $X = \{(i, j) \in M_{x,y} | x_{t-1} \leq i < x_t, x_{k-1} \leq j \leq y\}$ .

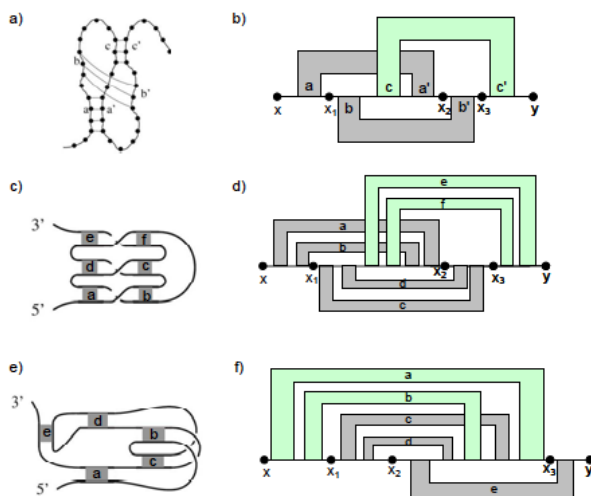
- For each  $(i, j) \in M_{x,y}$ , either  $(i, j) \in M_w(1 \leq w \leq k - 2)$  or  $(i, j) \in X$ .
- $M_w$  and  $X$  is a regular structure.

Type II simple non-standard pseudoknots (see Figure 3(c)) are symmetric to Type I simple nonstandard pseudoknots with the special region on the left end. As shown in Figure 4, the complex structures in Figure 1 and Figure 2 belongs to the simple non-standard pseudoknot structure. For the sake of simplicity, in the rest of the paper, we only consider Type I simple non-standard pseudoknots and simply refer it as simple non-standard pseudoknots. So, we omit the definition of Type II simple non-standard pseudoknots.


### 3 Problem Definition

In the following, we assume that the secondary structure of the given RNA sequence is of simple nonstandard pseudoknot. We now define a free energy model for a standard non-standard pseudoknot structure. We use a similar model as defined in [1] and [2]. The free energy of a structure is the sum of the free energy of all the stems and the free energy of all the unpaired bases which do not belong to any stem.

Roughly speaking, considering the structure of a simple nonstandard pseudoknot, a stem is a pair of maximal regions bounded by two base pairs such that no base pair with one end inside any of the two regions and another end outside the region. A stem is



**Fig. 4.** The complex pseudoknot structure (a) and (c) are type I simple non-standard pseudoknots shown in (b) and (d) respectively. Structure (e) is type II simple non-standard pseudoknot shown in (f). The gray region represents a set of base pairs and the green region represents a set of base pairs in the special region of simple non-standard pseudoknot structure.



**Fig. 5.** Dot lines represent base pairs. Two regions on the RNA sequence covered by the same color are a stem region. Say  $[6, 10] \cup [23, 27]$  represents a stem region. In this example, there are 4 stem regions.

defined formally as two non-overlapping regions  $[a_i a_{i+p}]$  and  $[a_{j-q} \dots a_j]$  such that (i)  $(a_i, a_j)$  and  $(a_{i+p}, a_{j-q})$  both are base pairs; (ii) all base pairs  $(a_r, a_s)$  where  $i \leq r \leq i + p$  and  $j - q \leq s \leq j$  do not cross each other; (iii) there does not exist any base pair such that one end is inside the regions, but another one is outside the regions; (iv) there does not exist any base pair  $(a_r, a_s)$  such that both ends are inside the same region (i.e.  $i \leq a_r < a_s \leq i + p$  or  $j - q \leq a_r < a_s \leq j$ ); (v) the values of  $p$  and  $q$  are maximum. Note that every base pair will belong to a stem of  $p, q \geq 1$ . Figure 5 illustrates the idea of stems inside the simple non-standard pseudoknot structure.

For the stem consists of two size-1 regions, we define the free energy of the stem as  $e_1$ . For the other stems (i.e. the sizes of both regions are at least 2), each of them can be partitioned into loops. The free energy of the stem is defined as the sum of free energy of all the loops inside a stem and the

corresponding energy are as follows. Let the minimum distance between  $a_i$  and  $a_j$  of any base pair  $(a_i, a_j)$  be  $M$ .

- A *stacking pair* represents two consecutive base pairs.  $[a_i \dots a_{i+1}]$  and  $[a_{j-1} \dots a_j]$  where  $i + 1 < j - 1 + M$  is a stacking pair if  $(a_i, a_j)$  and  $(a_{i+1}, a_{j-1})$  are both base pairs. The corresponding energy is defined as  $e_{stack}(a_i, a_j, a_{i+1}, a_{j-1})$ .
- A *left-side bulge* is a loop enclosed by two base pairs  $(a_i, a_j)$  and  $(a_{i'}, a_{j'})$  such that  $i' - i = 1$  and  $j - j' > 1$  and  $i' < j' + M$ . The corresponding energy is defined as  $\alpha_0 + (j - j' - 1)\alpha_1$ .
- Similarly, a *right-side bulge* is a loop enclosed by two base pairs  $(a_i, a_j)$  and  $(a_{i'}, a_{j'})$  such that  $i' - i > 1$  and  $j - j' = 1$  and  $i' < j' + M$ . The corresponding energy is defined as  $\beta_0 + (i' - i - 1)\beta_1$ .
- A *internal loop* is a loop enclosed by two base pairs  $(a_i, a_j)$  and  $(a_{i'}, a_{j'})$  such that  $i' - i > 1$  and  $j - j' > 1$  and  $i' < j' + M$ . The corresponding energy is defined as  $\gamma_0 + ((i' - i - 1) + (j - j' - 1))\gamma_1$ .

For the energy of the unpaired bases which do not belong to any stem, we use a simple function as follows: let  $L$  be the total number of these unpaired bases. The corresponding energy is defined as  $(L)e_0$ . Table 1 summarizes the parameters of the energy model. Our dynamic programming algorithm (which will be described in the next section) is designed to compute the minimum energy according to this energy model. The algorithm can be further extended to consider a more complicated energy model and include more parameters to increase the accuracy of the structure prediction.


**Table 1.** Parameters of the energy model

Parameters	Description
$e_{stack}(a_i, a_j, a_{i+1}, a_{j-1})$	energy of a stacking pair enclosed by $(a_i, a_j)$ and $(a_{i+1}, a_{j-1})$
$\alpha_0$	energy for initiation of a left-side bulge
$\alpha_1(k)$	energy for $k$ unpaired bases inside a left-side bulge
$\beta_0$	energy for initiation of a right-side bulge
$\beta_1(k)$	energy for $k$ unpaired bases inside a right-side bulge
$\gamma_0$	energy for initiation of an internal loop
$\gamma_1(k)$	energy for $k$ unpaired bases inside an internal loop
$e_0(k)$	energy for $k$ unpaired bases which do not belong to any stem
$e_1$	energy for a stem with two size-1 regions

The problem is defined as follows. Given an RNA sequence, compute a secondary structure which is a simple non-standard pseudoknot with minimum free energy.

## 4 Prediction Algorithm

We predict the optimal structure with minimum free energy using a dynamic programming algorithm. The core of our algorithm is based on the concept of a subregion so that we can find the optimal structure recursively. In the following, we first explain the concept of subregion, then we provide the details of the algorithm followed by the time and space complexity analysis.



**Fig. 6.** Subregion of a simple non-standard pseudoknot

### 4.1 Subregion in Simple Non-standard Pseudoknot

Since all known RNAs with simple non-standard pseudoknots are of degree at most 4, in the following, we only consider degree 4 simple non-standard pseudoknots. The subregion is defined in a way such that we do not have base pairs with one end inside the subregion and the other end outside the subregion, thus enabling us to use dynamic programming approach to solve the problem. To make it easy to understand what a subregion is, we redraw the pseudoknot structure as in Figure 6(b). Based on the way we draw the structure, it is easy to see that for the optimal secondary structure, base pairs can be ordered from top to bottom without crossing each other.

We define a subregion using four points  $p, q, r, s$  with  $x \leq p < q < r < s \leq y$ . An example is shown in Figure 6(c) in which the highlighted part is the subregion defined by the four points. Note that when we predict the secondary structure, we do not actually know the locations of  $x_1, x_2, x_3$ . So, we try all possible combination of  $p, q, r, s$  to define subregions. These points are added in Figure 6 to illustrate that there is always a way to define a subregion so that base pairs in the optimal structure will not have one end point inside the subregion while the other end outside the subregion. So, for each subregion we define in the dynamic programming algorithm, we will not have base pairs having one end point inside the subregion and the other end point outside the subregion.

The same definition of subregion cannot be applied when  $t$  is even. Figure 6(d) shows the problem by using the same definition as there can be base pair with one end point inside the subregion and the other end point outside the subregion, thus dynamic programming approach cannot be applied easily. Note that the two base pairs that appear to cross each other in Figure 6(d) is only due to the way we draw it, they do not actually cross each other, so the structure is still a simple non-standard pseudoknot. To solve the

problem, we use a different definition for subregions when  $t$  is even as shown in Figure 6(e). Formally speaking, we define a subregion as follows. Let  $A[x\dots y]$  be an RNA sequence. Let  $v = (p, q, r, s)$  be a quadruple with  $x \leq p < q < r < s \leq y$ . If  $t$  is odd, define the subregion  $R^{odd}(v) = [p\dots q] \cup [r\dots s]$ . For  $t$  is even, we need an additional parameter  $x_3$  and define the subregion  $R^{even}(x_3, v) = [p\dots q] \cup [r\dots x_3 - 1] \cup [s\dots y]$ .

## 4.2 Dynamic Programming Algorithm

We first show how to compute the minimum free energy of the optimal secondary structure for the case of  $R^{odd}$ .

Let  $S_Y^{odd}(p, q, r, s)$  be the minimum free energy of the optimal secondary structure in  $R^{odd}(p, q, r, s)$  where  $Y \in \{L_z, R_z, S_z, D\}$  and  $z \in \{start, pair, Lbulge, Rbulge, internal\}$  is one of the possible cases to consider for having the optimal secondary structure inside  $R^{odd}(p, q, r, s)$ . These cases are explained in the following. Note that according to the definition of simple non-standard pseudoknot (for  $t$  is odd), only  $(p, q)$ ,  $(q, r)$ , and  $(p, s)$  can form a base pair. If the bases  $p$  and  $q$  inside the same stem, then we have one of the following cases: (1)  $(p, q)$  is a base pair which may start a stem, or extend a stem; (2)  $(p, q)$  is not a base pair and belongs to either stacking pair, left-side bulge, right-side bulge or internal loop; The following shows all these cases for different value of  $Y$ .

$L_{start}$  refers to the case that  $(p, q)$  forms a base pair to start a stem;

$L_{pair}$  refers to the case that  $(p, q)$  form a base pair but it is not the start of the stem;

$L_{Lbulge}$  refers to the case that  $(p, q)$  is inside a left-side bulge but does not form a base pair;

$L_{Rbulge}$  refers to the case that  $(p, r)$  is inside a right-side bulge but does not form a base pair;

$L_{internal}$  refers to the case that  $(p, r)$  is inside an internal loop but does not form a base pair;

Similar for  $R_z$  and  $S_z$  where  $z \in \{start, pair, Lbulge, Rbulge, internal\}$ , which refers to the case for  $(q, r)$  and  $(p, s)$  respectively.

$D$  refers to the case that all bases  $x \in \{p, q, r, s\}$  may either (i) contribute to the last base pair in some stems, or (ii) do not form base pairs and do not belong to any stem.

For any two bases  $a_i, a_j$  in the RNA sequence, let  $v(a_i, a_j) = 0$  if  $a_i$  and  $a_j$  can form a base pair, otherwise let  $v(a_i, a_j) = +\infty$ . The following shows how to compute  $S_Y^{odd}(p, q, r, s)$  recursively by considering all possible cases.

**Recurrences.** We first consider the situation when  $p$  and  $q$  are in the same stem.

For the case when  $(p, q)$  is a base pair and starts a stem:

$$S_{L_{start}}^{odd}(p, q, r, s) = v(a_p, a_q) + S_D^{odd}(p + 1, q - 1, r, s)$$

For the case when  $p$  and  $q$  do not form base pair but belong to a left-side bulge:

$$S_{L_{Lbulge}}^{odd}(p, q, r, s) = \min \left\{ \begin{array}{l} S_{L_{start}}^{odd}(p + 1, q, r, s) + \alpha_0 + \alpha_1 \\ S_{L_{Lbulge}}^{odd}(p + 1, q, r, s) + \alpha_1 \end{array} \right.$$

For the case when  $p$  and  $q$  do not form base pair but belong to a right-side bulge:

$$S_{LRbulge}^{odd}(p, q, r, s) = \min \left\{ \begin{array}{l} S_{Lstart}^{odd}(p, q - 1, r, s) + \beta_0 + \beta_1 \\ S_{LRbulge}^{odd}(p, q - 1, r, s) + \beta_1 \end{array} \right.$$

For the case when  $p$  and  $q$  do not form base pair but belong to an internal loop:

$$S_{Linternal}^{odd}(p, q, r, s) = \min \left\{ \begin{array}{l} // \text{ for internal loop, each side has at least 1 unpaired base.} \\ S_{Lstart}^{odd}(p + 1, q - 1, r, s) + \gamma_0 + 2\gamma_1 \\ S_{Linternal}^{odd}(p + 1, q, r, s) + \gamma_1 \\ S_{Linternal}^{odd}(p, q - 1, r, s) + \gamma_1 \end{array} \right.$$

For the case when  $p$  and  $q$  is a base pair and inside a stem (i.e. not starting a stem):

$$S_{Lpair}^{odd}(p, q, r, s) = v(a_p, a_q) + \min \left\{ \begin{array}{l} S_{Lstart}^{odd}(p + 1, q - 1, r, s) + e_{stack}(a_p, a_q, a_{p+1}, a_{q-1}) \\ S_{Lpair}^{odd}(p + 1, q - 1, r, s) + e_{stack}(a_p, a_q, a_{p+1}, a_{q-1}) \\ S_{Lbulge}^{odd}(p + 1, q - 1, r, s) \\ S_{Rbulge}^{odd}(p + 1, q - 1, r, s) \\ S_{Linternal}^{odd}(p + 1, q - 1, r, s) \end{array} \right.$$

It is similar for the other situations when  $q, r$  are in the same stem, and  $p, s$  are in the same stem (i.e.  $S_{Rz}^{odd}$  and  $S_{Sz}^{odd}$  where  $z \in \{start, pair, Lbulge, Rbulge, internal\}$ ). For the cases when all bases  $x \in \{p, q, r, s\}$  may either (i) contribute to the last base pair in some stems, or (ii) do not form base pairs and do not belong to any stem:

$$S_D^{odd}(p, q, r, s) = \min\{CLOSE, BETWEEN\}$$

where *CLOSE* refers to the closing of a stem and *BETWEEN* refers to the case that it is not inside a stem.

$$CLOSE = \min \left\{ \begin{array}{l} S_{Lstart}^{odd}(p, q, r, s) + e_1 \\ S_{Lpair}^{odd}(p, q, r, s) \\ S_{Rstart}^{odd}(p, q, r, s) + e_1 \\ S_{Rpair}^{odd}(p, q, r, s) \\ S_{Sstart}^{odd}(p, q, r, s) + e_1 \\ S_{Spair}^{odd}(p, q, r, s) \end{array} \right.$$

$$BETWEEN = \min \left\{ \begin{array}{l} S_D^{odd}(p + 1, q, r, s) + e_0 \\ S_D^{odd}(p, q - 1, r, s) + e_0 \\ S_D^{odd}(p, q, r + 1, s) + e_0 \\ S_D^{odd}(p, q, r, s - 1) + e_0 \end{array} \right.$$

The recurrences for computing the minimum free energy of the optimal secondary structure for  $R^{even}(p, q, r, s, x_3)$  will be similar, but note the additional parameter required for this case due to the slightly different definition of subregions. Let  $S_Y^{even}(p, q,$

$r, s, x_3$ ) be the minimum free energy of the optimal secondary structure in  $R^{even}(p, q, r, s, x_3)$  where  $Y \in \{L_z, R_z, S_z, D\}$  and  $z \in \{start, pair, Lbulge, Rbulge, internal\}$  is one of the possible cases. The definitions of  $Y$  is the same as that for  $S_Y^{odd}(p, q, r, s)$  except with  $(q, s)$  replacing  $(p, s)$  as for this case  $(p, s)$  will not form a base pair, but  $(q, s)$  can form a base pair.

The minimum free energy of the optimal structure for the whole RNA sequence is the minimum value of  $\{\min_x \{S_D^{odd}(1, x, x+1, n)\}, \min_{y < x_3} \{S_D^{even}(1, y, y+1, x_3, x_3)\}\}$ .

From the real data, the distance between  $x_3$  and the end of the sequence is usually bounded by a small constant, so we assume that the number of different  $x_3$  values we need to consider is only a small constant. The time complexity of the above algorithm is  $O(m^4)$ . The memory complexity of the algorithm is also  $O(m^4)$ .

## 5 Conclusions

In this paper, we consider a new class of pseudoknots which include more complicated structures that none of the existing algorithms can handle. We then provide an  $O(m^4)$  time algorithm for predicting these a structure of degree 4 with minimum free energy which already covers all known secondary structures of this class in existing databases. We implemented our algorithm and the running time is reasonable, which takes about 70sec for a RNA of length about 100 and about 3 times faster than the one in [10]. We will evaluate the accuracy of the predicted structures once we can locate a set of appropriate parameters for the energy model. In fact, there are not many known RNAs with simple non-standard pseudoknots. One of the reasons may be due to the limitation of existing computational prediction tools. With our algorithm, we may be able to predict more RNAs with such a structure for follow-up verification. Although there are no other more complicated known pseudoknot structures, there is a high chance that there exist novel RNAs with more complicated structures, so designing efficient prediction algorithms for more complicated pseudoknot structures remains an important open problem.

## References

1. Akutsu, T.: Dynamic programming algorithms for rna secondary structure prediction with pseudoknots. *Discrete Applied Mathematics* 104, 45–62 (2000)
2. Chen, H., Condon, A., Jabbari, H.: An  $o(n^5)$  algorithm for mfe prediction of kissing hairpins and 4- chains in nucleic acids. *Journal of Computational Biology* 16(6), 803–815 (2009)
3. Deogun, J., Donis, R., Komina, O., Ma, F.: Rna secondary structure prediction with simple pseudoknots. In: *Proceedings of the second conference on Asia- Pacific Bioinformatics Conference (APBC 2004)*, pp. 239–246 (2004)
4. Dirks, R., Pierce, N.: A partition function algorithm for nucleic acid secondary structure including pseudoknots. *Journal of Comput. Chem.* 24(13), 1664–1677 (2003)
5. Frank, D., Pace, N.: Ribonuclease p: unity and diversity in a trna processing ribozyme. *Annu. Rev. Biochem.* 67, 153–180 (1998)
6. Gluick, T., Draper, D.: Thermodynamics of folding a pseudoknotted mrna fragment. *Journal of Molecular Biology* 241, 246–262 (1994)



7. Lyngso, R., Pedersen, C.: A dynamic programming algorithm for rna structure prediction including pseudoknots. In: Proc. of the Fourth Annual International Conferences on Computational Molecular Biology (RECOMB 2000). ACM Press, New York (2000)
8. Nguyen, V., Kiss, T., Michels, A., Bensaude, O.: 7sk small nuclear rna binds to and inhibits the activity of cdk9/cyclin t complexes. *Nature* 414, 322–325 (2001)
9. Reeder, J., Giegerich, R.: Design, implementation and evaluation of a practical pseudoknot folding algorithm based on thermodynamics. *BMC Bioinformatics* 5, 104 (2004)
10. Rivas, E., Eddy, S.: A dynamic programming algorithm for rna structure prediction including pseudoknots. *Journal of Molecular Biology* 285(5), 2053–2068 (1999)
11. Roland, E.: Pseudoknots in rna secondary structures: representation, enumeration, and prevalence. *Journal of Computational Biology* 13(6), 1197–1213 (2006)
12. Uemura, Y., Hasegawa, A., Kobayashi, S., Yokomori, T.: Tree adjoining grammars for rna structure prediction. *Theoretical Computer Science* 210, 277–303 (1999)
13. Yang, Z., Zhu, Q., Luo, K., Zhou, Q.: The 7sk small nuclear rna inhibits the cdk9/cyclin t1 kinase to control transcription. *Nature* 414, 317–322 (2001)

# Intra- and Inter-Molecular Coevolution: The Case of HIV1 Protease and Reverse Transcriptase

Patrick Boba, Philipp Weil, Franziska Hoffgaard, and Kay Hamacher

Bioinformatics & Theoretical Biology Group, Technische Universität Darmstadt, Germany  
bobapat@googlemail.com,  
{weil, hoffgaard, hamacher}@bio.tu-darmstadt.de  
<http://www.kay-hamacher.de>

**Abstract.** The stability, fold, and the function of proteins need to be maintained throughout the evolution of these molecules – inducing a selective pressure, that can be revealed in sequence data sets. The conservation of structure and function implies coevolution of amino acids within the protein. To understand such selective pressure in the evolution of the human immunodeficiency virus (HIV), we apply information theoretical measures to the two most important enzymes for the progression of viral infection: the reverse transcriptase and the protease. We computed the mutual information to derive insight into the selective pressure acting locally and globally on the enzymes. We found intra- and inter-protein co-evolution of residues in these enzymes and annotate important structural-evolutionary correlations. We discuss a signal indicating a potential co-evolution between the protease and the reverse transcriptase.

**Keywords:** Mutual information, HIV, Molecular evolution, Coevolution, Sequence analysis.

## 1 Introduction

The acquired immunodeficiency syndrome (AIDS) is induced by the human immunodeficiency virus (HIV). Its viral replication cycle depends on the virus own protease and several other enzymes such as the reverse transcriptase. Currently the anti-HIV drugs target these two enzymes to prevent the maturation of new virions [12].

Neutral evolution and drug resistance development have been under investigation for a long time: 1) the high mutation rate of HIV makes the virus an interesting evolutionary object in itself as it performs a large-scale mutagenesis study [3]; 2) a deeper knowledge of potential evolutionary barriers might lead to new therapeutics besides the HAART-procedure [4].

The theoretical understanding of the viral evolution has greatly improved over the recent years [5,6,7,8], even the biophysical annotation based on *in silico* models of the molecular dynamics is under way [9].

At the same time the wealth of information on HIV - in particular the large data sets of sequences - prompt for a deeper analysis on the sequence level alone. Here we leverage an available data set of 45,161 mutant sequences of the HIV-1 protease (PR) and reverse transcriptase (RT).

## 2 Materials and Methods

### 2.1 Sequence Data

The 45,161 positive selection mutant sequences have been collected by the Lee lab [10,11] and were made available on the net. The data set contains the genomic, nucleotide sequences from treated and untreated patients under various drug regimes. The individual entries are, however, not annotated by the drug treatment regime of the particular patient. We therefore find in this data set the diverse evolutionary dynamics, including effects such as neutral drift, drug resistance development, and other selective pressures on the two enzymes. Wherever a codon could not be mapped unequivocally to an amino acid we used a wild-card character, treating these cases similar to gaps.

### 2.2 Information Theoretical Measures on (Co-)evolution

The evolution of an amino acid at a position  $i$  means a change in the symbols  $S^i$  over time within a set of acceptable values  $\mathcal{S}$ . One way to quantify the information content of such collections of symbols is the Shannon entropy [12]

$$H^i := - \sum_{S^i \in \mathcal{S}} p(S^i) \cdot \log_2(p(S^i)) \quad (1)$$

where  $p(S^i)$  is the probability of the occurrence of the symbol  $S^i$  within the empirical or theoretical data set under investigation. For empirical data sets one usually sets this probability to the frequency of the symbol within the data set. Positions (in e.g. a sequence alignment) with high entropy are then amino acids with high variability during evolutionary times. Our choice of  $\mathcal{S}$  comprised the 20 standard amino acids and the above mentioned wild-card character.

The *correlated* change in the amino acid composition within a molecule or between molecules is now based on empirical found two-point probabilities  $p(S^i, S^j)$  for the co-evolution of positions  $i$  and  $j$ . We can define the Mutual Information (MI) between these positions as a relative entropy as follows [13]:


$$\begin{aligned} \text{MI}^{i,j} &:= \sum_{S^i, S^j \in \mathcal{S}} p(S^i, S^j) \cdot \log_2 \left( \frac{p(S^i, S^j)}{p(S^i) \cdot p(S^j)} \right) \\ &= H^{i,j} - H^i - H^j \end{aligned} \quad (2)$$

The value of the MI gives the amount of information that one position  $i$  conveys about the other position  $j$ . The MI can be derived from the Kullback-Leibler divergence as a relative entropy, which has - besides sequence based approaches - also attracted attention as a measure in *in silico* drug design and phylogenies [14,15,16,17]. We discuss the statistical relevance and possible normalization procedures elsewhere [18].

## 3 Results

### 3.1 Comparing Intra- and Inter-Protein Co-evolution of Residues

In figure 1 we show distributions of the naked MI-values from our study on both, the HIV-1 PR and the HIV-1 RT, as well as the inter-MI for a potential co-evolution of residues in these enzymes.



**Fig. 1.** Comparison of the MI-values for the intra-protein co-evolution within the HIV-1 Protease (black) and the HIV-1 Reverse Transcriptase (blue). We compare to the *inter*-MI for the co-evolution between residues of the HIV-1 Protease on the one hand and the HIV-1 Reverse Transcriptase on the other (red).

We observe similarity of MI results for the *intra*-co-evolution within the individual, isolated enzymes. Obviously the evolutionary dynamics gave rise to the same overall “mutual information picture”.


The dissimilarity of the MI-distributions for the RT/PR and the one for the inter-molecular MI comes as no surprise: *within* a molecule the evolutionary pressure on the co-evolving dynamics of amino acids can be regarded as quite different in the evolutionary dynamics between residues in *different* molecules, despite potential protein-protein-interactions or other implicit interdependencies resulting from cell biological effects or drug combinations.

Although the RT consists of four domains – namely the finger, palm, thumb, connection domains – the potential for co-evolution between sites distributed over the four domains runs approximately in parallel to the scenario of the protease, both of which - in turn - are constructed as a molecular phenotype in form of homodimers.

Our results on the HIV-1 PR are in agreement with an early study on this particular molecule [9]. The raw data is available from our web-site [19] for future analysis.

To analyze our MI results further and to overlay these with structural knowledge, we went on with a spectral decomposition of the MI matrices for the HIV-1 PR and HIV-1 RT. For the *inter*-MI values, that would indicate potential co-evolution between residues of different molecules, the MI matrix is, however, non-quadratic as the protein lengths are in general different. We therefore applied a singular value decomposition [20] to obtain a pseudo-spectral decomposition with respect to the singular values of the *inter*-MI matrix. We found a dominating eigenvalues/singular value [17]. If we now overlay the corresponding eigenvectors onto the structures of the molecules, we can immediately connect structural and evolutionary information. This is done in figures 2 and 3.


Figures 2 and 3 both show high mutual information for secondary structure elements. In particular the  $\beta$ -sheet in the PR needs to be maintained as a structural basis of the fold of this protein. This is achieved by co-evolution of the residues within this element. In the RT the  $\beta$ -sheet close to the reactive center, as well as the  $\alpha$ -helices forming the “finger” of the RT are structurally maintained by co-evolving the residues without giving them the freedom to independently mutate.



**Fig. 2.** a) sequence entropy of the HIV-1 PR as in eq. 1; b) absolute values of the entries of the 1<sup>st</sup> eigenvector for the MI matrix of HIV-1 PR; c) absolute values of the entries of the 2<sup>nd</sup> eigenvector. We rescaled all values so that blue=maximum value, red=minimum value.

In figure 3 we show in addition relevant residues as small spheres. The colors indicate: yellow=three catalytic aspartic acids; green=residues that enhance the excision reaction; red=non-nucleoside inhibitor binding pocket; black=residues involved in NRTI-resistance. This mapping was done in accordance with previous work [21].

We note in passing that the level of conservation needs to be taken into account: an absolutely conserved position shows a MI of zero, always, as the knowledge about the identity of a residue here does not convey any information on any other position. We therefore decided to also display the sequence entropy of equation 1 as a measure




**Fig. 3.** I) a) sequence entropy of the HIV-1 RT as in eq. 1; b) absolute values of the entries of the 1<sup>st</sup> eigenvector for the MI of HIV-1 RT. II) viewed from orthogonal projection & rescaled as in fig. 2. The black part is the C-terminus for which we had only insufficient statistics; we omitted it from our analysis. The small spheres indicate functional sites as discussed in the text. Domains are indicated as follows: light blue="fingers", green="thumb", gray="palm".

of local sequence conservation in the figures 2 and 3. We return to this issue in the discussion section.

Furthermore one can see the “correlation” of low sequence entropy and therefore the necessarily low mutual information in figure 4. In this figure we motivate the classification of an amino acid by its two evolutionary/co-evolutionary measures, that is the sequence variability as expressed by the entropy of eq. 1 and the contribution to mutual information correlation expressed by the respective entry in the leading eigenvector.

Class I (low sequence entropy, high MI) must be empty. Classes II and IV are the important ones for evolution: amino acids found in class IV are subject to extensive selective pressure to maintain their identity (low sequence entropy, thus small sequence variability). Evolution acts here *locally* to force *sequence stabilization*. Amino Acids in class II on the other hand can vary extensively (large sequence entropy), but at the same time convey information about other amino acids, thus show correlation with other sites within the protein. Accordingly the high MI reflects a selective pressure to maintain not a particular amino acid character, but instead to maintain some “interaction”. This “interaction” might be a physical, direct interaction such as steric repulsion or charge interactions, but could also reflect, e.g. folding properties of a monomer or recognition capabilities in protein-protein-binding mechanisms. One might hypothesize about the origin of this correlation or “interaction”, but a high MI indicates always a selective pressure to connect residues.

Class III on the other hand consists of those amino acids, which are highly variable (high sequence entropy), but at the same time show low dependence on other sites within the protein, and thus low connection via MI to these other positions.



**Fig. 4.** a) The correlation of MI contributions as found in the first eigenvector of the MI-matrix of PR vs. the respective sequence entropy of the amino acids in the PR. We applied an intuition driven classification scheme to decompose the results into four classes, numbered by Roman letters and illustrated by the red and green line. b) same as a), but for the RT.

### 3.2 Co-evolution between Residues in the PR and the RT?

In figure 1 we have seen that the *inter*-MI is some two orders of magnitude smaller than the overall MI of the *intra*-MI. Combining this insight with the typical values found

in MI studies [17] we nevertheless find a basal co-evolution between the two enzymes under investigation.

Although such a co-evolution is counter-intuitive at first sight, there might be some small or even unknown interdependencies between the two molecules. For example one effect might be due to the packing of RNA in the viral capsid and the genes coding for the enzymes are in close vicinity of the RNA. Such packing is highly susceptible to local charges and balances thereof - probably leading to long-range correlations along the genomic sequence. Additional potential effects are discussed in the discussion section [4].

In figure 5 we show our results for the pseudo-spectral reconstruction of the *inter*-MI between the PR and the RT.

## 4 Discussion and Summary

In this paper we have analyzed two of the most important enzymes for the progression of viral infection by the human immunodeficiency virus (HIV-1 protease and the HIV-1 reverse transcriptase) in an information theoretical setting to investigate evolutionary dynamics and extract positions under exceptional selective pressure.

A first insight is possible by looking solely at the sequence variability, which reveals selective pressure to maintain *local* properties within the molecule - local is meant here in the sense of an individual position. Sites of enzymatic action are prone examples of such findings.

Nevertheless molecular evolution provides for an extended selective pressure, which we label as *non-local* as it involves several amino acids at the same time. Despite individual amino acids being variable, pairs of residues are connected or correlated. This is revealed by the mutual information they carry.

We have shown that our sampling statistics is sufficient and a standard normalization procedure usually applied is not necessary in our case - due to large sample size and absence of gaps in aligned sequences.


A particularly interesting result is the high sequence variability in the  $\beta$ -sheets of the PR, as shown in figure 2 a). At the same time we find these residues also to be relevant for the high MI (parts b & c of the same figure). This was recently discussed and annotated in a biophysical simulation setting [9].

At the same time, we find one residue (I54 in the wild-type) in the flaps to be highly variable and well correlated to other parts of the PR, see the blue residue in the upper strand of the  $\beta$ -sheet forming the flaps in figure 2.

Interestingly in figure 2 the dimerization interface of the PR in the lower part of the molecule shows over a larger range high sequence variability as well as large contributions to the mutual information. This indicates HIV-1's ability to vary the composition of the binding interface to dimerize the PR-monomers to become the PR-homodimer. Obviously maintaining recognition capabilities for binding is of paramount importance for the virus, revealing itself in the high MI.

The implications of relating sequence variability and mutual information can be seen in figure 4. An intuitive classification scheme can be justified on grounds of selective pressure induced by the ongoing evolution of these molecular phenotypes and divided in accordance with this classification procedure.





**Fig. 5.** Absolute values of the entries of the 1<sup>st</sup> left- and right-singular vector for a) the HIV-1 PR and b)+c) for the HIV-1 RT. Again we rescaled all values so that blue=maximum value, red=minimum value, and show the non-analyzed parts of the RT in black.

**Table 1.** The most pronounced members of the classes as introduced in figure 4 for both enzymes. For class I no points exist that fulfill the particular requirement. The enumeration is in accordance with p66 monomer of the RT dimer. We used the numbering convention of [22,21] for RT.

<u>HIV-1 PR</u>				<u>HIV-1 RT</u>			
I	II	III	IV	I	II	III	IV
	63	41	49		334	102	349
	10	69	78		335	49	348
	71	57	29		329	108	347
	12	60	56		333	106	344
none	20	70	28	none	339	249	346
	7	61	27		338	48	343
	90	16	86		324	165	345
	82	39	52		322	100	342
	46	67	98		326	90	341
	54	92	51		311	4	152

In table 1 we extracted the most pronounced residues under these classification schemes - the ones that correspond the most to the three existing classes. To this end we have chosen visually those residues most distant from the intersection of the red and green lines in figure 4.

We find in table 1 the L10 and M46 for the protease to be of class II. Correlated mutations in these positions are known to reduce binding of well-known protease inhibitors, such as JE-2147 by an order of magnitude or even more [23,24]. This makes the acquisition of mutations relatively easy: these amino acids are not to be preserved, they only need to maintain “interactions” or correlations, thus opening the path to change the sequence locally in a correlated fashion to reduce drug efficiency while maintaining the structure, function, and thus the infectious outcome of the protease.

In class IV we found some of the amino acids building the flaps of the PR (res. no. 52-58 are usually labeled to be part of the flaps). As is known from extensive simulations [25] the flaps need to be most flexible to embrace the substrate of the PR. This - as indicated by our findings - is achieved evolutionary to strictly conserve the overall sequence composition of the flaps.

For the RT we find in table 1 the class III very interesting: residues of the binding pocket for the non-nucleoside inhibitors are to be found here. Class III contains, however, those positions that vary a lot, but do not show high correlation to other positions in the molecule. This implies that the amino acids binding the inhibitor can more or less freely mutate, because they are not correlated to other positions and thus there is not need for correlated mutations, which turned out to be necessary for the resistance development of the protease (see above).

We found some indications for a potential co-evolution between the PR and the RT. We can think of three reasons to this end:

The weak co-evolution between the proteins might be – as speculated in the results section – induced by implicit interactions of the coding genes during packing of the viral RNA into the capsid. Obviously charge distributions play a prominent role during these events and that might correlate (slightly) nucleotides and therefore also the coded

amino acids. It is, however, reasonable to assume this effect to be distributed all over the proteins and not localized on particular residues.

Another selective pressure on both proteins is collectively induced by application of protease and reverse transcriptase inhibitors at the same time or in temporal proximity, as in e.g. HAART treatment [4], combining both types of inhibitors, for example lopinavir, ritonavir, tenofovir and emtricitabine. We note that we at least in the RT structure no contribution from portions of the complex that bind RT inhibitors can be observed, making this explanation less likely.

And finally one cannot completely neglect the possibility of functional protein-protein-interactions between the RT and the PR. Although there are currently no indications to this effect and we doubt that they exist, we mention this possibility for the sake of completeness here.

**Acknowledgements.** KH gratefully acknowledges financial support by the Fonds der Chemischen Industrie through the program Sachkostenzuschuß für den Hochschul-lehrernachwuchs.

Molecular Structures were visualized by VMD [26,27]. VMD was developed by the Theoretical and Computational Biophysics Group in the Beckman Institute for Advanced Science and Technology at the University of Illinois at Urbana-Champaign.

## References

1. Tsygankov, A.Y.: Current developments in anti-HIV/AIDS gene therapy. *Curr. Opin. Investig Drugs* 10(2), 137–149 (2009)
2. Wlodawer, A., Erickson, J.: Structure-based inhibitors of HIV-1 protease. *Annu. Rev. Biochem.* 62(1), 543–585 (1993)
3. Perelson, A.S., Neumann, A.U., Markowitz, M., Leonard, J., Ho, D.: HIV-1 dynamics in vivo: virion clearance rate, infected cell life-span, and viral generation time. *Science* 271, 1582–1586 (1996)
4. Richman, D., Margolis, D., Delaney, M., Greene, W.C., Hazuda, D., Pomerantz, R.J.: The challenge of finding a cure for HIV infection. *Science* 323, 1304–1307 (2009)
5. Rong, L., Gilchrist, M.A., Feng, Z., Perelson, A.S.: Modeling within-host HIV-1 dynamics and the evolution of drug resistance: Trade-offs between viral enzyme function and drug susceptibility. *J. Theo. Biol.* 247, 804–818 (2007)
6. Chen, L., Lee, C.: Distinguishing HIV-1 drug resistance, accessory, and viral fitness mutations using conditional selection pressure analysis of treated versus untreated patient samples. *Biology Direct* 1(1), 14 (2006)
7. Trylska, J., Tozzini, V., Chang, C., McCammon, J.A.: HIV-1 protease substrate binding and product release pathways explored with coarse-grained molecular dynamics. *Biophys. J.* 92, 4179–4187 (2007)
8. Hamacher, K., McCammon, J.A.: Computing the amino acid specificity of fluctuations in biomolecular systems. *J. Chem. Theory Comput.* 2(3), 873–878 (2006)
9. Hamacher, K.: Relating sequence evolution of HIV1-protease to its underlying molecular mechanics. *Gene* 422, 30–36 (2008)
10. Pan, C., Kim, J., Chen, L., Wang, Q., Lee, C.: The hiv positive selection mutation database. *Nuc. Acids Res.* 35(1), D371–D375 (2007)

11. Chen, L., Perlina, A., Lee, C.J.: Positive Selection Detection in 40,000 Human Immunodeficiency Virus (HIV) Type 1 Sequences Automatically Identifies Drug Resistance and Positive Fitness Mutations in HIV Protease and Reverse Transcriptase. *J. Virol.* 78(7), 3722–3732 (2004)
12. Shannon, C.E.: Prediction and entropy of printed english. *The Bell System Technical Journal* 30, 50–64 (1951)
13. Lund, O., Nielsen, M., Lundegaard, C., Brunak, C.K.S.: *Immunological Bioinformatics*. MIT Press, Cambridge (2005)
14. Hamacher, K.: Information theoretical measures to analyze trajectories in rational molecular design. *J. Comp. Chem.* 28(16), 2576–2580 (2007)
15. Hamacher, K.: Protein domain phylogenies - information theory and evolutionary dynamics. In: Fred, A., Filipe, J., Gamboa, H. (eds.) *BIOINFORMATICS 2010*, pp. 114–122 (2010)
16. Pape, S., Hoffgaard, F., Hamacher, K.: Distance-dependent classification of amino acids by information theory. *Proteins: Structure, Function, and Bioinformatics* 78, 2322–2328 (2010)
17. Boba, P., Weil, P., Hoffgaard, F., Hamacher, K.: Co-evolution in HIV enzymes. In: Fred, A., Filipe, J., Gamboa, H. (eds.) *BIOINFORMATICS 2010*, pp. 39–47 (2010)
18. Weil, P., Hoffgaard, F., Hamacher, K.: Estimating sufficient statistics in co-evolutionary analysis by mutual information. *Computational Biology and Chemistry* 33(6), 440–444 (2009)
19. Boba, P., Hamacher, K. (2009), <http://bioserver.bio.tu-darmstadt.de/HIV>
20. Press, W.H., et al.: *Numerical Recipes in C*. Cambridge University Press, Cambridge (1995)
21. Sarafianos, S.G., Das, K., Hughes, S.H., Arnold, E.: Taking aim at a moving target: designing drugs to inhibit drug-resistant hiv-1 reverse transcriptases. *Current Opinion in Structural Biology* 14(6), 716–730 (2004)
22. Prajapati, D.G., Ramajayam, R., Yadav, M.R., Giridhar, R.: The search for potent, small molecule nrtis: A review. *Bioorganic & Medicinal Chemistry* 17(16), 5744–5762 (2009)
23. Yoshimura, K., Kato, R., Yusa, K., Kavlick, M.F., Maroun, V., Nguyen, A., Mimoto, T., Ueno, T., Shintani, M., Falloon, J., Masur, H., Hayashi, H., Erickson, J., Mitsuya, H.: JE-2147: A dipeptide protease inhibitor (PI) that potently inhibits multi-PI-resistant HIV-1. *Proc. Natl. Acad. Sci.* 96, 8675–8680 (1999)
24. Reiling, K., Endres, N., Dauber, D., Craik, C., Stroud, R.: Anisotropic dynamics of the JE-2147-HIV protease complex: Drug resistance and thermodynamic binding mode examined in a 1.09 Å structure. *Biochemistry* 41, 4582 (2002)
25. Perryman, A.L., Lin, J.H., McCammon, J.A.: Restrained molecular dynamics simulations of hiv-1 protease: The first step in validating a new target for drug design. *Biopolymers* 82(3), 272–284 (2006)
26. Stone, J.: *An Efficient Library for Parallel Ray Tracing and Animation*. Master's thesis, Computer Science Department, University of Missouri-Rolla (April 1998)
27. Humphrey, W., Dalke, A., Schulten, K.: VMD – Visual Molecular Dynamics. *Journal of Molecular Graphics* 14, 33–38 (1996)

# Structured Output Prediction of Novel Enzyme Function with Reaction Kernels

Katja Astikainen<sup>1</sup>, Liisa Holm<sup>2</sup>, Esa Pitkänen<sup>1</sup>, Sandor Szedmak<sup>3</sup>, and Juho Rousu<sup>1</sup>

<sup>1</sup> Department of Computer Science, University of Helsinki, PO Box 68, Helsinki, Finland

<sup>2</sup> Institute of Biotechnology and Department of Biological Sciences,  
University of Helsinki, P.O. Box 56, Helsinki, Finland

<sup>3</sup> Electronics and Computer Science, University of Southampton, SO17 1BJ,  
Southampton, U.K.

juho.rousu@cs.helsinki.fi

**Abstract.** Enzyme function prediction is an important problem in post-genomic bioinformatics, needed for reconstruction of metabolic networks of organisms. Currently there are two general methods for solving the problem: annotation transfer from a similar annotated protein, and machine learning approaches that treat the problem as classification against a fixed taxonomy, such as Gene Ontology or the EC hierarchy. These methods are suitable in cases where the function of the new protein is indeed previously characterized and included in the taxonomy. However, given a new function that is not previously described, these approaches are not of significant assistance to the human expert. The goal of this paper is to bring forward structured output learning approaches for the case where the exactly correct function of the enzyme to be annotated may not be contained in the training set. Our approach hinges on fine-grained representation of the enzyme function via the so called reaction kernels that allow interpolation and extrapolation in the output (reaction) space. A kernel-based structured output prediction model is used to predict enzymatic reactions from sequence motifs. We bring forward several choices for constructing reaction kernels and experiment with them in the remote homology case where the functions in the test set have not been seen in the training phase.

## 1 Introduction

Enzymes are the workhorses of living cells, producing energy and building blocks for cell growth as well as participating in maintaining and regulation of the metabolic states of the cells. Reliable assignment of enzyme function, that is, the biochemical reactions catalyzed by the enzymes, is a prerequisite of high-quality metabolic reconstruction [14,10], pathway analysis [13] and metabolic flux estimation [16].

In literature, the enzyme function prediction problem comes in two general formulations: annotation transfer or classification by machine learning. In the first approach, given an unannotated protein, a similar annotated protein with experimentally verified function is searched for in databases, and the annotation is transferred to the new protein. In the second approach, a model is trained to classify the new protein into one of the predefined functional classes such as four-level hierarchical EC classification of enzymatic functions.

The success of the above approaches depends on the set of previously characterized and catalogued enzymatic functions. If the new protein belongs to the existing function classes, annotation transfer or classification learning may work. If the new protein, however, possesses a function that is not pre-existing, correct function cannot be predicted even in principle.

Given the diversity of the tree of life, it is likely that completely new functions are encountered as sequencing and annotation efforts widen. Tools, which can give accurate predictions of what the new functions might be, could expedite these efforts. In this paper, we develop a structured output prediction approach that, to our knowledge is the first enzyme function prediction tool to possess the capability of prediction previously unseen functions. The key component of our method is the representation of enzyme function in fine-grained fashion with the so called reaction kernels, that allow interpolation and extrapolation in the space of enzymatic function.

The organization of the paper is the following. In section 2 we briefly describe main existing approaches in enzyme function prediction. In section 3 we review structured output prediction approaches, in particular Kernel Density Estimation and Maximum Margin Regression which are applied in the subsequent sections. In section 4, we describe representations for structured output prediction of enzyme function. We put forward two reaction kernel variants that allow us to interpolate and extrapolate in the space of enzymatic reactions. Section 5 describes experiments validating our approach. Section 6 discusses the relative merits of the current and competing methods, and outlines directions for future work.

## 2 Enzyme Function Prediction

Protein function prediction is recognized as one of the key problems in bioinformatics, and hence there is a large number of approaches to tackle this problem. Most enzyme function prediction methods are instantiations of the more general protein function prediction problem. Here we give a brief overview of protein function prediction approaches. For more information, we refer the interested reader to the recent survey of [15].

### 2.1 Annotation Transfer Approaches

The most widely used function prediction approach is still annotation transfer based on sequence similarity: given an unannotated protein, using a sequence comparison tool such as BLAST, search for an annotated sequence homolog with an experimentally verified function, and transfer the annotation to the new protein. This approach has well-known pitfalls: sequence similarity does not equate to homology, function is typically determined by a small group of residues whose contribution in the overall similarity may fail to be detected, and the danger of the propagation of the annotation errors.

Sequence motifs or signatures are used to overcome shortcomings of overall sequence similarity. As the protein function is typically dependent on a small region of the sequence (e.g. for enzymes the residues forming the active center), a significant amount of research has been conducted to derive sequence motifs that are predictive

of the function [15]. In this paper, we apply the Global Trace Graph [9] features that can be interpreted as predicted conserved residues. The GTG features are derived from a global alignment of all known protein sequences. In this alignment, GTG features correspond to residues that align consistently within a group of proteins.

Information about the 3D structure is known to be a powerful aid in function prediction, due to the fact that it is ultimately the three-dimensional structure that determines the protein function. Structural similarity of two proteins may indicate common evolutionary origin even in the absence of significant sequence similarity. Numerous structural alignment methods (e.g. [11]) have been developed to make use of the 3D structures. Structural motifs are an analogous concept to sequence motifs: a local constellation of residues in the active center of an enzyme may be highly predictive of the function. In this paper, we do not apply 3D information, but leave this as future work.

## 2.2 Machine Learning Approaches

Machine learning methods are potentially useful in cases where the new protein does not possess significant sequence (or structure) similarity to existing proteins. Given large enough data, machine learning methods are able to distill non-trivial associations between the input features and the function.

In the machine learning setting, enzyme function prediction has been generally defined as a classification problem. The works by Lanckriet et al. [12] and Borgwardt et al. [4] use kernel methods to predict the main categories in MIPS and EC taxonomies, respectively.

Other works aim to predict the membership in the whole taxonomy. These include the work by Clare and King [5] who use decision trees to predict the membership in the MIPS taxonomy. Barutcuoglu et al. [2] combine Bayesian networks with a hierarchy of support vector machines to predict Gene Ontology classification. Blockeel et al. [3] use multilabel decision tree approaches to functional class classification according to the MIPS FunCat taxonomy.

Structured output approaches (see below) for hierarchical multilabel classification (c.f. [17]) have been applied to enzyme function prediction by Astikainen et al. [1] and Sokolov and Ben-hur [19]. In this paper, we take the hierarchical classification against the EC hierarchy [1] as one of the comparison methods to the reaction kernel approach.

## 3 Structured Output Learning

Our objective is to learn a function that, given (a feature representation of) a sequence, can predict (a feature representation of) an enzymatic reaction.

Learning algorithms that are designed for structured prediction tasks like the above, are many. We concentrate on kernel methods, that let us utilize high-dimensional feature spaces without computing the feature maps explicitly. Structured SVM [22], Max-Margin-Markov networks [21][17], Kernel Density Estimation (KDE) and Maximum-Margin Regression (MMR) [20] are learning methods falling into this category.

We consider a training set of (sequence, reaction)-pairs  $D^m = \{(x_i, y_i) | x_i \in \mathcal{X}, y_i \in \mathcal{Y}\}_{i=1}^m$  drawn from an unknown joint distribution  $\mathcal{P}(\mathcal{X}, \mathcal{Y})$ .

For sequences and reactions, respectively, we assume feature mappings  $\phi: \mathcal{X} \mapsto \mathcal{F}_X$  and  $\psi: \mathcal{Y} \mapsto \mathcal{F}_Y$ , mapping the input and output objects into associated inner product spaces  $\mathcal{F}_X$  and  $\mathcal{F}_Y$ . The kernels  $K_X(x, x') = \langle \phi(x), \phi(x') \rangle$  and  $K_Y(y, y') = \langle \psi(y), \psi(y') \rangle$  defined by the feature maps are called the input and output kernel, respectively. Above  $\langle \cdot \rangle$  denotes the inner product. Subsequently, we discuss particular choices for the feature mappings and the kernels suitable for the enzyme function prediction task.

### 3.1 Joint Kernels

In structured prediction models based on kernels, the associations between the inputs and outputs are typically represented by a *joint* kernel, defined by some feature map joint for inputs and outputs. In this paper, we use a joint feature map

$$\varphi(x, y): \mathcal{X} \times \mathcal{Y} \mapsto \mathcal{F}_{X \otimes Y},$$

where  $\varphi(x, y) = \phi(x) \otimes \psi(y)$  is the tensor product of input and output feature maps, thus consisting of all pairwise products  $\phi_j(x)\psi_k(y)$  between input and output features. This choice gives us the joint kernel representation as elementwise product of the input and output kernels

$$K_{XY}(x, y; x', y') = K_X(x, x')K_Y(y, y').$$

The tensor product kernel is suitable in situations where there is no prior alignment information of input and output features available, but the learning machine is expected to learn the alignments. This is the case in our enzyme function prediction setup.

### 3.2 Learning Task

Most structured prediction models [21,22,20,17] take the form of a linear score function

$$F_w(x, y) = \langle w, \varphi(x, y) \rangle = \langle w, \phi(x) \otimes \psi(y) \rangle$$

in the joint feature space. The model's prediction  $\hat{y}(x)$  corresponds to highest scoring output  $y$ :

$$\hat{y}(x) = \mathbf{argmax}_y F_w(x, y).$$

For the model learning we use two computational methods. The first method is Kernel Density Estimation (KDE) which uses the joint kernel density function

$$F_w(x, y) = \sum_i K_{XY}(x, y; x_i, y_i) \tag{1}$$

for scoring. This is the simplest model we use for prediction, since there is no weighting vector  $w$  for the training examples and all the datapoints are thus equally important.

The second method, Max-Margin Regression (MMR) [20] aims to separate the training data  $\varphi(x_i, y_i)$  from the origin of the joint feature space with maximum margin, thus



it can be seen analogous to the one-class SVM [18]. The primal form of the MMR optimization problem can be written as

$$\begin{aligned} \min \quad & \frac{1}{2} \|w\|^2 + C \sum_i \xi_i \\ \text{s.t.} \quad & \langle w, \varphi(x_i, y_i) \rangle \geq 1 - \xi_i \\ & \xi_i \geq 0, i = 1, \dots, m. \end{aligned}$$

The dual form of the MMR problem can be expressed as

$$\begin{aligned} \max \quad & \sum_{i=1}^m \alpha_i - \frac{1}{2} \sum_{i,j=1}^m \alpha_i \alpha_j K_X(x_i, x_j) K_Y(y_i, y_j) \\ \text{s.t.} \quad & 0 \leq \alpha_i \leq C, i = 1, \dots, m. \end{aligned} \quad (2)$$

MMR, due to its simple form, can be optimized very efficiently which makes, for example, the optimization of kernel parameters a feasible task on medium sized datasets ( $10^3$ - $10^4$  examples), which is not true for most competing approaches [21,22,17].

Furthermore, as the output representation is kernelized, it is possible to learn in very complex output spaces, as we will demonstrate subsequently.

### 3.3 Preimage Problem

In all structured output prediction approaches, the prediction of the model needs to be extracted by solving the preimage problem

$$\hat{y}(x) = \mathbf{argmax}_{y \in \mathcal{Y}} F_w(x, y).$$

Depending on the output space, solving the preimage exactly can be computationally challenging or intractable.

Using kernelized outputs, as in the case of dual MMR (2), the preimage takes an even more challenging form

$$\hat{y}(x) = \mathbf{argmax}_{y \in \mathcal{Y}} \sum_i \alpha_i K_X(x, x_i) K_Y(y, y_i),$$

for which efficient algorithms are hard to come by. However, a difference between MMR and most structured output prediction methods is that there is no need to solve the preimage problem as part of the training, only during prediction. Thus, the computational complexity of the preimage is not as a crucial issue.


In the experiments reported in this paper, we use a brute-force preimage algorithm: we enumerate the set of outputs contained in our whole dataset (training and test examples included)  $\mathcal{Y}^n = \{y | (x, y) \in D^n\}$ . This approach will give us an approximate solution to the preimage problem, that is, the globally best scoring prediction may lie outside the set  $\mathcal{Y}^n$ . This approach is sufficient for first evaluation of the proposed prediction methods. We note that the set of putative reactions  $\mathcal{Y}$  could be much larger, e.g. all chemical reactions among metabolites below certain size.

## 4 Kernels for Chemical Reactions


In this section, we consider how to build kernels for chemical reactions, using molecule graph kernels as the building blocks.

Let us first introduce some notation used in this section. We denote a basic set of reactions  $\mathcal{R}$ , where a reaction  $\rho(S(\rho), P(\rho)) \in \mathcal{R}$  is given by a set substrates  $S(\rho) \subset \mathcal{M}$  and products  $P(\rho) \subset \mathcal{M}$ <sup>1</sup>. The set of reactants is simply the union of substrates and products  $R(\rho) = S(\rho) \cup P(\rho)$ . A feature vector describing a reaction  $\rho$  is denoted by  $\psi(\rho)$  and the feature vector describing a molecule  $M$  is denoted by  $\phi(M)$ .

For illustration, consider a chemical reaction  $\rho = (\{S_1\}, \{P_1, P_2\})$  converting a substrate molecule  $S_1$  into two product molecules  $P_1$  and  $P_2$ , thus defined by the reaction equation



Consider now a second reaction  $\rho' = (\{S'_1, S'_2\}, \{P'_1, P'_2\})$ , converting substrates  $S'_1, S'_2$  into products  $P'_1 + P'_2$ , and back, expressed as



How can we measure the similarity of these reactions via kernels? The approach in this paper is to consider pairwise similarities of the constituent molecules and compute an aggregate on them. While there are many ways how this could be done in principle, two important considerations arise from the (bio)chemical reality:


- **Similarity of Reaction Events vs. Reactants.** We should make a distinction between the similarity of the reaction events versus the similarity of the reactant molecules. For example, enzymes belonging to the amino-transferase group are similar to each other in that they transfer a certain functional group (the amino group) from a reactant molecule to another. However, the reactant molecules need not be similar.

Conversely, there are many different transformations which can be performed on the same molecule. For example, pyruvate, an important hub metabolite in the central metabolism of all living cells, participates in many reactions. The transformations applied by the reactions may be very different from each other, although they work on the same substrate molecule pyruvate.

Thus, depending on the application, our kernel should be designed to measure one of these similarity notions, or measure both of them in some proportion.


- **Directionality of Reactions.** The reactions may be defined as unidirectional or bidirectional. As the direction of a reaction depends on thermodynamical conditions, this may or may not be a relevant issue. For example, most enzymatic reactions are bidirectional in principle, but the conditions inside a living cell force unidirectionality.

When the directionality of reactions is of importance, each bidirectional reaction can be divided into forward and backward reactions. In our example, we would obtain



<sup>1</sup> To fully represent chemical reaction equations, we would also need to consider the stoichiometric coefficients for each reactant; However, we ignore this modelling aspect here.

and



In this case we would like our kernel to be sensitive to the direction so that forward and backward directions of the same reaction can be discriminated in the feature space.

However, when reaction direction is of no importance, the forward and backward directions of a bidirectional reaction should be treated the same by our kernel.

Below, we will describe a molecule graph kernel matrix  $K_{\mathcal{M}}$  which constitutes the basic component of the two alternative reaction kernels described next. For both reaction kernels we also show the underlying feature map which will suffice to show that both of the reaction kernels below are valid Mercer kernels if the underlying molecule kernel is a valid Mercer kernel.

Both of the reaction kernels described below are very fast to compute, given that the molecule kernel  $K_{\mathcal{M}}$  is pre-computed: the time complexity of the reaction kernel computation is then linear in the number of the elements in the kernel matrix.

#### 4.1 Kernels for Molecule Graphs

As the molecule kernel  $K_{\mathcal{M}}$  underlying the reaction kernels we use a subgraph kernel restricted to small subgraphs (10 nodes or less). The kernel computes the product graph of the two molecule graphs and counts its connected subgraphs. The kernel constructed in this way may in general not be a valid Mercer kernel. However, on our dataset, the kernel matrix was observed to be positive semidefinite.

Enumerating the subgraphs up to the maximum subgraph size  $d$  takes  $O(m^d)$  time, where  $m$  is the number of edges in the product graph. Thus the kernel is quite time-consuming to compute. In practice, we were able to compute the common connected subgraphs of 1767 KEGG LIGAND [7] molecules up to subgraph size 10 in a week with approximately 50 Pentium 4 class computers. Considering the computational resources available nowadays in research labs, and the time available to solve a typical problem involving molecular data, the computational complexity hardly presents a prohibitive constraint.

We note that it would also be possible to use a more quickly computable graph kernel based on common walks [6], that is, sequences of labeled atoms and bonds, which can be thought to approximate common subgraphs (each common subgraph induces a set of common walks). However, we leave exploring this direction as future work.


#### 4.2 Sum-of-Reactants Kernel

A simple kernel, called the Sum-of-Reactants (SoR) kernel, is obtained by defining

$$K_{SoR}(\rho, \rho') = \mathbf{m}(\rho)^T K_{\mathcal{M}} \mathbf{m}(\rho'),$$

where the vector  $\mathbf{m}(\rho)$  consists of indicators  $m_j(\rho) = \mathbf{1}_{\{M_j \in R_\rho\}}$  for the presence or absence of a molecule  $M_j$  in the set of reactants of  $\rho$ . The corresponding feature vector is simply the sum of feature vectors of molecule graphs in  $R_\rho$ :

$$\psi(\rho) = \sum_{M \in R_\rho} \phi(M)$$



**Fig. 1.** Examples of matchings induced by the SoR kernel. Top left and top right pictures represent a valid and a spurious matching for a unidirectional reaction, respectively. Bottom left and bottom right pictures represent a spurious and a valid matching for a bidirectional reaction, respectively.

Intuitively, the kernel measures the similarity of reactions in terms of how similar the molecules manipulated by the reactions are on average, rather than the similarity of reaction events. The reaction representation and the kernel can be considered bidirectional as the different roles of reactant molecules are not considered.

### 4.3 Reactant-Matching Kernels

In the SoR kernel there is an underlying all-against-all matching between the substrate sets  $(S_\rho, S'_{\rho'})$ , product sets  $(P_\rho, P'_{\rho'})$  and between the cross-pairs  $(S_\rho, P'_{\rho'})$  and  $(P_\rho, S'_{\rho'})$ . This measure implicitly contains spurious matches of two kinds:

1. Given two unidirectional reactions, matching substrates of one reaction with the products of another (e.g. Fig 1 top right) is not justified. However, it may be perfectly sensible to match the first substrate of one reaction to the second substrate of another as the order of declaring the reactants on one side of the reaction is merely a convention (e.g. Fig 1 top left).
2. Matching one substrate  $s_1 \in S_\rho$  against a substrate  $s' \in S'_{\rho'}$  and another  $s_2 \in S_\rho$  against a product  $p' \in P'_{\rho'}$  has no biological justification (Fig 1 bottom left). However, for bidirectional reactions, matching all substrates of one reaction against the products of another is perfectly sensible as it corresponds to flipping the nominal direction of one of the reactions (Fig 1 bottom right).

We can filter out the above spurious matches by defining a feature map via the tensor product

$$\psi(\rho) = \sum_{M \in S_\rho} \phi(M) \otimes \sum_{M \in P_\rho} \phi(M),$$

which gives us the Reactant-Matching (RM) kernel

$$K(\rho, \rho') = K(S_\rho, S'_{\rho'})K(P_\rho, P'_{\rho'}),$$

where we use the shorthand

$$K(S, S') = \sum_{M \in S} \sum_{M' \in S'} K_{\mathcal{M}}(M, M').$$

The above kernel is obviously unidirectional as it matches the reactions in the forward direction. To obtain a bidirectional kernel we compute the backward direction by taking the cross terms

$$K(\rho, \rho') = \frac{1}{2}(K(S_{\rho}, S_{\rho'})K(P_{\rho}, P_{\rho'}) + K(S_{\rho}, P_{\rho'})K(P_{\rho}, S_{\rho'}))$$

We note that the bidirectional kernel still filters out the above mentioned spurious matches, in the second term the other reaction is just flipped around.

## 5 Experiments

### 5.1 Data

The dataset is a sample(sequence, reaction) pairs from the KEGG LIGAND database [7]. As the input (sequence) representation, we use Global Trace Graph (GTG, [9]) features that can be interpreted as predicted conserved amino acids.

We have two separate datasets: the parameter validation set of 1481 enzymes and testing set of 8112 enzymes, which do not have overlapping EC numbers. Parameter validation set is yet divided into two folds, training set of 930 and test set of 551 enzymes. Testing dataset is divided into five folds with average of 1622 enzymes. Members of the folds are chosen such that each of the different EC number exist only in one of the folds, so the training sets have no enzymes with the test set EC number appear. This is to simulate setting where a previously unseen functions are to be predicted.

Both the input (GTG) kernel and the output (reaction) kernels are fed to a polynomial kernel  $K_{poly}(x, z) = (K(x, z) + 1)^d$  and normalized. The restricted size subgraph kernel is used as the molecule kernel underlying all the reaction kernel variants.

### 5.2 Compared Methods

We compare the following methods:

- NN(BLAST): This is the baseline annotation transfer method: given a test sequence, find the nearest sequence neighbor in the training set and transfer the annotation to the new protein. Sequence similarity is taken from pre-computed Blast scores from the Pairs-DB server [8].
- NN(GTG): This is the annotation transfer methods using the GTG data. Given a test sequence, find the training sequence with the most common GTG features with the test sequence, and transfer the annotation.
- MMR(GTG,Hierarchical): The hierarchical structured output prediction from [1]. The method predicts the membership of the new protein in the EC hierarchy; generally the prediction is a root-to-leaf path in the EC hierarchy.

- MMR(GTG, RM): MMR with GTG as input kernel and Reactant-Matching as output kernel.
- KDE(GTG, RM): KDE with GTG as input kernel and Reactant-Matching as output kernel.

We have beforehand made an experiment where we compared the function prediction accuracy with both of the reaction kernels using degree-6 polynomial kernel over the inputs and degree-20 polynomial kernel over the outputs. F1 score for RM was 27.9% and for SoR it was 25.9%. Since the RM outperformed SoR, we use the RM as output kernel in all the following experiments.

### 5.3 Measure of Success

To measure accuracy of prediction, for each test instance  $(x, y)$ , we first compute the set of top-scoring functions  $\hat{\mathcal{Y}}(x) = \{y_i \in \mathcal{Y}^n | F(\alpha, x, y_i) \geq F(\alpha, x, y'), \forall y' \in \mathcal{Y}^n\}$ , that is the reactions that the prediction model considers the (equally) best. This set is considered as the prediction of the model.

For each function  $y' \in \hat{\mathcal{Y}}(x)$ , we check how many consecutive digits starting from the left of the EC number associated with  $y'$  coincide with digits of the EC number associated with the reference function  $y$ . Each such correctly predicted EC digits counts as a true positive, rest of the EC digits counts as a false positive. For example, if the reference function  $y$  is 3.1.1.1 and prediction set  $\hat{\mathcal{Y}}(x)$  contains two members **3.1.2.1** and **3.1.1.10**, there are five true positives (marked bold) and three false positives out of 8 EC digits. The EC digit F1 is then the F1 score taken over all EC digit predictions in the test set.


### 5.4 Results

**Effect of Polynomial Kernel Degree.** In the first experiment we illustrate the behaviour of the structured output learning of MMR in very high-dimensional joint feature space. We use the GTG kernel (predicted conserved residues) as the base input kernel and the RM kernel as the base output kernel.

In this experiments we use two sets: one for training and second for testing. Figure 2 shows a heat map of the EC digit F1 score. The F1 score improves when the degree either the input, the output or both input and output polynomial kernels increases. The optimum reaches a plateau at input degrees 1-4 and output degrees 8-16 indicating robustness with respect to changes in parameter values.

Applying a high-degree polynomial to the base kernel makes the resulting output kernel more sparse, which suggests that the reactant matching kernel alone is too smooth for optimum performance. We note that optimizing the input and output kernels independently can be useful in other structured prediction settings as well.


**Prediction under Remote Homology.** In the final experiment, we demonstrate the generalization ability of the structured output prediction methods. We measure how many EC digit are correctly predicted in testing over a five fold set of enzyme families where the four digit EC numbers are not overlapping between folds. Thus the training



**Fig. 2.** The EC digit F1 score plotted as the function of the degrees of the input and output kernels. The best results are obtained with degree 2 polynomial over the inputs and degree 8 or higher over the outputs.

set contains no enzyme that has exactly the same EC number, but families that have three matching EC digits typically appear in the training.

In this setup it should be clear that the nearest neighbor classifier or the hierarchical classifier cannot ever predict four-digit EC number correctly, as the methods have not seen any examples of that particular family. The reaction kernel approach, however, does not suffer from this limitation: as all possible reactions can be represented in the output space, it is in principle possible to predict the correct function.



**Fig. 3.** The cumulative distribution of correctly predicted EC digits in the test set (bottom chart). Each member of the top ranking predictions  $\hat{Y}(x)$  contributes one item in the distribution.

Figure 3 shows the results of this experiment. Here, we used a degree 8 polynomial kernel over the RM kernel and degree 2 polynomial kernel over the inputs. In the bottom chart is the cumulative chart depicting the number of enzyme families that have at least certain number of correctly predicted EC digits.

It can be seen that the methods relying on the GTG features (NN(GTG), KDE(GTG, RM) and both MMR methods) are more effective in predicting more than one EC digits correctly. The KDE reaction kernel and MMR hierarchical approach is slightly better in predicting two or more EC digits correctly than the competing approaches. Finally, we note that the reaction kernel approach is the only method that, at times, can get the whole EC number correct. In other words, the set of top-ranking reactions  $\hat{Y}(x)$  contain reactions that possess the exactly correct EC number.

## 6 Discussion and Future Work

The present experiments show the potential of structured output prediction using reaction kernels: given a novel, previously unseen enzymatic function, the reaction kernel approach is significantly more accurate than the annotation transfer approach and also compares with a hierarchical classifier trained with structured output learning.

Also we note that the reaction kernel approach is an *enabling* technique: it is possible, albeit not easy, to predict the new function exactly correctly. Interestingly best results are obtained with a highly complex output representation: a high-degree polynomial kernel over reactant matching kernel.

As the result show, using the reaction kernel methods for enzyme function prediction is encouraging way to go, even if the prediction accuracy is still very low for all of the methods used. There are many areas where the methods can be improved. First, we only used predicted conserved residues (GTG) as inputs. Although they work well, augmenting them with other types of data, e.g. structural information should be helpful. Second, the presented reaction kernels certain can be improved and completely different kinds of encodings of enzyme function can be imagined.

Third, a better preimage algorithm will be needed for efficient prediction, brute-force enumeration of reactions will not be feasible for huge collection of putative reactions needed for truly novel function prediction. Also, as simpler output representations may provide more efficient preimage algorithms, it would be tempting to simplify the representations. However, in our view this should not be done at the expense of predictive accuracy.

**Acknowledgements.** This paper has benefited from discussions with Craig Saunders and John Shawe-Taylor. The work has been supported by Academy of Finland grants 110514 (UR-ENZYMES) and 118653 (ALGODAN), and in part by the IST Programme of the European Community under PASCAL2 Network of Excellence, IST-2007-216886. This publication only reflects the authors' views.

## References

1. Astikainen, K., Holm, L., Pitknen, E., Szedmak, S., Rousu, J.: Towards structured output prediction of enzyme function. In: BMC Proceedings, vol. 2(S4), S2 (2008)
2. Barutcuoglu, Z., Schapire, R., Troyanskaya, O.: Hierarchical multi-label prediction of gene function. *Bioinformatics* 22(7), 830–836 (2006)



3. Blockeel, H., Schietgat, L., Struyf, J., Džeroski, S., Clare, A.: Decision trees for hierarchical multilabel classification: A case study in functional genomics. In: Fürnkranz, J., Scheffer, T., Spiliopoulou, M. (eds.) PKDD 2006. LNCS (LNAI), vol. 4213, pp. 18–29. Springer, Heidelberg (2006)
4. Borgwardt, K.M., Ong, C.S., Schnauer, S., Vishwanathan, S.V.N., Smola, A.J., Kriegel, H.P.: Protein function prediction via graph kernels. *Bioinformatics* 21(1), 47–56 (2005)
5. Clare, A., King, R.: Machine learning of functional class from phenotype data. *Bioinformatics* 18(1), 160–166 (2002)
6. Gartner, T.: A survey of kernels for structured data. *SIGKDD Explorations* 5 (2003)
7. Goto, S., Okuno, Y., Hattori, M., Nishioka, T., Kanehisa, M.: Ligand: database of chemical compounds and reactions in biological pathways. *Nucleic Acids Research* 30(1), 402 (2002)
8. Heger, A., Korpelainen, E., Hupponen, T., Mattila, K., Ollikainen, V., Holm, L.: Pairsdb atlas of protein sequence space. *Nucl. Acids Res.* 36, D276–D280 (2008)
9. Heger, A., Mallick, S., Wilton, C., Holm, L.: The global trace graph, a novel paradigm for searching protein sequence databases. *Bioinformatics* 23(18) (2007)
10. Heinonen, M., Lappalainen, S., Mielikäinen, T., Rousu, J.: Computing Atom Mappings for Biochemical Reactions without Subgraph Isomorphism. *Journal of Computational Biology* (to appear 2011)
11. Holm, L., Sander, C.: Dali/fssp classification of three-dimensional protein folds. *Nucleic Acids Research* 25(1), 231–234 (1996)
12. Lanckriet, G., Deng, M., Cristianini, N., et al.: Kernel-based data fusion and its application to protein function prediction in yeast. In: PSB 2004 (2004)
13. Pitkänen, E., Jouhten, P., Rousu, J.: Inferring branching pathways in genome-scale metabolic networks. *BMC Systems Biology* 3(1), 103 (2009)
14. Pitkänen, E., Rousu, J., Ukkonen, E.: Computational methods for metabolic reconstruction. *Current Opinion in Biotechnology* 21, 70–77 (2010)
15. Punta, M., Ofran, Y.: The rough guide to in silico function prediction, or how to use sequence and structure information to predict protein function. *PLoS Computational Biology* 4(10) (2008)
16. Rantanen, A., Rousu, J., Jouhten, P., Zamboni, N., Maaheimo, H., Ukkonen, E.: An analytic and systematic framework for estimating metabolic flux ratios from 13 C tracer experiments. *BMC bioinformatics* 9(1), 266 (2008)
17. Rousu, J., Saunders, C., Szedmak, S., Shawe-Taylor, J.: Kernel-based learning of hierarchical multilabel classification models. *JMLR* 7 (2006)
18. Schlkopf, B., Platt, J.C., Shawe-Taylor, J., Smola, A.J., Williamson, R.C.: Estimating the support of a high-dimensional distribution. *Neural Computation* 13(7), 1443–1471 (2001)
19. Sokolov, A., Ben-Hur, A.: A structured-outputs method for prediction of protein function. In: Proceedings of the 3rd International Workshop on Machine Learning in Systems Biology (2008)
20. Szedmak, S., Shawe-Taylor, J., Parado-Hernandez, E.: Learning via linear operators: Maximum margin regression. Tech. rep., Pascal (2005)
21. Taskar, B., Guestrin, C., Koller, D.: Max-margin markov networks. In: NIPS 2003 (2004)
22. Tsochantaridis, I., Hofmann, T., Joachims, T., Altun, Y.: Support vector machine learning for interdependent and structured output spaces. In: ICML (2004)

# A Parallel Algorithm for Counting Subgraphs in Complex Networks

Pedro Ribeiro, Fernando Silva, and Luís Lopes

CRACS & INESC-Porto LA

Faculdade de Ciências, Universidade do Porto, Porto, Portugal

{pribeiro, fds, lblopes}@dcc.fc.up.pt

**Abstract.** Many natural and artificial structures can be represented as complex networks. Computing the frequency of all subgraphs of a certain size can give a very comprehensive structural characterization of these networks. This is known as the subgraph census problem, and it is also important as an intermediate step in the computation of other features of the network, such as network motifs. The subgraph census problem is *computationally hard* and most associated algorithms for it are sequential. Here we present several increasingly efficient parallel strategies for, culminating in a scalable and adaptive parallel algorithm. We applied our strategies to a representative set of biological networks and achieved almost linear speedups up to 128 processors, paving the way for making it possible to compute the census for bigger networks and larger subgraph sizes.

**Keywords:** Complex networks, Graph mining, Parallel algorithms, Subgraph census.

## 1 Introduction

A broad range of biological structures can be represented as complex networks. The study of such networks is relatively recent and has received increased attention by the scientific community [2]. A large number of concepts and techniques appeared to analyze and understand complex networks, leading to an impressive panoply of different measurements used to mine interesting data from them [6].

One important measure is the frequency in which subgraphs appear in a network. Sometimes we are just interested in determining frequent patterns [13], while in others we need to determine a full count of all different classes of isomorphic subgraphs [4]. This last option is also known as a *subgraph census* and can provide a very accurate structural characterization of a network. This is typically applied for subgraphs of a specific size and it is normally limited to small sizes, mostly for efficiency reasons. This has been done not only on biological networks [16], but also on other domains, such as social networks analysis, where the triad census is very common [25].

Subgraph census also plays a major role as an intermediate step in the calculation of other important measures, such as *network motifs* [17], which are basically subgraphs that are statistically over-represented in the network (and conjectured to have some functional significance). Network motifs have applications on several biological domains, like protein-protein interaction [1] and brain networks [22], or in other domains, like electronic circuits [10]. The practical available algorithms and tools for network

motifs all use a census to discover the frequency in the original network and then calculate it again for a series of similar randomized networks [17,27]. This is a computationally hard problem that is closely related to the problem of graph isomorphism [15]. Some techniques were developed to speedup the calculations, like sampling [12], but they normally trade accuracy for speed.

In all these applications, having a more efficient way to calculate the census is highly desirable. As we increase the size of the subgraphs, their frequency increases exponentially and it becomes unfeasible to count all of them using traditional approaches. Moreover, to date, almost all algorithms for complete subgraph census are sequential. Some exceptions exist, particularly in the area of network motifs, but they are scarce and still limited (c.f. section 2.3). One reason is that present network motifs methods still resort to the generation of hundreds of random networks to measure significance. This puts the obvious opportunity for parallelism not in the census itself but in the generation of random networks and their respective census. However, analytical methods to estimate the significance are now appearing [14,19] and once they are fully developed the burden of the calculation will then reside on the census of the original network.

Considering the relevance of calculating exhaustive census and the computational complexity involved, resorting to parallel algorithms to speedup the calculation is, in our view, an approach that will impact in many application areas, particularly in the study of biological networks. The use of parallelism can not only speed up the calculation of census, but also allow the calculation of the census for subgraph sizes that were until now unreachable.

This paper focuses on strategies for solving the subgraph census problem in parallel. With this objective in mind we start with an efficient sequential algorithm, ESU [27], and progressively modify it to accommodate scalable parallel execution and data-structures. This process led us to the formulation of a novel adaptive parallel algorithm for subgraph census that features a work sharing scheme that dynamically adjusts to the available search-space. The results obtained show that the algorithm is efficient and scalable.

The remainder of this paper is organized as follows. Section 2 establishes a network terminology, formalizes the problem we want to tackle and gives an overview of related work. Section 3 details all the followed parallel strategies and the algorithm we developed. Section 4 discusses the results obtained when applied to a set of representative biological networks. Section 5 concludes the paper, commenting on the obtained results and suggesting possible future work.

## 2 Preliminaries

### 2.1 Network Terminology

In order to have a well defined and coherent network terminology throughout the paper, we first review the main concepts and introduce some notation that will be used on the following sections.

A network can be modeled as a *graph*  $G$  composed of the set  $V(G)$  of *vertices* or *nodes* and the set  $E(G)$  of *edges* or *connections*. The *size* of a graph is the number of vertices and is written as  $|V(G)|$ . A  $k$ -graph is graph of size  $k$ . The *neighborhood* of

a vertex  $u \in V(G)$ , denoted as  $N(u)$ , is composed by the set of vertices  $v \in V(G)$  that are adjacent to  $u$  ( $u$  is not included). All vertices are assigned consecutive integer numbers starting from 0, and the comparison  $v < u$  means that the index of  $v$  is lower than that of  $u$ .

A *subgraph*  $G_k$  of a graph  $G$  is a graph of size  $k$  in which  $V(G_k) \subseteq V(G)$  and  $E(G_k) \subseteq E(G)$ . This subgraph is said to be *induced* if for any pair of vertices  $u$  and  $v$  of  $V(G_k)$ ,  $(u, v)$  is an edge of  $G_k$  if and only if  $(u, v)$  is an edge of  $G$ . The neighborhood of a subgraph  $G_k$ , denoted by  $N(G_k)$  is the union of  $N(u)$  for all  $u \in V(G_k)$ . The *exclusive neighborhood* of a vertex  $u$  relative to a subgraph  $G_k$  is defined as  $N_{excl}(u, G_k) = \{v \in N(u) : v \notin G_k \cup N(G_k)\}$ .

A *mapping* of a graph is a bijection where each vertex is assigned a value. Two graphs  $G$  and  $H$  are said to be *isomorphic* if there is a one-to-one mapping between the vertices of both graphs where two vertices of  $G$  share an edge if and only if their corresponding vertices in  $H$  also share an edge.

## 2.2 Subgraph Census

We give a rigorous definition for the subgraph census problem:

**Definition 1 (*k*-subgraph Census).** *A  $k$ -subgraph census of a graph  $G$  is determined by the exact count of all occurrences of isomorphic induced subgraph classes of size  $k$  in  $G$ , where  $k \leq |V(G)|$ .*

Note that this definition is very broad and can be applied to all kinds of networks, whether they are directed or undirected, colored or not and weighted or unweighted. Also note that here, unlike in [12], we are concerned with an exact result and not just an approximation.

A crucial concept that we have not yet completely defined is how to distinguish two different occurrences of a subgraph. Given that we are only interested in finding induced subgraphs, we can allow an arbitrary overlap of vertices and edges or have some constraints such as no edge or vertex sharing by two occurrences. The several possibilities that we can have for the frequency are considered and discussed in [21]. Here we focus on the most widely used definition that we formalize next:

**Definition 2 (Different Occurrences of  $k$ -subgraphs).** *Two occurrences of subgraphs of size  $k$ , in a graph  $G$ , are considered different if they have at least one vertex or edge that they do not share. All other vertices and edges can overlap.*

Note that this has a vital importance on the number of subgraphs we find and consequently to the tractability of the problem.

## 2.3 Related Work

There exists a vast amount of work on graph mining. Particularly, the field of frequent subgraph mining has been very prolific, producing sequential algorithms like Gaston [18]. Although related, these algorithms differ substantially in concept from our approach since their goal is to find the most frequent subgraphs that appear in a set of graphs, while we try to find the frequency of all subgraphs on a single graph.

Regarding subgraph census itself, most of the work on social networks is based on small sized subgraphs - mostly triads [25,8] - and therefore does not focus on efficiency, but rather on the interpretation of the results. However, for network motifs, efficiency does play an important role and much importance is given to the algorithm for generating the census. Increasing the speed may lead to detection of bigger patterns and even an increase in size of just one can yield scientifically important results because a new previously unseen pattern with functional significance may be discovered.

The three best known production tools for finding motifs are all based on serial algorithms. Mfinder [17] was the first and it is based on a recursive backtracking algorithm that generates all  $k$ -subgraphs. It may generate the same subgraph several times because it initiates a search procedure in each of its nodes. Fanmod [27] uses an improved algorithm called ESU, that only allows searches being initiated on the nodes with an index higher than the root node and therefore each subgraph is found only once. MAVisto [21] does not improve efficiency except when it uses a different concept for frequency.

Work on parallel algorithms for subgraph census is scarce. [23] propose an algorithm for finding frequent subgraphs but do not count all of them. [20] focuses on network motifs and how to parallelize queries of individual subgraphs and not on how to enumerate all of them. [24] takes the closest approach to our work. Their algorithm relies on finding a neighborhood assignment for each node that avoids overlap and redundancy on subgraph counts, as in [27], and tries to statically balance the workload “a priori” based only on each node degree (no details are given on how this is done and how it scales). Another distinctive characteristic of their approach is that they do not do isomorphism tests during the parallel computation, they wait until the end to check all the subgraphs and compute the corresponding isomorphic classes. As we will see, our approach differs significantly from this one as it contributes with dynamic and adaptive strategies for load balancing, thus attaining higher efficiency.

## 3 Parallel Algorithms

### 3.1 Core Sequential Unit

Given that we are interested in having an exact count of all classes of isomorphic subgraphs, we must enumerate all subgraphs. The ESU algorithm [27] is a key component of the fastest network motif tool available and as far as we know it is one of the most efficient algorithms for subgraph enumeration. Thus we chose the ESU algorithm as our starting point and modified its recursive part to create a procedure that given a graph  $G$ , a size  $k$ , a vertex minimum index  $min$ , a partially constructed subgraph  $G_{subgraph}$ , and a list of possible extension nodes  $V_{ext}$ , enumerates all  $k$ -subgraphs that contain  $G_{subgraph}$  and no nodes with index lower than  $min$ . This procedure is depicted in algorithm 1. It recursively extends the subgraph  $G_{subgraph}$  by first adding the new node  $u$ . If the new subgraph has size  $k$ , then it determines a unique identification and saves it in a dictionary. Otherwise, it expands the set of possible extension nodes,  $V_{ext}$ , with the nodes that are in the exclusive neighborhood of  $u$  relative to the subgraph  $G_{subgraph}$  and also satisfy the property of being numerically bigger than  $min$ . If the extension set of nodes is not null then a new node is removed from extended  $V_{ext}$  and recursion is made.

---

**Algorithm 1.** Extending a partially enumerated subgraph


---

```

1: procedure EXTEND( $G, k, min, u, G_{subgraph}, V_{ext}$ )
2:    $G'_{subgraph} \leftarrow G_{subgraph} \cup \{u\}$ 
3:   if  $|V(G'_{subgraph})| = k$  then
4:      $str \leftarrow \text{CanonicalString}(G'_{subgraph})$ 
5:     Dictionary.AddAndIncrement( $str$ )
6:   else
7:      $V'_{ext} \leftarrow V_{ext} \cup \{v \in N_{excl}(u, G_{subgraph}) : v > min\}$ 
8:     while  $V'_{ext} \neq \emptyset$  do
9:       Remove an arbitrarily chosen  $v \in V'_{ext}$ 
10:      Extend( $G, k, min, v, G'_{subgraph}, V'_{ext}$ )

```

---



**Fig. 1.** Example of how `Extend()` calls generate all subgraphs

Calling `Extend( $G, k, u, u, \{\}, \{\})$`  for every  $u \in V(G)$  is exactly the equivalent to the original ESU algorithm. Therefore, as long as we call it on all nodes, we can be certain that it will produce complete results, as shown in [27]. Moreover, `Extend()` guarantees that each existent subgraph will only be found once on the call of its lowest index, as exemplified in figure 1. This avoids redundant calculations as in [24] and is crucial to achieve an efficient census.

Before going into the details of the parallelism two additional notes are needed. First, isomorphism (line 4 of the procedure) is taken care of by using the canonical string representation of the graphs, defined as the concatenation of the elements of the adjacency matrix of the canonical labeling. In our case we use McKay’s *nauty* algorithm [15], a widely known fast and practical implementation of isomorphism detection. Second, in order to store the results found within one call to our procedure (line 5), we use a string dictionary structure. This can be implemented in many ways, for example using a hash table or a balanced red-black tree. We implement the later (using STL *map* from C++).

### 3.2 Initial Parallel Approaches

Each of the aforementioned calls to `Extend( $G, k, u, u, \{\}, \{\})$`  is completely independent from each other and we call it a *primary* work unit. A possible way of parallelizing subgraph census is then to distribute these work units among all CPUs<sup>1</sup>. The problem is that these units have a computational cost with a huge variance, as the inherent substructure and the number of subgraphs each one enumerates are also quite different.

---

<sup>1</sup> From now on we will refer to processors in computational nodes as CPUs or workers to avoid confusion between them and graph nodes.

We experimented several strategies for the distribution in order to obtain the desired load balance. The first one was to statically allocate the units to workers before starting the census computation. In order to obtain good results this would need accurate estimates of the time that each unit takes to compute. We were unable to do that with the desired accuracy, since calculating this is almost as difficult as enumerating the subgraphs themselves.

We then took the path of a more dynamic approach using a master-worker architecture. The master maintains a list of unprocessed primary work units. Workers ask the master for a unit, process it and repeat until there is nothing more to compute. Each worker maintains its own dictionary of frequencies. When all work units have been computed the master is responsible for collecting and merging all results, summing up the frequencies found.


The position of the work units on the master's list will determine the total time needed and we tried several strategies. Initially we just added all work units to the list in chronological order of the nodes. This proved to be a bad strategy since it is the same as a random assignment, which is in principle the worst possible. We then experimented giving the work units sorted to an estimated cost, using *LPTF (Largest Processing Time First)* strategy. If the estimate was perfect, it is known that we would achieve at least  $\frac{3}{4}$  of the optimum [9]. We only had an approximation (based on the number of nodes achievable in  $k - 1$  steps) and therefore that boundary is not guaranteed. However, since our heuristic function maintained a reasonable ordering of the nodes, the performance was vastly improved.

We still had the problem that the call to a few primary work units (potentially even just one) could consume almost all the necessary compute time. This prevents good load balance strategies, given that each work units runs sequentially. This problems occurs very often in reality because typical complex networks are scale free [3]. Whenever their hubs are the starting nodes of a work unit, very large neighborhoods are induced and a huge amount of subgraphs is generated. No matter what we do, there will always be a worker computing the largest sequential work unit and therefore the total compute time needed cannot be smaller than that. On some of our experiments with real biological networks, this largest atomic unit could consume more than 25% of the total execution time, which limits our scalability.

Considering that a work unit only calculates subgraphs containing nodes with indices greater than the index of the initial node, we devised a novel strategy in which we give higher index numbers to the potentially more time consuming nodes (those with larger degrees). This reduces the number of subgraphs spawning from these nodes, thus reducing the granularity of the work units induced by those nodes. To accomplish this strategy, we implemented a node relabeling algorithm in which the nodes are sorted increasingly by their degree. This improved our results, but we could not reduce enough the granularity, and we need a strategy that can divide the work units further.

### 3.3 Adaptive Parallel Enumeration

By inspecting the computation flow of a primary work unit, we can observe that there are several recursive calls to `Extend()`, as exemplified in figure 2 (for simplicity, we do not show  $G$  and  $k$  since these are fixed arguments).



**Fig. 2.** The computation flow of a primary work unit

With our formulation of `Extend()`, all recursive calls are independent with no need for information of previous data on the recursion stack besides the arguments it was called with. One way to divide a primary work unit is therefore to partitionate it in its recursive calls. A tuple  $(min, u, G_{subgraph}, V_{ext})$  completely defines the resulting call to `Extend()` and we will now call work unit to a tuple like this, with primary work units being only a particular case.


Our new strategy to reduce the granularity of the work units uses a threshold parameter to indicate the point in the computation at which we split the execution of the current work unit into smaller work units. Instead of really computing subsequent recursive calls, we encapsulate their arguments into new smaller work units and send them to the master to be added to the list of unprocessed work, effectively dividing our previously atomic sequential units. This leads to a simpler, yet elegant, solution when compared to more common adaptive strategies that need a queue in each computation node [7]. Figure 3 illustrates our strategy at work. Remember that the new work units are still independent and we do not need to be concerned with locality. All subgraphs will be found and added to the respective worker’s dictionary of frequencies, being merged in the end of the whole computation to determine the global resulting census.

Our algorithm is able to adjust itself during execution using this division strategy. It splits large work units into new smaller work units ensuring that their grain-size will never be larger than the size of work units executed up to the threshold value. In doing so, we are able to improve the load balancing and thus achieve an effective dynamic and adaptive behavior.

The splitting threshold parameter is central in our adaptive algorithm. If it is set too high, the work units will not be sufficiently divided in order to adequately balance the work among all CPUs. If it is too low, work will be divided too soon and the communication costs will increase. As a proof of concept our current implementation uses a threshold that limits the computation time spent on a work unit to a maximum value, but other measures could be used like for example the number of subgraphs already enumerated.

One aspect not yet discussed, but that is orthogonal to all discussed strategies, concerns the aggregation of results at the master. If a naive approach was taken, then each worker would be sending their results to the master sequentially. This would be highly inefficient and therefore we devised a parallel approach for this final step. We use an hierarchical binary tree to organize the aggregation of results, where each worker receives the results of two other *child* workers, updates its own frequency dictionary accordingly,





```

Extend(1, 1, {}, {});  $V'_{ext} = \{2, 3, 4\}$ 
Extend(1, 2, {1}, {3, 4});  $V'_{ext} = \{3, 4, 5\}$ 
Extend(0, 3, {1, 2}, {4, 5}) → Found {1, 2, 3}
Extend(0, 4, {1, 2}, {5}) → Found {1, 2, 4}
----- Splitting Threshold -----
Extend(0, 5, {1, 2}, {})
⇒ New work unit with these arguments
Extend(0, 3, {1}, {4});  $V'_{ext} = \{4\}$ 
⇒ New work unit with these arguments
Extend(0, 4, {1}, {});  $V'_{ext} = \{5\}$ 
⇒ New work unit with these arguments

```

**Fig. 3.** The computation flow of a primary work unit

and then in turn sends the aggregated results to its parent. This has the potential to logarithmically reduce the total time needed to accomplish this step.

All the ideas described are the basis for our main algorithm that we called *Adaptive Parallel Enumeration* (APE). Algorithms 2 and 3 describe in detail our APE master and worker procedures.

---

**Algorithm 2.** APE master node

---

```

1: procedure MASTER( $G, k$ )
2:    $L_{WorkUnits}.add(AllPrimaryWorkUnits)$ 
3:   while  $CPU_sWorking \neq \emptyset$  do
4:      $msg \leftarrow ReceiveMessage(AnyWorker)$ 
5:     if  $msg.type = RequestForWork$  then
6:       if  $L_{WorkUnits}.notEmpty()$  then
7:          $W \leftarrow L_{WorkUnits}.pop()$ 
8:          $newMsg \leftarrow EncapsulateWorkUnit(W)$ 
9:          $SendMessage(msg.Sender, newMsg)$ 
10:      else
11:         $IdleWorkers.push(msg.Sender)$ 
12:      else if  $msg.type = NewWorkUnit$  then
13:        if  $IdleWorkers.notEmpty()$  then
14:           $worker \leftarrow IdleWorker.pop()$ 
15:           $SendMessage(worker, msg)$ 
16:        else
17:           $W \leftarrow ExtractWorkUnit(msg)$ 
18:           $L_{WorkUnits}.push(W)$ 
19:       $BroadcastMessage(Terminate);$ 
20:       $ReceiveResults(LeftChild, RightChild)$ 

```

---

The master starts by adding all primary work units to the list of unprocessed work units ( $L_{WorkUnits}$ ). Then starts its main cycle where it waits for a message from a worker. If the message indicates that the worker needs more work, the master sends it the next unprocessed work unit  $L_{WorkUnits}$ . If the list is empty, the master signals

**Algorithm 3.** APE worker node

---

```

1: procedure WORKER( $G, k$ )
2:   while  $msg.type \neq Terminate$  do
3:      $msg \leftarrow ReceiveMessage(Master)$ 
4:     if  $msg.type = NewWorkUnit$  then
5:        $W = (G, k, min, u, G_{subgraph}, V_{ext}) \leftarrow ExtractWorkUnit(msg)$ 
6:        $Extend'(W)$ 
7:      $ReceiveResults(LeftChild, RightChild)$ 
8:      $SendResults(ParentWorker)$ 

9: procedure EXTEND'( $W$ )
10:  if SplittingThresholdAchieved() then
11:     $msg \leftarrow EncapsulateWorkUnit(W)$ 
12:     $SendMessage(Master, msg)$ 
13:  else
14:    lines 2 to 9 of algorithm 1
15:     $Extend'(W' = (G, k, min, v, G'_{subgraph}, V'_{ext}))$ 
16:    lines 11 and 12 of algorithm 1

```

---

the worker as being idle. If the message indicates that the worker is splitting work and thus sending a new unprocessed work unit, then the master adds that new work unit to  $L_{WorkUnits}$ . If there is an idle worker, then this unit is sent right away to that worker. When all workers are idle, the subgraph enumeration is complete and the master ends its main cycle, broadcasting to all workers that event. What remains is then to collect the results and following the explained hierarchical aggregation process, the master receives the results of two workers and merges them in an unified global dictionary of the frequencies of each isomorphic class of  $k$ -subgraphs.

The worker has a main cycle where it waits for messages from the master. If the message signals a new work unit to be processed, than it calls a modified version of the  $Extend()$  procedure to compute it. If the message signals termination, then it exits the cycle, receiving and merging the results from two other workers with its own dictionary. It then send those results to a single parent processor, that depending on the worker rank number may be other worker or the master itself, completing the hierarchical aggregation phase. Regarding the modified version of the  $Extend()$  procedure, it is exactly the same as the version depicted on algorithm [1](#) except the fact than when the splitting threshold is achieved, the computation is stopped and all subsequent calls consist now in encapsulating the arguments into a new work unit and sending it to the master.

There are two issues that we would like to clarify. First, we decided to use a dedicated master because it is a central piece in the architecture and we needed the highest possible throughput in the assignment of new work units to idle workers. Second, APE was originally created having in mind homogeneous resources but its dynamic and adaptive design makes it also suited for heterogeneous environments.

## 4 Results

All experimental results were obtained on a dedicated cluster with 12 SuperMicro Twin-view Servers for a total of 24 nodes. Each node has 2 quad core Xeon 5335 processors and 12 GB of RAM, totaling 192 cores, 288 GB of RAM, and 3.8TB of disk space, using Infiniband interconnect. The code was developed in C++ and compiled with gcc 4.1.2. For message passing we used OpenMPI 1.2.7. All the times measured were *wall clock times* meaning real time from the start to the end of all processes.

In order to evaluate our parallel algorithms we used four different representative biological networks from different domains: Neural [26], Gene [17], Metabolic [11] and Protein [5]. The networks present varied topological features that are summarized in Table 1.

**Table 1.** Networks used for experimental testing of the algorithms

Network	Nodes	Edges	Avg. Degree	Description
Neural	297	2345	7.90	Neural network of <i>C. elegans</i>
Gene	688	1079	1.57	Gene regulation network of <i>S. cerevisiae</i>
Metabolic	1057	2527	2.39	Metabolic network of <i>S. pneumoniae</i>
Protein	2361	7182	3.04	Protein-protein interaction network of <i>S. cerevisiae</i>

We first studied the computational behaviour of each network using the equivalent to the ESU algorithm, sequentially calling all primary work units (with no MPI overhead). This measures how much time a serial program would take to calculate a subgraph census. We took note of what was the maximum possible subgraph size  $k$  achievable in a reasonable amount of time (we chose one hour as the maximum time limit). We calculated the average growth ratio, that is, by which factor did the execution time grew up as we increased  $k$  by one. Finally, we also calculated the total number of different occurrences of  $k$ -subgraphs and the number of different classes of isomorphism found on those subgraphs. The results obtained can be seen in table 2.

**Table 2.** Maximum achievable subgraph sizes  $k$  using a serial program

Network	$k$	Time spent (s)	Average Growth	Total nr of subgraphs	Isomor. classes
Neural	6	10,982.9	47.6±0.4	$1.3 \times 10^{10}$	286,376
Gene	7	4,951.0	19.0±1.4	$4.2 \times 10^9$	4,089
Metabolic	6	14,000.1	46.5±2.7	$1.9 \times 10^{10}$	1,696
Protein	6	10,055.2	31.2±3.1	$1.3 \times 10^{10}$	231,620

Note the relatively small subgraph sizes achievable. This is not caused by our implementation, since using the FANMOD tool [27], the fastest available for network motifs calculation, we also were only able to achieve the same maximum  $k$  in one hour. The cause is that, as expected, the computing time grows exponentially as the subgraph size




Fig. 4. Speedups obtained with several parallel approaches

increases. We also observe that although different graphs present very different average growths, the growth rate for a given graph seems fairly constant (note the standard deviation).

For the next set of results we decided to fix the respective  $k$  for each graph to the values depicted in table 2 in order to have more comparable results. We evaluated the parallel strategies described in section 3. We compared the speedup obtained on all three graphs for the dynamic strategy with chronological order in the work units list (DYN-CRON), with LPTF ordering (DYN-LPTF), with graph relabeling followed by LPTF (DYN-RELABEL) and finally with the APE strategy.

For the APE algorithm it is necessary to explain how we chose the value for the splitting threshold parameter. We chose to employ the time spent in the same work unit as a proof of concept for the usefulness of APE and we empirically experimented several values for this time limit, reducing it while verifying that the speedup was being improved. This value controls the granularity of the work units. We want it as small as possible, as long as the increase in communication costs does not overcome the effect of increased sharing. We found that for our context 2 seconds appeared to be a good and balanced value (the time spent in communications during the enumeration of the subgraphs was always smaller than 2% of the total time spent), and we measured the speedup with that particular value chosen as the threshold (APE-2s).

We used a minimum of 8 CPUs because each computation node in the cluster had precisely that number of processors. With less CPUs the nodes would not be exclusively dedicated to the subgraph census. The results obtained up to 128 processors are depicted in figure 4. The results obtained clearly show different performance levels for the different strategies. Generally speaking, the strategies based on the atomic primary work units do not scale well, although the incremental strategies used show some improvements in the speedup. Overall, as expected, the adaptive strategy, APE-2s, outperforms all others and clearly achieves scalability.

Next, we further analyze the performance of APE-2s on all networks. Table 3 details the performance of the APE-2s up to 128 processors, and show the percentage of time spent in the final step of aggregating all frequency results in the master CPU.

We can observe that for Gene and Metabolic, APE-2s obtains almost perfect linear speedup, with a reduced amount of time spent in the aggregation phase. In Neural

**Table 3.** Detailed APE behavior with splitting threshold set to 2s

Network	$k$	#CPUs: speedup (% time spent in aggregating results)				
		8	16	32	64	128
Neural	6	7.0 (0.2%)	14.8 (1.2%)	30.1 (3.4%)	58.7 (8.0%)	107.0 (16.9%)
Gene	7	7.0 (0.1%)	15.0 (0.2%)	30.7 (0.4%)	62.0 (0.8%)	125.0 (0.7%)
Metabolic	6	6.9 (0.1%)	14.9 (0.1%)	30.8 (0.3%)	62.4 (0.6%)	125.5 (1.3%)
Protein	6	6.6 (0.2%)	13.7 (1.2%)	28.0 (3.3%)	54.6 (7.4%)	96.9 (18.8%)

and *Protein*, despite the good results, there is still some room for improvement. The time spent communicating the results in the end of the computation, more than 15% of the execution time with 128 processors, is the cause for the loss in the speedup. In fact, with more than 200,000 classes of isomorphic subgraphs in the network (see table 2), each worker has to communicate all of the frequencies it finds in its respective computation. On average, the number of different classes of isomorphism discovered on the same CPU, is larger than 150,000. Each of these classes has to be encapsulated (uniquely identifying the class) in order for the receiver to be able to decode it and merge the results. Even with our hierarchical method for aggregating the results, this still takes a considerable amount of time. This effect is not so noticeable in the other networks since the number of different classes found is much lower (due to inherent network topology, with a smaller average degree per node).

As a final demonstration of the relevance of our algorithm, consider the average growth as we increase  $k$  (shown in Table 2). As long as the number of processors we have available is larger than the average growth, we should be able to compute the  $(k + 1)$ -census in the same amount of time we initially were able to compute the  $k$ -census sequentially. For example, using the average growth, we can estimate that calculating a 7-census of *Metabolic* would take more than one week, if done sequentially. Using 128 processors and APE-2s, we were able to calculate it in less than 1h30m, spending even less than half of the time a sequential 6-census takes.

## 5 Conclusions

We presented several strategies for calculating subgraph census of biological networks in parallel. Our approaches are based on an efficient sequential algorithm called ESU that we parallelized by initially modifying it to a version capable of producing independent and dividable work units. We started with a dynamic master-worker strategy and subsequently improved it with an LPTF order of processing and a smart relabeling of the nodes in the graph. We also presented APE, a novel adaptive load balancing algorithm, which includes an hierarchical aggregation of the results found in each worker. APE proved to be an acceptable and scalable solution for the set of representative networks studied, successfully reducing the time needed to calculate the subgraph census and achieving larger subgraph sizes than were before possible.

The main drawback of APE seems to be the final aggregation of results. We plan to research and improve this step in the future. One way of doing it would be to use a more compact and compressed representation of the results. We also plan to research

the splitting threshold parameter in order to better understand on what does it depend, exactly how does it affect the computation and how could it be automatically determined by the algorithm. We are collaborating with neuroinformatics scientists in order to apply the described strategies on real neural networks to obtain new and interesting results on previously unfeasible subgraph census.

**Acknowledgements.** We thank Enrico Pontelli for the use of Inter Cluster in the New Mexico State University. Pedro Ribeiro is funded by an FCT Research Grant (SFRH/BD/19753/2004). This work was also partially supported by project CALLAS of the FCT (contract PTDC/EIA/71462/2006).

## References

1. Albert, I., Albert, R.: Conserved network motifs allow protein-protein interaction prediction. *Bioinformatics* 20(18), 3346–3352 (2004)
2. Alm, E., Arkin, A.P.: Biological networks. *Current Opinion in Structural Biology* 13(2), 193–202 (2003)
3. Barabasi, A.L., Albert, R.: Emergence of scaling in random networks. *Science* 286(5439), 509–512 (1999)
4. Bordino, I., Donato, D., Gionis, A., Leonardi, S.: Mining large networks with subgraph counting. In: *Proc of the 8th IEEE International Conference on Data Mining (ICDM)*, 6 pp.+ (December 2008)
5. Bu, D., Zhao, Y., Cai, L., Xue, H., Zhu, X., Lu, H., Zhang, J., Sun, S., Ling, L., Zhang, N., Li, G., Chen, R.: Topological structure analysis of the protein-protein interaction network in budding yeast. *Nucl. Acids Res.* 31(9), 2443–2450 (2003)
6. Costa, L., Rodrigues, F.A., Travieso, G., Villas Boas, P.R.: Characterization of complex networks: A survey of measurements. *Advances In Physics* 56, 167 (2007)
7. Eager, D.L., Lazowska, E.D., Zahorjan, J.: Adaptive load sharing in homogeneous distributed systems. *IEEE Trans. Softw. Eng.* 12(5), 662–675 (1986)
8. Faust, K.: Very local structure in social networks. *Sociological Methodology* 37(1), 209–256 (2007)
9. Hall, L.A.: Approximation algorithms for scheduling. In: *Approximation Algorithms for NP-Hard Problems*, USA, pp. 1–45. PWS Publishing Co. (1997)
10. Itzkovitz, S., Levitt, R., Kashtan, N., Milo, R., Itzkovitz, M., Alon, U.: Coarse-graining and self-dissimilarity of complex networks. *Phys. Rev. E Stat. Nonlin. Soft. Matter. Phys.* 71(1 Pt 2) (January 2005)
11. Jeong, H., Tombor, B., Albert, R., Oltvai, Z.N., Barabási, A.L.: The large-scale organization of metabolic networks. *Nature* 407(6804), 651–654 (2000)
12. Kashtan, N., Itzkovitz, S., Milo, R., Alon, U.: Efficient sampling algorithm for estimating subgraph concentrations and detecting network motifs. *Bioinformatics* 20(11), 1746–1758 (2004)
13. Kuramochi, M., Karypis, G.: Frequent subgraph discovery. In: *IEEE International Conference on Data Mining*, p. 313 (2001)
14. Matias, C., Schbath, S., Birmel, E., Daudin, J.-J., Robin, S.: Network motifs: mean and variance for the count. *REVSTAT* 4, 31–35 (2006)
15. McKay, B.: Practical graph isomorphism. *Congressus Numerantium* 30, 45–87 (1981)
16. Middendorff, M., Ziv, E., Wiggins, C.: Inferring network mechanisms: The drosophila melanogaster protein interaction network. *PNAS* 102, 3192 (2004)

17. Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D., Alon, U.: Network motifs: simple building blocks of complex networks. *Science* 298(5594), 824–827 (2002)
18. Nijssen, S., Kok, J.N.: Frequent graph mining and its application to molecular databases. In: *SMC*, vol. (5), pp. 4571–4577. IEEE, Los Alamitos (2004)
19. Picard, F., Daudin, J.-J.J., Koskas, M., Schbath, S., Robin, S.: Assessing the exceptionality of network motifs. *J. Comput. Biol.* (February 2008)
20. Schatz, M., Cooper-Balis, E., Bazinet, A.: Parallel network motif finding (2008)
21. Schreiber, F., Schwobbermeyer, H.: Towards motif detection in networks: Frequency concepts and flexible search. In: *Proc. of the Int. Workshop on Network Tools and Applications in Biology (NETTAB 2004)*, pp. 91–102 (2004)
22. Sporns, O., Kotter, R.: Motifs in brain networks. *PLoS Biology* 2 (2004)
23. Wang, C., Parthasarathy, S.: Parallel algorithms for mining frequent structural motifs in scientific data. In: *ACM International Conference on Supercomputing (ICS) 2004* (2004)
24. Wang, T., Touchman, J.W., Zhang, W., Suh, E.B., Xue, G.: A parallel algorithm for extracting transcription regulatory network motifs. In: *IEEE Int. Symposium on Bioinformatic and Bioengineering*, pp. 193–200 (2005)
25. Wasserman, S., Faust, K., Iacobucci, D.: *Social Network Analysis: Methods and Applications (Structural Analysis in the Social Sciences)*. Cambridge University Press, Cambridge (November 1994)
26. Watts, D.J., Strogatz, S.H.: Collective dynamics of 'small-world' networks. *Nature* 393(6684), 440–442 (1998)
27. Wernicke, S.: Efficient detection of network motifs. *IEEE/ACM Trans. Comput. Biol. Bioinformatics* 3(4), 347–359 (2006)

# Efficient Searching for Motifs in DNA Sequences Using Position Weight Matrices

Nikola Stojanovic

Department of Computer Science and Engineering  
University of Texas at Arlington  
Arlington, TX 76019, U.S.A.  
nick@cse.uta.edu

**Abstract.** Searching genomic sequences for motifs representing functionally important sites is a significant and well-established subfield of bioinformatics. In that context, Position Weight Matrices are a popular way of representing variable motifs, as they have been widely used for describing the binding sites of transcriptional proteins. However, the standard implementation of PWM matching, while not inefficient on shorter sequences, is too expensive for whole-genome searches. In this paper we present an algorithm we have developed for efficient matching of PWMs in long target sequences. After the initial pre-processing of the matrix it performs in time linear to the size of the genomic segment.

**Keywords:** DNA motifs, Position weight matrices, Genome-wide analysis, Algorithms, Genomics, Pattern matching.

## 1 Introduction

Eukaryotic genes are generally regulated in complex ways, through networks of protein–DNA and protein–protein interactions, which direct chromatin remodeling, histone modifications, formation of transcriptional initiation complexes and RNA Polymerase elongation. The prevailing opinion, corroborated by some studies [8,22] but being increasingly questioned [11,16], especially in the light of the somewhat surprising findings of the ENCODE Consortium [19], is that most of these interactions take place within a few hundred bases upstream of the transcription start sites. Even as regions important for the regulation of genes have been found at distal loci, around alternative first exons, in introns and sequences located downstream of the transcription start sites or even genes themselves, core promoter sites are still considered the most important for the gene expression. It is also generally accepted that the binding of transcriptional enzymes is in large part directed by specific motifs in DNA sequence. However, while there are proteins which bind only to exact layouts of bases, most transcription factors are rather non-specific in their choice of a binding site.

Consequently, the search for transcription factor binding sites has become one of the most popular subfields of bioinformatics, and many algorithms have been developed over about two decades of intensive research. The early approaches relied on a rather naive assumption that the motifs at the target sites of proteins must feature information content sufficient for their recognition, but disillusionment soon followed, as any



attempt to isolate these and other functional elements in DNA following this logic resulted in a very large number of false positives. Recent methods have thus concentrated on the incorporation of additional information to the raw sequence data, although they have so far neglected many important biochemical aspects. The additional information often relied on clustering of important motifs in putative regulatory modules, phylogenetic conservation or matching the sequence to databases of experimentally confirmed sites, such as TRANSFAC [21] or Jaspar [3].

Since many proteins important for transcriptional regulation bind with low specificity, the experimentally determined target motifs in DNA, of the same transcription factor, can be substantially different. Nevertheless, they often feature a well defined consensus sequence, with specific loci varying only marginally from the consensus. The extent of that variation is usually captured by Position Weight Matrices (PWMs), described in more detail below. Briefly, the PWMs record the information about the permissible variation of letters over the 4-letter DNA alphabet, assigning a weight to each letter at each position in the motif in accordance with how often that letter has been seen at that position within the experimentally determined binding sequences of the corresponding factor. The information about functional motifs in DNA recorded in the databases is nowadays usually in the form of PWMs, and they can be used for searches in previously uncharacterized DNA fragments, for motifs sufficiently similar to the consensus of a particular protein binding site. Over the years many programs have been written for matching PWMs, including MATCH [7] by the TRANSFAC database team, and variants such as combinations of PWMs using mixture model [5].

With the advances in microarray technology large sets of putatively co-expressed genes became available, stimulating the development of methods to detect conserved motifs in their upstream regions, such as [6], as well as the search for putatively co-regulated genes by the identification of shared regulatory modules, such as in [13]. It is intuitive that if a group of genes is coordinately regulated, it should be controlled by similar sets of transcription factors. From the hypothesis that protein binding is largely directed by target DNA sequence motifs it follows that same (or similar) motifs should be present in regulatory sequences of co-expressed genes, moreover as a cluster, or clusters. Consequently, this led to further exploitation of motif over-representation in sets of target sequences [2,20].

In the course of our work on the study of these genomic environments, we have developed software which extracts the most significant shared short (5–25 bases) approximate motifs found within the upstream sequences (i.e. putative promoter regions) of genes postulated to be co-expressed by microarray and other experiments [14]. We have recently expanded that work to genome-wide searches for similar layouts, as reflected by conglomerations of a statistically significant number of motifs previously discovered to be shared within the promoters of a training set of co-expressed genes [15]. Unfortunately, even as we were able to record the consensus of the over-represented motifs of interest in the form of PWMs, our genome-wide search had to be executed using string representation (standing for multiple exact patterns), since the current methods for PWM matching were not efficient. When the matching needs to be performed in limited environments representing gene upstream sequences, genomic domains or gene clusters, naive approaches work well, however for whole-genome scans any performance worse

than linear is not practical having in mind that, for instance, the human genome features more than three billion bases, and that it is not even the largest genome around.


In this paper we describe an algorithm we have devised to efficiently match patterns represented by Position Weight Matrices. It relies on a somewhat expensive pre-processing step, however the cost of that pre-processing is well compensated by the efficiency gains once it is applied to long target sequences, such as chromosomes or genomes. Before proceeding with the algorithm itself, we shall present some basic ideas governing the construction and use of the PWMs.

## 2 Position Weight Matrices

If all sites where a same protein binds would feature identical bases it would be a matter of simple pattern matching to find them all. Unfortunately, for gene regulatory proteins that is usually not the case, and one has to deal with approximate matching, consensus sequences and ambiguity codes. This is more than a technical complication, as it introduces substantial unreliability to motif recognition. The ambiguity codes do not distinguish between the bases they stand for, so rare or even impossible (concerning protein binding) combinations of characters in the represented motifs may match, leading to a large fraction of false positives. On the other hand, if a character which infrequently appears in the binding sequences is completely eliminated from the consensus, sites which do bind the protein may be skipped, adding poor sensitivity to already compromised specificity. Furthermore, binding motifs often contain irrelevant positions, and trying to match any particular character at these loci would obviously be futile.

A partial remedy to this problem can be an assignment of weights to characters of the consensus sequence, so that for each position where more than one base is possible there is a probability of occurrence associated with each choice. Stormo [18] thus proposed a matrix-based approach, where a position in a protein binding site is represented by a column of a matrix, with one row for each nucleotide. The guiding idea of the method was that if every character possible at a given position is assigned a score such that the addition of scores for all positions provides an estimate about how close the sequence is to the known protein binding site patterns, one can decide whether to accept or reject the site based on whether the cumulative score is above or below a certain predefined threshold. At any scanned sequence position  $j$  aligned with matrix column  $i$ , a simple lookup at the matrix entry at column  $i$  and row corresponding to the letter found at  $j$  would then provide the score for that position.

The threshold can be empirically determined as the critical value of the ratio of probabilities  $P_F(s)/P_G(s)$ , where  $s$  is the examined string, subscript  $F$  denotes that the probability is based on the distribution of strings within the population of binding sites, and subscript  $G$  denotes that the probability is based entirely on chance, given the overall distribution of nucleic acids in the DNA sequence under consideration. As the probabilities of the occurrence of  $s$  are based on two population models, this amounts to the likelihood ratio, as defined in statistics. It is also (somewhat unrealistically) assumed that character distributions at different positions are independent. That way, for a string  $s = c_1c_2 \dots c_L$ , of length  $L$ , the probability of its random occurrence would be  $P_G(s) = \prod_{i=1}^L p_g(c_i)$ , where  $p_g(c_i)$  stands for the general probability of an individual character  $c_i$ .



**Fig. 1.** An example of a Position Weight Matrix: a) A set of motifs on which the matrix is built; b) Letter counts at each position in the motif, in matrix representation; c) Full matrix, obtained by taking  $\log_2$  of the probability of the character occurrence at the position, divided by the general probability of the character, here assumed to be 0.25 for each of the ‘A’, ‘C’, ‘G’, ‘T’. ‘—’ indicates an infinitesimally small value.

In order to estimate  $P_F$  for a given string, one also assumes the independence of probabilities for character occurrences at different positions. For any position  $i$ , we estimate the probability of occurrence of any base  $p_{f,i}(b)$  based on how frequently it occurred at the particular position  $i$  in known binding patterns. Thus, having a string  $s = c_1c_2 \dots c_L$ , of length  $L$ , the probability that it occurred under the distribution specific for the particular protein binding sites would be  $P_F(s) = \prod_{i=1}^L p_{f,i}(c_i)$ . This approach can suffer from artifacts (for instance, if only combinations “AG” or “TC” would occur within motif instances used to construct the matrix, once it is constructed it may assign equally high scores to combinations of “AC” and “TG” which have never been observed), which has led to the development of di-nucleotide [4] and even more elaborate models. It can also be imprecise due to incomplete experimental data or overly permissible thresholds. However, over many years of PWM application this method has been shown to yield better results than raw pattern matching, either exact or approximate. If it is only a relatively small number of motifs used for the construction of a PWM, it may be more appropriate to construct a finite state automaton which would recognize *all* these motifs [1], rather than approximating through a PWM, but that then introduces the risk of over-fitting.

Expanded, the likelihood ratio  $P_F(s)/P_G(s)$  is then  $\prod_{i=1}^L p_{f,i}(c_i) / \prod_{i=1}^L p_g(c_i) = \prod_{i=1}^L [p_{f,i}(c_i)/p_g(c_i)]$ . One can take a logarithm of this formula and convert it to additive form  $\sum_{i=1}^L \log_2 [p_{f,i}(c_i)/p_g(c_i)]$ , although some correction factors are necessary in practice, in order to avoid taking a logarithm of zero when a character does not appear at a given position in any of the motifs used to construct the matrix. The matrix associated with the protein binding pattern of length  $L$  has  $L$  columns and 4 rows, one for each nucleotide ‘A’, ‘C’, ‘G’ or ‘T’. If a base  $b$  occurs  $m_b$  times at the position corresponding to matrix column  $i$  within  $M$  known binding sites, then  $p_{f,i}(b) = m_b/M$ , and if it occurs  $n_b$  times within  $N$  nucleotides in the genome, then  $p_g(b) = n_b/N$ . The matrix entry for base  $b$  in column  $i$  thus contains  $\log_2(m_bN/n_bM)$ . An illustration of a set of motifs used to determine a PWM, and the resulting matrix, is shown in Figure 1. In practice there are several variants of PWMs, however they are all based on the same principle.

Once a PWM is constructed, it can be used for scanning genomic sequences, for motifs which are sufficiently similar to these captured by the matrix. Looking at a string  $s = c_1c_2 \dots c_L$  one locates the corresponding row of the matrix for each  $c_i$ , in column  $i$ , and adds these values to obtain the likelihood ratio  $P_F/P_G$ . If that likelihood ratio is above the pre-determined (and inevitably heuristic) threshold, a match is declared. Since the shift for just one position can result in a dramatically different score, PWM matching programs advance for one position, and start the matching process over the entire matrix every time after the shift has been made. As matrices usually do not feature too many columns, reflecting the fact that they model short motifs, this does not lead to intractability, yet it slows down any search for an order of magnitude, and on the genomic scale that presents a problem. Pruning techniques have been explored, although one can only stop further matching attempts from a position if it has been determined that the score-so-far is insufficient for a match. The other option, recognizing a match before all characters have been examined is not possible, since some letters at some positions may contribute (large) negative scores. Other algorithmic solutions have also been applied, such as building indexing schemes to facilitate matching of multiple matrices in the same sequence [10].


### 3 Algorithm

We here present an algorithm which matches a single Position Weight Matrix against a genomic sequence in linear time.

We start by pre-processing the matrix  $\mathcal{M}$ , of dimension  $4 \times L$ , into a tree-like structure, following the basic ideas of classic algorithms of Knuth–Morris–Pratt [9] and Aho–Corasick [1]. For every position  $p$  our algorithm keeps track of its longest possible suffix  $[k..p]$ ,  $1 < k < p$ , in  $\mathcal{M}$ , which can still lead to a match, and thus the amount of shift which can be made after  $\mathcal{M}$  has been successfully matched, or when it has been determined that a match at the current starting position, represented by the root, is not possible.

In order to enable the tracking, the matrix  $\mathcal{M}$  is thus converted into a finite state automaton, implementing a tree  $\mathcal{T}$  with cross-linked branches. At every node of the tree a structure is maintained, with the following fields:

- level:** Keeps track of the length of the branch starting at the root and ending at the current node, following the most direct (i.e. downward) path. This information is used to determine whether a match along the current branch is still possible, and to check whether a match has already been successful (when a node whose level equals  $L$  is reached).
- score:** Keeps track of the score achieved so far along the path from the root through the current node. If the score is less than required for the current node level it is an indication that no match is possible at the position currently aligned with the root and a shift to a suffix has to be made; if a match has been achieved, this score can be reported along with match data.
- suffix:** A pointer to another node in  $\mathcal{T}$ , positioned on the branch which represents a suffix of the branch ending in the current node (the longest suffix which can still



**Fig. 2.** An example of an FSA created for a simple PWM: a) Position weight matrix; b) Minimal scores which have to be reached at every position in order for threshold value of 9 to be met at the end (a match); c) FSA constructed based on the matrix. Only three strings, “CAA”, “CTA” and “CTT”, would score sufficiently high to meet the threshold, as reflected by three paths from the root to the reporting nodes (level 3), shown as solid circles. Scores at each node are shown inside the circles. Transitions are guided by the letters labeling the edges.

lead to a viable match from its starting position). The target node of this pointer follows the transition on the same letter as the current node<sup>1</sup>.

**move:** A four-element array, providing a pointer to another node in  $\mathcal{T}$ , directing the move on each of the four letters of the DNA alphabet (‘A’, ‘C’, ‘G’ and ‘T’). It can lead to forward motion if a match along the current branch is still possible, or backward to another path starting from the root (a suffix) if the match at the currently examined position cannot be achieved.

The preprocessing starts with the establishment of a score array *minimum*, whose dimension  $L$  equals the length of the motif represented by  $\mathcal{M}$ . Each entry in *minimum* records the minimal score which must be achieved at that position in order to lead to a possible match. Since a match is defined as a sequence scoring at *threshold* value or higher, it must be  $minimum[L] = threshold$ . If  $max[i]$  is the maximal score recorded in the  $i$ -th column of  $\mathcal{M}$ , then  $minimum[i - 1] = minimum[i] - max[i]$ . Whenever a position  $p$  in  $\mathcal{M}$  is reached it is checked whether the score so far is still greater than or equal to  $minimum[p]$ .

The FSA  $\mathcal{T}$  is constructed in the breadth-first fashion, after the root node has been formed, and its immediate children have been placed in the queue (if their score warranted the placement). Every time a node is dequeued, it is first checked if its level  $l$  equals  $L$ . If that is the case (indicating time to report a match and move on), its

<sup>1</sup> In our implementation we have kept this field outside of the FSA, as it is used only once during the pre-processing of a matrix, in situations when the branch needs to be changed on an occurrence of a letter.

**Algorithm 3.1.** PREPROCESS(*matrix*, *dimension*, *threshold*)

```

// Find the minimal values necessary at positions
minimum[dimension] ← threshold
for i ← dimension - 1 downto 0
  do { max ← Maximal value in matrix column i + 1
      { minimum[i] ← minimum[i + 1] - max
//Process the root
root.level ← 0
root.score ← 0
root.suffix ← root
for letter ← A,C,G,T
  do { score ← matrix[letter][1]
      { if score < minimum[1]
        { then root.move[letter] ← root
          { child.score ← score
            { child.level ← 1
              { else { child.suffix ← root
                    { root.move[letter] ← child
                      { ENQUEUE(child)
// Process the remaining nodes, breadth-first
while Queue not empty
  { current ← DEQUEUE()
  { if current.level = dimension
    { then { for letter ← A,C,G,T
            { do current.move[letter] ← current.suffix.move[letter]
              { for letter ← A,C,G,T
                { do { score ← current.score + matrix[letter][current.level + 1]
                    { if score < minimum[current.level + 1]
                      { then current.move[letter] ← current.suffix.move[letter]
                        { child.score ← score
                          { child.level ← current.level + 1
                            { else { child.suffix ← current.suffix.move[letter]
                                  { current.move[letter] ← child
                                    { ENQUEUE(child)
return (root)

```

**Algorithm 3.2.** MATCH(*matrix*, *text*, *threshold*)

```

track ← PREPROCESS(matrix, matrix_column_count, threshold)
for i ← 1 to text_length - matrix_column_count + 1
  do { track ← track.move[text[i]]
      { if track.level = matrix_column_count
        { then Report a match at i - matrix_column_count + 1, scoring track.score

```

**Fig. 3.** Pseudo-code of the PWM matching algorithm

transitions on all four letters are determined by following the transition on that letter from the (suffix) node pointed to by the *suffix* pointer. If  $l < L$ , forward transitions (and the increment of the level) are considered. If the score of the transition on a letter falls below the minimum required at  $l$ , a back pointer is created to the extension from the node pointed to by the *suffix*. Otherwise, a child node is created and enqueued, defining a forward transition on the letter. An example of an FSA created for a simple PWM is provided in Figure 2 and the pseudo-code of the preprocessing algorithm is provided in the function `Preprocess` in Figure 3. The complete matching function is shown as function `Match`.

After the preprocessing is completed, the matrix matching itself is rather straightforward. The algorithm maintains two pointers, one to the currently examined position  $i$  in the DNA sequence (initialized at 1), and the other to the current node in  $\mathcal{T}$  (initialized at the root). A transition in  $\mathcal{T}$  is done depending on the character at position  $i$ , and if the level of the reached node equals  $L$  a match can be reported (with appropriate score). This process continues until the end of the sequence is reached.

## 4 Algorithm Performance

Since the matching done by this algorithm is completely guided by the FSA  $\mathcal{T}$ , it can be proven correct by observing that every matching substring would lead to a traversal of a path from the root to a node at level  $L$ , and that no mismatching substring would reach level  $L$ . Since every time an insufficient score is obtained it would result in a stagnation or decrease in the current level, and every time a sufficient score is obtained it results in an increase of the current level, the above properties will hold if it can be proven that the back pointers indeed lead to a node  $N$  such that the direct path from the root to  $N$  represents the longest feasible suffix of the currently examined substring.

The formal proof can be derived by mathematical induction, showing that the following invariants hold every time a node is dequeued and processed during the matrix pre-processing phase:

1. If there is a viable continuation from the position represented by the node, on a particular letter, that leads to the creation of a forward link, i.e. a node at the next deeper level.
2. If there is no continuation from the position represented by the node, for each letter, that leads to the creation of a backward link, to a node at the same level as current, or closer to the root.
3. The *suffix* pointer of a newly enqueued node points to another node in the FSA, at some level closer to the root, which is the end of a path from the root representing the longest suffix of the string represented by the path to the enqueued node, and which can still yield a match.
4. If a path cannot be continued to a further level, on some letter, then the transition on that letter is made to a node representing the end of the longest suffix of the current path which can still yield a match.

For the induction base, it is trivial to show that properties 1 through 4 hold at the time the root is processed (initialization), and its children are being placed in the queue.

We can assume that they are also satisfied when  $k^{th}$  node is being processed, and show that they hold after the  $k + 1^{th}$  has been handled. It should be noted that the *suffix* pointers are already fixed at the time of enqueueing the node.

Invariant 1 holds trivially, as it is the decision made in the code while processing the dequeued node. Since the tree is processed breath-first, that means that all nodes at depths smaller than that of the current one have already been processed. If there is no viable continuation of a match on a currently considered letter, that leads to following the transition from the node pointed to by the *suffix* pointer, which is, by invariant 3, at a level closer to the root, and thus, by invariants 1 and 2 being already satisfied at earlier nodes, cannot lead to level deeper than the current. This proves that invariant 2 also holds after the current node is processed.

If there is a viable continuation from the current node on any particular letter, then the *suffix* pointer assigned to the newly created child is to the link of the node at the other end of the current *suffix* pointer, and thus at the level closer to the root than the newly created child. If the node pointed to by the current *suffix* featured a viable forward move on the letter, that extends that suffix which can yield a match. If not, it extends its longest suffix which can, by the induction hypothesis, and thus represents a new longest suffix which can still yield a match. Therefore, invariant 3 also holds.

When a forward move (one continuing the match) cannot be done on a particular letter, it is being done from the node pointed to by the *suffix* pointer, which, by invariant 3, represents the longest suffix of the current path which can still yield a match. Therefore, the transition is being made to the node extending the other end of the *suffix* pointer on that letter, and the invariant 4 holds.

During the search phase the traversal of the FSA  $\mathcal{T}$  proceeds in the forward direction (i.e. to deeper levels) for as long as a match is still possible (invariant 1). When it is no longer the case, invariant 4 guarantees that the jump is being made to the next position from which there may be a possible match. As a determined mismatch cannot lead to a forward move in  $\mathcal{T}$  no mismatching positions can lead to reaching level  $L$  (and thus be reported), and as the jump is being made to the beginning of the longest suffix of the currently examined substrings which can still yield a match (invariant 4), no matching positions can be skipped. We thus conclude that the algorithm is correct.

## 4.1 Space Performance

This algorithm can have large space requirements. For a matrix of length  $L$  it can theoretically build an FSA with  $O(4^L)$  nodes, since every possible matching substring has to be represented by a path through  $\mathcal{T}$  reaching level  $L$ . However, all internal nodes, and thus their children as well, for which it has been established that they could not lead to a match, are promptly pruned. Since in practice PWMs are constructed, and their thresholds set, so that they accept a very limited number of variants of the target site, it is expected that in most practical situations pruning will be quite dramatic. The only situation when there can exist a large conglomeration of nodes towards the top of the tree, the most space-consuming setting, is when there is a variety of possible characters occurring in relatively high percentages at the left hand side of the pattern, thus dictating extensive branching near the root of the tree, and a more inclusive threshold value. However, this situation practically never happens, since tails exhibiting



substantial variations normally do not get included in the patterns captured by PWMs. Towards the bottom of the tree the only branches which have not been pruned would be these leading to a pattern variant matching the matrix (as set by the given threshold).

Apart from the FSA encoding the matrix, the only space requirement for the matching process itself is the space needed to store the sequence in which the matching is to be done.

## 4.2 Time Performance

Since, during the matching, no character in the text is examined more than once (as signified by the single `for`-loop with constant-time body in the algorithm in Figure 3), this process clearly executes in linear time. The construction of the FSA (pre-processing) takes time proportional to the number of nodes which are being created (as each node is being processed in constant time), which, under the worst-case scenario of a large number of matching strings, can be exponential. However, since the children of the pruned nodes never get processed, the exponential blowout is expected to take place only at shallow depth, towards the top of the tree. Furthermore, whereas pre-processing of the long PWMs would take the longest time, these are precisely the cases where the gain from subsequent linear-time matching would be the largest.

Even as the PWMs are generally short (rarely exceeding 15 columns, and almost never exceeding 25) and pruning extensive (as the thresholds are usually very stringent), the pre-processing step has the potential of annihilating the gains from the subsequent linear-time processing. This algorithm is thus recommended for use only in very large scale analyses, such as our whole-genome scans [15]. In our recent tests, done on a genomic segment of about 500 million bases (corresponding to a large chromosome), the naive match was taking an average of 144.33 seconds per matrix, on an Apple Mac Pro 2.5 GHz Intel Core Duo with 4 Gb 667 MHz DDR2 SDRAM memory laptop computer, while an implementation of our algorithm was taking an average of 16.58 seconds per matrix. On whole-genome scale the average gain would be about 10–20 fold per matrix.

## 5 Discussion

The space requirements of this algorithm can be further reduced by a simple extension of the pre-processing step, alas at the price of an increase in the computational time. This would require the maintenance of two additional fields at the nodes of the FSA: *parent pointer* and *number of forward links*. An additional traversal pass through the FSA can then eliminate all partial paths which do not extend until the bottom (match), reducing the number of nodes to less than the sum of the lengths of all strings which would score above the threshold. However, we have not implemented this modification, as the space requirements of the FSA have never been a limiting factor in practice, and time was critical.

Often a PWM will feature letters which must not be found at a specific position, with the corresponding score in the PWM of (theoretically)  $-\infty$ . This situation is particularly favorable to our algorithm, as it leads to immediate pruning, especially when these

positions are concentrated towards the beginning of the matrix. If they are concentrated towards the end, the least favorable setting, one can do the matching in the inverse complement of the sequence. For long matrices one can also attempt the match of their most specific core, then attempt the full match, using the naive approach, only around the positions where the core has matched.

Despite of its relatively expensive pre-processing, we have found this algorithm useful for whole-genome scans, such as our search for the conglomerations of variable motifs, with a potential of reducing days of computation to just a few hours. This can be of particular importance for the tools implemented as a part of a web server. Our earlier version of the matcher [15] implemented at <http://bioinformatics.uta.edu/toolkit/motifs/> used direct pattern matching (i.e. not based on PWMs), and the development of this algorithm has allowed us to consider the matrix-based approach, too.

After the original version of this manuscript has been reported [17], it has come to our attention that at around the same time we have developed this algorithm another group has published essentially the same result, although in a broader context [12]. However, it is our opinion that our presentation, if not approach, is better structured, yielding itself to a straightforward software construction. This is unlike the exposition of Pizzi *et al.*, which, although sound, may require some “reading between the lines”.

**Acknowledgements.** The author is grateful to Abanish Singh, whose effort on motif finding and the implementation of the whole-genome motif search has made us aware of the need for this algorithm. This work has been partially supported by NIH grant 5R03LM009033-02.

## References

1. Aho, A., Corasick, M.: Efficient string matching: an aid to bibliographic search. *Comm. Assoc. Comput. Mach.* 18, 333–340 (1975)
2. Apostolico, A., Bock, M., Lonardi, S., Xu, X.: Efficient detection of unusual words. *J. Comput. Biol.* 7, 71–94 (2000)
3. Bryne, J., Valen, E., Tang, M., Marstrand, T., Winther, O., da Piedade, I., Krogh, A., Lenhard, B., Sandelin, A.: JASPAR, the open access database of transcription factor-binding profiles: new content and tools in the 2008 update. *Nucleic Acids Res.* 36, D102–D106 (2008)
4. Gershenzon, N.I., Stormo, G.D., Ioshikhes, I.P.: Computational technique for improvement of the position-weight matrices for the DNA/protein binding sites. *Nucleic Acids Res.* 33, 2290–2301 (2005)
5. Hannenhalli, S., Wang, L.S.: Enhanced position weight matrices using mixture models. *Bioinformatics* 21, i204–i212 (2005)
6. Hughes, J., Estep, P., Tavazoie, S., Church, G.: Computational identification of *cis*-regulatory elements associated with groups of functionally related genes in *Saccharomyces cerevisiae*. *J. Mol. Biol.* 296, 1205–1214 (2000)
7. Kel, A.E., Gössling, E., Reuter, I., Chermushkin, E., Kel-Margoulis, O.V., Wingender, E.: Match: A tool for searching transcription factor binding sites in dna sequences. *Nucleic Acids Res.* 31(13), 3576–3579 (2003), <http://dx.doi.org/10.1093/nar/gkg585>
8. Khambata-Ford, S., Liu, Y., Gleason, C., Dickson, M., Altman, R., Batzoglou, S., Myers, R.: Identification of promoter regions in the human genome by using a retroviral plasmid library-based functional reporter gene assay. *Genome Res.* 13, 1765–1774 (2003)

9. Knuth, D., Morris, J., Pratt, V.: Fast pattern matching in strings. *SIAM J. Computing* 6, 323–350 (1977)
10. Liefvooghe, A., Touzet, H., Varré, J.S.: Large Scale Matching for Position Weight Matrices. In: Lewenstein, M., Valiente, G. (eds.) *CPM 2006*. LNCS, vol. 4009, pp. 401–412. Springer, Heidelberg (2006)
11. Nelson, C., Hersh, B., Carroll, S.B.: The regulatory content of intergenic DNA shapes genome architecture. *Genome Biol.* 5, R25 (2004)
12. Pizzi, C., Rastas, P., Ukkonen, E.: Finding significant matches of position weight matrices in linear time. *IEEE/ACM Transactions on Computational Biology and Bioinformatics E-publication ahead of print* (2009)
13. Qin, Z., McCue, L., Thompson, W., Mayerhofer, L., Lawrence, C., Liu, J.: Identification of co-regulated genes through Bayesian clustering of predicted regulatory binding sites. *Nature Biotechnology* 21, 435–439 (2003)
14. Singh, A., Stojanovic, N.: An efficient algorithm for the identification of repetitive variable motifs in the regulatory sequences of co-expressed genes. In: Levi, A., Savaş, E., Yenigün, H., Balcısoy, S., Saygın, Y. (eds.) *ISCIS 2006*. LNCS, vol. 4263, pp. 182–191. Springer, Heidelberg (2006)
15. Singh, A., Stojanovic, N.: Genome-wide search for putative transcriptional modules in eukaryotic sequences. In: *Proceedings of BIOCOMP 2009*, pp. 848–854 (2009)
16. Stojanovic, N.: A study on the distribution of phylogenetically conserved blocks within clusters of mammalian homeobox genes. *Genetics and Molecular Biology* 32, 666–673 (2009)
17. Stojanovic, N.: Linear-time matching of position weight matrices. In: *Proceedings of the First International Conference on Bioinformatics, BIOINFORMATICS 2010*, pp. 66–73 (2010)
18. Stormo, G.: Consensus patterns in DNA. *Methods Enzym.* 183, 211–221 (1990)
19. The ENCODE Project Consortium: The ENCODE pilot project: Identification and analysis of functional elements in 1% of the human genome. *Nature* 447, 799–816 (2007)
20. van Helden, J.: Metrics for comparing regulatory sequences on the basis of pattern counts. *Bioinformatics* 20, 399–406 (2004)
21. Wingender, E.: The TRANSFAC project as an example of framework technology that supports the analysis of genomic regulation. *Briefings in Bioinformatics* 9, 326–332 (2008)
22. Young, J.E., Vogt, T., Gross, K.W., Khani, S.C.: A short, highly active photoreceptor-specific enhancer/promoter region upstream of the human rhodopsin kinase gene. *Investigative Ophthalmology and Visual Science* 44, 4076–4085 (2003)

# Author Index

- Alvarez, Sergio A. 29  
Amaro Jr., Edson 278  
Ansari, Mohd. Zahid 166  
Astikainen, Katja 367
- Bates, Marsha E. 224  
Biswal, Devendra Kumar 309  
Bleiker, Andreas 111  
Boba, Patrick 356  
Brown, Edward 70  
Buckman, Jennifer F. 224  
Busch, Clemens 111
- Cabitzza, Federico 44  
Cardoso, Ellison Fernando 278  
Castellano, Gabriela 278  
Cho, Chongdu 166  
Clifton, David A. 199  
Cruz, Antonio Barrientos 152  
Cunha, João Paulo Silva 290
- Das, Amar K. 97  
da Silva, Elvis Lira 278  
DeBoer, Friso 266  
de Lannoy, Gael 212  
Delbeke, Jean 212
- Eggenberger Hotz, Peter 141  
Elgendi, Mohamed 266
- Faliagka, Evanthia 15  
Faubel, Raquel 3  
Floyd, Stuart 29  
Fraile, Rubén 238  
François, Damien 212  
French, Paddy 189
- Geraci, F. 322  
Gesso, Iade 44  
Godino-Llorente, J. Ignacio 238  
Goodyear, Jamie 70  
Grandi, Fabio 57  
Gutiérrez-Arriola, Juana M. 238
- Hadorn, Maik 141  
Hamacher, Kay 356  
Hoffgaard, Franziska 356  
Holm, Liisa 367  
Hugueny, Samuel 199
- Johannesson, Paul 123  
Jonkman, Mirjam 266
- Klompmaker, Florian 111  
Kob, Malte 238  
Koufi, Vassiliki 83
- Lopes, Luís 380
- Mailund, Thomas 334  
Malamateniou, Flora 83  
Malsawmtluangi, Chenkual 309  
Matsopoulos, George 15  
Meyer, Jean-Arcady 252
- Nebe, Karsten 111  
N'Guyen, Steve 252
- O'Connor, Martin J. 97  
Osma-Ruiz, Víctor 238
- Pandina, Robert J. 224  
Pandraud, Gregory 189  
Pedersen, Christian N.S. 334  
Pellegrini, M. 322  
Pérez, José Luis Martínez 152  
Perjons, Erik 123  
Pirim, Patrick 252  
Pitkänen, Esa 367  
Prasad, Pramod Kumar 309  
Purniawan, Agung 189
- Quintela, Bruno 290
- Renevey, Philippe 177  
Ribeiro, Pedro 380  
Rousu, Juho 367  
Ruiz, Carolina 29

- Sáenz-Lechón, Nicolás 238  
Sarro, Pasqualina M. 189  
Sato, João Ricardo 278  
Schleusing, Olaf 177  
Schweizer, Valérie 177  
Silva, Fernando 380  
Simonsen, Martin 334  
Stojanovic, Nikola 394  
Szedmak, Sandor 367
- Tandon, Veena 309  
Tarassenko, Lionel 199  
Traver, Vicente 3  
Tsakalidis, Athanasios 15  
Tsaknakis, John 15  
Tseng, Jennifer 29  
Tzimas, Giannis 15
- Vaschillo, Bronya 224  
Vaschillo, Evgeny G. 224  
Vassilacopoulos, George 83  
Verleysen, Michel 212  
Vesin, Jean-Marc 177  
Vetter, Rolf 177
- Weil, Philipp 356  
Whalen, Giles 29  
Willemsen, Detlev 111  
Wong, Thomas K.F. 345
- Yiu, S.M. 345
- Zúquete, André 290