

Bruno Apolloni
Simone Bassis
Anna Esposito
Francesco Carlo Morabito
Editors

SMART INNOVATION,
SYSTEMS AND TECHNOLOGIES ■ 19



Neural Nets and Surroundings

22nd Italian Workshop on Neural Nets,
WIRN 2012, May 17–19, Vietri sul Mare,
Salerno, Italy



 Springer

Editors-in-Chief

Prof. Robert J. Howlett
KES International
PO Box 2115
Shoreham-by-sea
BN43 9AF
UK
E-mail: rjhowlett@kesinternational.org

Dr. Lakhmi C. Jain
Adjunct Professor
University of Canberra
ACT 2601
Australia
and
University of South Australia
Adelaide
South Australia SA 5095
Australia
E-mail: Lakhmi.jain@unisa.edu.au

Bruno Apolloni, Simone Bassis, Anna Esposito,
and Francesco Carlo Morabito (Eds.)

Neural Nets and Surroundings

22nd Italian Workshop on Neural Nets,
WIRN 2012, May 17–19, Vietri sul Mare,
Salerno, Italy

 Springer

Editors

Prof. Bruno Apolloni
Department of Computer Science
University of Milano
Milano
Italy

Dr. Simone Bassis
Department of Computer Science
University of Milano
Milano
Italy

Prof. Anna Esposito
Department of Psychology
Second University of Naples
Caserta
Italy

and
Institute for Advanced Scientific
Studies (IIASS)
Vietri sul Mare Salerno
Italy

Prof. Francesco Carlo Morabito
Department of Mechanics and Materials
Mediterranea University of Reggio Calabria
Reggio Calabria
Italy

ISSN 2190-3018

e-ISSN 2190-3026

ISBN 978-3-642-35466-3

e-ISBN 978-3-642-35467-0

DOI 10.1007/978-3-642-35467-0

Springer Heidelberg New York Dordrecht London

Library of Congress Control Number: 2012953656

© Springer-Verlag Berlin Heidelberg 2013

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

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

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

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

Preface

This volume collects a selection of contributions which has been presented at the 22nd Italian Workshop on Neural Networks, the yearly meeting of the Italian Society for Neural Networks (SIREN). The conference was held in Italy, Vietri sul Mare (Salerno), during May 17–19, 2012. The annual meeting of SIREN is sponsored by International Neural Network Society (INNS), European Neural Network Society (ENNS) and IEEE Computational Intelligence Society (CIS).

The workshop, and thus this book, is organized in three main components, two special sessions and a group of regular sessions featuring different aspects and point of views of artificial neural networks and natural intelligence, also including applications of present compelling interest.

More than 60 papers were presented at the Workshop, and most of them are reported here. The review process has been carried out in two steps, one before and one after the workshop in order to meet Publisher's requirements. The selection of the papers was made through peer-review process, where each submission was evaluated by at least two reviewers. The submitted papers were authored by peer scholars from different countries (the Italian component was anyway preponderant). The acceptance rate is thus high also because most of the attendees are involved in SIREN research and organization activities for more than 20 years. In addition to regular papers, the technical program featured keynote plenary lectures by some worldwide renowned scientist (Soo Young Lee, South Korea; Ganesh K. Venayagamoorthy, USA; Jacek Zurada, USA; Günther Palm, Germany; Alessandro Vinciarelli, UK; Danilo Mandic, UK). One of the two special sessions was supported by the EU-sponsored COST Action 2102 that closed his work on February 2011 even though the Members of the Action are still networking and collaborating in scientific activities.

The first Special Session explored the new frontiers and challenges in Smart Grid research and proposed a proficient discussion table for scientists joining the WIRN conference, whose expertise typically cover the research fields addressed in Smart Grid technology, as electrical and electronic engineering, computational intelligence, digital signal processing and telecommunications. The Session included two invited contributions and seven regular ones. The Session was particularly relevant because it introduced

some aspects of neural network applications not commonly known at the community in a field of growing interest.

The second Special Session was titled Computational Intelligence in Emotional or Affective Systems and was given in honour of John Taylor, the Editor-in-chief of the journal Neural Networks recently died. The Session featured two keynote lectures and 10 regular contributions. Computational Intelligence (CI) methods have shown great capabilities in modelling, prediction, and recognition tasks and a mature degree of understanding has been achieved in many application areas, in particular in complex multimodal systems supporting human-machine or human-human interaction. At the same time, the emotional issue has recently gained increasing attention in such complex systems due to its relevance in most common human tasks (like cognitive processes, perception, learning, communication and even “rational” decision-making) and therefore is highly relevant for the goal of human-like interaction with machines. The real challenge is taking advantage of the emotional characterization of humans to make the computer interfacing with them more natural and therefore useful. The scope of the session was to assess to what extent and how sophisticated computational intelligence tools developed so far might support the multidisciplinary research on the characterization of an appropriate system reaction to human emotions and expression in interactive scenarios.

We would like to thank all of the special sessions organizers, namely: Stefano Squartini, Rosario Carbone, Michele Scarpiniti, Francesco Piazza, Aurelio Uncini, Anna Esposito, Günther Palm.

The organization of an International Conference gathers for the efforts of several people involved. We would like to express our gratitude to everyone that has cooperate to the organization, by offering their commitment, energy and spare time to make this event a successful one.

May 2012

Bruno Apolloni
Simone Bassis
Anna Esposito
Francesco Carlo Morabito

Organization

WIRN 2012 is organized by the Italian Society of Neural Networks (SIREN) in cooperation with the International Institute for Advanced Scientific Studies (IIASS) of Vietri S/M (Italy).

Executive Committee

Bruno Apolloni	University of Milano, Italy
Simone Bassis	University of Milano, Italy
Anna Esposito	University Federico II of Napoli, Italy
Francesco Masulli	University of Genova, Italy
Francesco Carlo Morabito	University Mediterranea of Reggio Calabria, Italy
Francesco Palmieri	Second University of Napoli, Italy
Eros Pasero	Polytechnic of Torino, Italy
Stefano Squartini	Polytechnic University of Marche, Italy
Roberto Tagliaferri	University of Salerno, Italy
Aurelio Uncini	University "La Sapienza" of Roma, Italy
Salvatore Vitabile	University of Palermo, Italy

Program Committee

Conference Chair

Francesco Carlo Morabito	University Mediterranea of Reggio Calabria, Italy
--------------------------	--

Conference Co-Chair

Simone Bassis	University of Milan, Italy
---------------	----------------------------

Program Chair

Bruno Apolloni	University of Milan, Italy
----------------	----------------------------

Organizing Chair

Anna Esposito

Second University of Napoli, Italy

Special Tracks

Anna Esposito

Second University of Napoli, Italy

Stefano Squartini

Polytechnic University of Marche, Italy

Referees

G. Albano

S. Funari

M. Re

B. Apolloni

C. Furlanello

A. Rizzi

S. Bassis

G. L. Galliani

P. M. Ros

A. Borghese

S. Giove

S. Rovetta

F. Camastra

G. Ippoliti

A. Rozza

W. Capraro

F. La Foresta

M. Russolillo

R. Carbone

G. Lombardi

S. Scarpetta

M. Cardin

M. Lucchese

M. Scarpiniti

A. Ciaramella

D. Malchiodi

R. Serra

C. Claudio

U. Maniscalco

G. Spagnuolo

D. Communiello

C. Marco

S. Squartini

V. d'Amato

F. Masulli

A. Staiano

R. de Rosa

L. Menconi

A. Uncini

F. Epifania

A. Micheli

G. Valentini

A. M. Esposito

F. C. Morabito

L. Valerio

A. Esposito

G. Palm

M. Villani

M. Faundez-Zanuy

F. Palmieri

S. Vitabile

A. Filisetti

E. Pasero

Q. Wei

M. Frasca

F. Piazza

A. Zippo

Sponsoring Institutions

International Institute for Advanced Scientific Studies (IIASS) of Vietri S/M (Italy)

Department of Psychology, Second University of Napoli (Italy)

Provincia di Salerno (Italy)

Comune di Vietri sul Mare, Salerno (Italy)

Contents

Part I: Algorithms

Probability Learning and Soft Quantization in Bayesian Factor Graphs	3
<i>Francesco A.N. Palmieri, Alberto Cavallo</i>	
Rival-Penalized Competitive Clustering: A Study and Comparison	11
<i>Alberto Borghese, William Capraro</i>	
An Interpretation of the Boundary Movement Method for Imbalanced Dataset Classification Based on Data Quality	21
<i>Dario Malchiodi</i>	
Genetic Algorithm Modeling with GPU Parallel Computing Technology . . .	29
<i>Stefano Cavuoti, Mauro Garofalo, Massimo Brescia, Antonio Pescapé, Giuseppe Longo, Giorgio Ventre</i>	
An Experimental Evaluation of Reservoir Computation for Ambient Assisted Living	41
<i>Davide Bacciu, Stefano Chessa, Claudio Gallicchio, Alessio Micheli, Paolo Barsocchi</i>	
Balancing Recall and Precision in Stock Market Predictors Using Support Vector Machines	51
<i>Marco Lippi, Lorenzo Menconi, Marco Gori</i>	
Measures of Brain Connectivity through Permutation Entropy in Epileptic Disorders	59
<i>Domenico Labate, Giuseppina Inuso, Gianluigi Occhiuto, Fabio La Foresta, Francesco C. Morabito</i>	
A New System for Automatic Recognition of Italian Sign Language	69
<i>Marco Fagiani, Emanuele Principi, Stefano Squartini, Francesco Piazza</i>	

Fall Detection Using an Ensemble of Learning Machines 81
Simon Bulotta, Hassan Mahmoud, Francesco Masulli, Ernesto Palummeri, Stefano Rovetta

Part II: Signal Processing

PM₁₀ Forecasting Using Kernel Adaptive Filtering: An Italian Case Study 93
Simone Scardapane, Danilo Comminiello, Michele Scarpiniti, Raffaele Parisi, Aurelio Uncini

A Collaborative Filter Approach to Adaptive Noise Cancellation 101
Michele Scarpiniti, Danilo Comminiello, Raffaele Parisi, Aurelio Uncini

Waveform Variation of the Explosion-Quakes as a Function of the Eruptive Activity at Stromboli Volcano 111
Antonietta M. Esposito, Luca D’Auria, Flora Giudicepietro, Marcello Martini

Artificial Neural Network (ANN) Morphological Classification of Magnetic Resonance Imaging in Multiple Sclerosis 121
Alessia Bramanti, Lilla Bonanno, Placido Bramanti, Pietro Lanzafame

Neural Moving Object Detection by Pan-Tilt-Zoom Cameras 129
Alessio Ferone, Lucia Maddalena, Alfredo Petrosino

Control of Coffee Grinding with General Regression Neural Networks 139
Luca Mesin, Diego Alberto, Eros Pasero

Defects Detection in Pistachio Nuts Using Artificial Neural Networks 147
Paolo Motto Ros, Eros Pasero

Part III: Applications

LVQ-Based Hand Gesture Recognition Using a Data Glove 159
Francesco Camastra, Domenico De Felice

Investigation of Single Nucleotide Polymorphisms Associated to Familial Combined Hyperlipidemia with Random Forests 169
Antonino Staiano, Maria Donata Di Taranto, Elena Bloise, Maria Nicoletta D’Agostino, Antonietta D’Angelo, Gennaro Marotta, Marco Gentile, Fabrizio Jossa, Arcangelo Iannuzzi, Paolo Rubba, Giuliana Fortunato

A Neural Procedure for Gene Function Prediction 179
Marco Frasca, Alberto Bertoni, Andrea Sion

Handwritten Digits Recognition by Bio-inspired Hierarchical Networks 189
Antonio G. Zippo, Giuliana Gelsomino, Sara Nencini, Gabriele E.M. Biella

Forecasting Net Migration by Functional Demographic Model	201
<i>Valeria D'Amato, Gabriella Piscopo, Maria Russolillo</i>	
Simulation Framework in Fertility Projections	209
<i>Valeria D'Amato, Gabriella Piscopo, Maria Russolillo</i>	
Building a Global Performance Indicator to Evaluate Academic Activity Using Fuzzy Measures	217
<i>Marta Cardin, Marco Corazza, Stefania Funari, Silvio Giove</i>	
Testing the Weak Form Market Efficiency: Empirical Evidence from the Italian Stock Exchange	227
<i>Giuseppina Albano, Michele La Rocca, Cira Perna</i>	
Part IV: Special Session on “Smart Grids: New Frontiers and Challenges”	
Real Time Techniques and Architectures for Maximizing the Power Produced by a Photovoltaic Array	239
<i>Giovanni Petrone, Francisco Jose Sánchez Pacheco, Giovanni Spagnuolo</i>	
Sustainable Energy Microsystems for a Smart Grid	259
<i>Maria Carmen Falvo, Luigi Martirano, Danilo Sbordone</i>	
SVM Methods for Optimal Management of a Virtual Power Plant	271
<i>Emanuele Crisostomi, Mauro Tucci, Marco Raugi</i>	
Active Power Losses Constrained Optimization in Smart Grids by Genetic Algorithms	279
<i>Gian Luca Storti, Francesca Possemato, Maurizio Paschero, Silvio Alessandrini, Antonello Rizzi, Fabio Massimo Frattale Mascioli</i>	
Solar Irradiation Forecasting for PV Systems by Fully Tuned Minimal RBF Neural Networks	289
<i>Lucio Ciabattini, Gianluca Ippoliti, Sauro Longhi, Matteo Pirro, Matteo Cavalletti</i>	
Ontology-Based Device Configuration and Management for Smart Homes	301
<i>Michele Nucci, Marco Grassi, Francesco Piazza</i>	
A Comparison between Different Optimization Techniques for Energy Scheduling in Smart Home Environment	311
<i>Francesco De Angelis, Matteo Boaro, Danilo Fuselli, Stefano Squartini, Francesco Piazza</i>	

**Part V: Special Session on “Computational Intelligence
in Emotional or Affective Systems”**

Towards Emotion Recognition in Human Computer Interaction	323
<i>Günther Palm, Michael Glodek</i>	
Towards Causal Modeling of Human Behavior	337
<i>Matteo Campo, Anna Polychroniou, Hugues Salamin, Maurizio Filippone, Alessandro Vinciarelli</i>	
How Social Signal Processing (SSP) Can Help Assessment of Bonding Phenomena In Developmental Psychology?	345
<i>Emilie Delaherche, Sofiane Boucenna, Mohamed Chetouani, David Cohen</i>	
Emotion and Complex Tasks: Writing Abilities in Young Graders	357
<i>Michaël Fartoukh, Lucile Chanquoy, Annie Piolat</i>	
A Preliminary Study of Online Drawings and Dementia Diagnose	367
<i>Marcos Faundez-Zanuy, Enric Sesa-Nogueras, Josep Roure-Alcobe, Josep Garre-Olmo, Jiri Mekyska, Karmele Lopez-de-Ipiña, Anna Esposito</i>	
Hand-Based Gender Recognition Using Biometric Dispersion Matcher	375
<i>Xavier Font-Aragones, Marcos Faundez-Zanuy</i>	
Revisiting AVEC 2011 – An Information Fusion Architecture	385
<i>Martin Schels, Michael Glodek, Friedhelm Schwenker, Günther Palm</i>	
Discriminating Human vs. Stylized Emotional Faces: Recognition Accuracy in Young Children	395
<i>Anna Esposito, Maria Teresa Riviello, Vincenzo Capuano</i>	
Emotional Status Determination in HCI Interface for the Paralyzed	405
<i>Rytis Maskeliunas, Vidas Raudonis, Paulius Lengvenis</i>	
Emoticons Signal Expertise in Technical Web Forums	415
<i>Liliana Mamani Sanchez, Carl Vogel</i>	
Machine Learning and Soft Computing Methodologies for Music Emotion Recognition	427
<i>Angelo Ciaramella, Giuseppe Vettigli</i>	
Homo-Machina Visual Metaphors, Representations of Consciousness and Scientific Thinking	437
<i>Mauro Maldonato, Ilaria Anzoise</i>	
Author Index	453

Part I

Algorithms

Probability Learning and Soft Quantization in Bayesian Factor Graphs

Francesco A.N. Palmieri and Alberto Cavallo

Dipartimento di Ingegneria Industriale e dell'Informazione
Seconda Università di Napoli (SUN)
via Roma 29, 81031 Aversa (CE), Italy
{francesco.palmieri,alberto.cavallo}@unina2.it

Abstract. We focus on learning the probability matrix for discrete random variables in factor graphs. We review the problem and its variational approximation and, via entropic priors, we show that soft quantization can be included in a probabilistically-consistent fashion in a factor graph that learns the mutual relationship among the variables involved. The framework is explained with reference the "Tipper" example and the results of a Matlab simulation are included.

Keywords: Machine Learning, Factor Graphs, Bayesian Methods.

1 Introduction

Probability propagation on graphs is a very promising emerging paradigm for building intelligent signal processing systems [12]. Algorithms and applications are under development in many areas of research that range from communication and coding to signal processing and control. However, full use and development of artificial intelligence systems that operate with probability propagation techniques require refinements on a number of critical issues. Some of these are: 1. Propagation in graphs with cycles [1]; 2. Parameter learning [8]; 3. Graph-structure learning [13]; 4. Propagation and learning in hybrid graphs with both continuous and discrete variables; etc. In this paper we focus on learning the probability matrix in discrete-variable factor graphs [7][6] pointing to a connection to variational learning [5][3][2][18][19]. We apply the idea to a generic block where the whole probability matrix is learned from examples. Recent development on inference based on entropic priors [15][14] allows the introduction of soft quantization within the Bayesian graph framework much like in fuzzy logic [17]. Entropic priors allow to translate some of the successful heuristics typical of the fuzzy framework, into a probabilistically-consistent Bayesian learning paradigm on factor graphs. Soft logic formulated within standard probability theory [10] coupled with belief propagating on factor graphs represents a very promising framework to bring to a higher cognitive level many of the current signal processing problems. In our formulation we use factor graphs in Forney's normal

form [11], because they are easier to handle in comparison to more traditional Bayesian graphs [16].

In this paper we first review the problem of learning the probability matrix pointing to a connection with variational message passing. Then we briefly introduce soft quantization with entropic priors and finally we apply the ideas to the well-known Tipper example. The results of a simulation show how this framework implements a very natural dynamic merge of inference and learning.

2 Learning the Probability Matrix

Probabilistic inference in factor graphs via message propagation is a relatively mature technique, at least in graphs with no cycles, when the conditional probability functions that make up the model are known [12]. A much harder problem is learning the model parameters on line, i.e. performing inference and learning at the same time. To focus on the specifics of this issue we start with the simplest (non trivial) factor graph of Figure 1 that models N independent realizations of two random variables $X \in \mathcal{X} = \{\xi_1, \dots, \xi_d\}$ and $Y \in \mathcal{Y} = \{\eta_1, \dots, \eta_m\}$. The variables are discrete and take values in the two alphabets \mathcal{X} and \mathcal{Y} and are related via the unknown conditional probability matrix

$$P(Y|X\Theta) = \begin{pmatrix} p(\eta_1|\xi_1) & \dots & p(\eta_m|\xi_1) \\ p(\eta_1|\xi_2) & \dots & p(\eta_m|\xi_2) \\ \vdots & \dots & \vdots \\ p(\eta_1|\xi_d) & \dots & p(\eta_m|\xi_d) \end{pmatrix} = \Theta = \begin{pmatrix} \Theta_{11} & \dots & \Theta_{1m} \\ \Theta_{21} & \dots & \Theta_{2m} \\ \vdots & \dots & \vdots \\ \Theta_{d1} & \dots & \Theta_{dm} \end{pmatrix}, \quad (1)$$

with $0 \leq \Theta_{ij} \leq 1$, $i = 1, \dots, d$, $j = 1, \dots, m$; $\sum_{j=1}^m \Theta_{ij} = 1$, $i = 1, \dots, d$. The unknown parameters make up the matrix $\Theta \in \mathcal{T}$, where \mathcal{T} denotes the set of all $d \times m$ stochastic matrices. Since the structure of Figure 1 may be part of a more complex network, we assume that information on $X[n]$ and $Y[n]$ is available in




Fig. 1. The factor graph for N independent realizations of $(X[n], Y[n])$

soft form via forward and backward distributions $f_{X[n]}(x)$, $b_{X[n]}(x)$, $f_{Y[n]}(y)$ and $b_{Y[n]}(y)$, with $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. Also information about matrix Θ is carried by forward and backward messages $f_{\Theta[n]}(\theta)$ and $b_{\Theta[n]}(\theta)$ which are matrix functions. These messages are related to each other via marginalization as

$$\begin{aligned} f_{Y[n]}(y) &\propto \int_{\theta \in \mathcal{T}} \sum_{x \in \mathcal{X}} P(y|x\theta) f_{X[n]}(x) f_{\Theta[n]}(\theta) d\theta; \\ b_{X[n]}(x) &\propto \int_{\theta \in \mathcal{T}} \sum_{y \in \mathcal{Y}} P(y|x\theta) b_{Y[n]}(y) f_{\Theta[n]}(\theta) d\theta; \\ b_{\Theta[n]}(\theta) &\propto \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} P(y|x\theta) b_{Y[n]}(y) f_{X[n]}(x). \end{aligned}$$

As usual in factor graphs, the notation \propto means that the expressions are distributions except for proper normalization. The complete model is hybrid because $X[n]$ and $Y[n]$ are discrete and Θ is continuous and multi-dimensional. In a more compact matrix representation, forward and backward messages for $X[n]$ and $Y[n]$ are the column vectors

$$\begin{aligned} \mathbf{f}_{X[n]} &= (f_{X[n]}(\xi_1), \dots, f_{X[n]}(\xi_d))^T; \quad \mathbf{b}_{X[n]} = (b_{X[n]}(\xi_1), \dots, b_{X[n]}(\xi_d))^T; \\ \mathbf{f}_{Y[n]} &= (f_{Y[n]}(\eta_1), \dots, f_{Y[n]}(\eta_m))^T; \quad \mathbf{b}_{Y[n]} = (b_{Y[n]}(\eta_1), \dots, b_{Y[n]}(\eta_m))^T. \end{aligned}$$

Therefore we can write

$$\begin{aligned} \mathbf{f}_{Y[n]} &\propto \int_{\theta \in \mathcal{T}} \theta^T \mathbf{f}_{X[n]} f_{\Theta[n]}(\theta) d\theta = F_{\theta[n]}^T \mathbf{f}_{X[n]}; \\ \mathbf{b}_{X[n]} &\propto \int_{\theta \in \mathcal{T}} \theta \mathbf{b}_{Y[n]} f_{\Theta[n]}(\theta) d\theta = F_{\theta[n]} \mathbf{b}_{Y[n]}, \end{aligned}$$

where $F_{\theta[n]} = \int_{\theta \in \mathcal{T}} \theta f_{\Theta[n]}(\theta) d\theta$ is the mean forward matrix for $\Theta[n]$. The backward message for $\Theta[n]$ is the matrix function

$$b_{\Theta[n]}(\theta) \propto \mathbf{f}_{X[n]}^T \theta \mathbf{b}_{Y[n]} = \mathbf{f}_{X[n]}^T \begin{pmatrix} \theta_{11} & \dots & \theta_{1m} \\ \theta_{21} & \dots & \theta_{2m} \\ \dots & \dots & \dots \\ \theta_{d1} & \dots & \theta_{dm} \end{pmatrix} \mathbf{b}_{Y[n]}. \quad (2)$$

Messages for $\Theta[n]$ and $\Theta'[n]$ in the other branches are formally the result of the product rule $f_{\Theta[n]}(\theta) \propto f_{\Theta'[n]}(\theta) b_{\Theta'[n-1]}(\theta)$; $b_{\Theta'[n]}(\theta) \propto b_{\Theta[n]}(\theta) b_{\Theta'[n-1]}(\theta)$; $f_{\Theta'[n]}(\theta) \propto b_{\Theta[n+1]}(\theta) f_{\Theta'[n+1]}(\theta)$. Each message is a product of the type

$$\mu_{\Theta}(\theta) \propto \prod_l \mathbf{f}_{X[l]}^T \begin{pmatrix} \theta_{11} & \dots & \theta_{1m} \\ \theta_{21} & \dots & \theta_{2m} \\ \dots & \dots & \dots \\ \theta_{d1} & \dots & \theta_{dm} \end{pmatrix} \mathbf{b}_{Y[l]} = \prod_l \sum_{i=1}^d \sum_{j=1}^m b_{Y[l]}(\eta_j) f_{X[l]}(\xi_i) \theta_{ij} \quad (3)$$

If variables $X[n]$ and $Y[n]$ of block n are *instantiated*, i.e. forward and backward messages are delta functions, $f_{X[n]}(x) = \delta(x - \xi_i)$, $b_{Y[n]}(y) = \delta(y - \eta_j)$, backward information from block n is simply $b_{\Theta[n]}(\theta) \propto \theta_{ij}$. If also all variables from all n are instantiated, information exchanged among the blocks (except possibly for the prior on Θ) are exactly products of Dirichlet distributions

$$\mu_{\Theta}(\theta) \propto \prod_{i=1}^d \prod_{j=1}^m \theta_{ij}^{n_{ij}} \propto \prod_{i=1}^d \text{Dir}(\theta_{i1}, \dots, \theta_{im}; n_{i1} + 1, \dots, n_{im} + 1), \quad (4)$$

where n_{ij} are the integer numbers that represent the cumulative counts of the occurrences of pair (i, j) (*hard scores*). Unfortunately, in the general case we are

interested in with forward and backward messages carrying soft information, expression (3) becomes intractable. Hence we resort to a variational approximation [5] [3] [18] for $b_{\Theta[n]}(\theta)$ that gives

$$b_{\Theta[n]}^V(\theta) \propto e^{\sum_{i=1}^d \sum_{j=1}^m b_{Y[n]}(\eta_j) f_{X[n]}(\xi_i) \log \theta_{ij}} = \prod_{i=1}^d \prod_{j=1}^m \theta_{ij}^{b_{Y[n]}(\eta_j) f_{X[n]}(\xi_i)} \quad (5)$$

$$\propto \prod_{i=1}^d \text{Dir}(\theta_{i1}, \dots, \theta_{im}; f_{X[n]}(\xi_i) b_{Y[n]}(\eta_1) + 1, \dots, f_{X[n]}(\xi_i) b_{Y[n]}(\eta_m) + 1),$$

which is again the product of d Dirichlet distributions. This is particularly interesting because the Dirichlet distribution, sometimes used as an assumption [7] [19], is exactly the variational approximation. Assuming that also the prior distribution π_{Θ} is a product of Dirichlet functions

$$\pi_{\Theta} \propto \prod_{i=1}^d \text{Dir}(\theta_{i1}, \dots, \theta_{im}; \alpha_{i1} + 1, \dots, \alpha_{im} + 1).$$

A generic message in the upper branches has the form

$$\mu_{\Theta} \propto \prod_{i=1}^d \text{Dir}(\theta_{i1}, \dots, \theta_{im}; \alpha_{i1} + \sum_l f_{X[l]}(\xi_i) b_{Y[l]}(\eta_l) + 1, \dots, \alpha_{im} + \sum_l f_{X[l]}(\xi_i) b_{Y[l]}(\eta_l) + 1). \quad (6)$$

A priori knowledge about the rule that maps X into Y can also be easily included in the coefficients of π_{Θ} . The exponential form for the variational approximation suggests that matrix variables $\Theta[n]$ and $\Theta'[n]$ could be replaced with *soft score* matrix variables $O[n]$ and $O'[n]$. Backward message from block n becomes matrix $b_{O[n]} = \mathbf{f}_{X[n]} \mathbf{b}_{Y[n]}^T$. Also all messages in the upper branches become $d \times m$ matrices with combination rules $f_{O[n]} = f_{O'[n]} + b_{O'[n-1]}$; $b_{O'[n]} = b_{O[n]} + b_{O'[n-1]}$; $f_{O'[n]} = b_{O[n+1]} + f_{O'[n+1]}$. Forward and backward messages for $Y[n]$ and $X[n]$ are respectively $\mathbf{f}_{Y[n]} \propto F_{O[n]}^T \mathbf{f}_{X[n]}$; $\mathbf{b}_{X[n]} \propto F_{O[n]} \mathbf{b}_{Y[n]}$, where $F_{O[n]}$ is the row-normalized version of $f_{O[n]}$. Note that these propagation rules represent the learning steps for Θ as inference and learning happen at the same time. Recall that the various stages in the graph represent time-unfolded versions of the same block. More details and proofs will be reported in a longer paper.

3 Soft Quantization

Manipulation of discrete quantities in machine learning, also when the problem involves continuous variables, may be particularly handy, because a priori qualitative information can be more easily injected into the system. Fuzzy methods [17] have shown great success in merging soft knowledge with hard functions especially in control [9]. In [15] we have shown how the use entropic priors in the Bayesian framework allowing the introduction of soft membership information in a way that is consistent within standard probability theory. This is a crucial step to allow soft quantization and coherent use of probability propagation for inference and learning in systems that contain both continuous and discrete variables.

Figure 2 shows a quantization scheme for a continuous variable S_a . All the likelihoods are triangular, complementary and centered on the M nodes ξ_1, \dots, ξ_M . Denoting the triangular function on a, b, c with $\Lambda(s_a; a, b, c)$, the M pdfs are

$$\left\{ \frac{2}{\xi_2 - \xi_1} \Lambda(s_a; \xi_1, \xi_1, \xi_2), \frac{2}{\xi_3 - \xi_1} \Lambda(s_a; \xi_1, \xi_2, \xi_3), \dots, \frac{2}{\xi_M - \xi_{M-2}} \Lambda(s_a; \xi_{M-2}, \xi_{M-1}, \xi_M), \frac{2}{\xi_M - \xi_{M-1}} \Lambda(s_a; \xi_{M-1}, \xi_M, \xi_M) \right\}, \quad (7)$$

and are shown in Figure 2(a). The differential entropy [4] of $\Lambda(s_a; a, b, c)$ is easily computed to be $h(S_a) = \frac{1}{2} + \log \frac{c-a}{2}$. With entropic priors $\pi_i \propto e^{h(S_a|i)}$ [15], the prior-likelihood products, become equivalent to a set of functions with same height as in Figure 2(b). We recall that entropic priors are the distribution that maximize the joint entropy $H(S_a, S)$ for fixed likelihoods $(p_{S_a}(s_a|1), \dots, p_{S_a}(s_a|M))$ [15]. The node distribution can be chosen according to the data points density, but the complementarity of the likelihoods guarantees that no information is lost after soft quantization. Figure 2(b) shows also how this kind of soft quantization can be drawn as a generative factor graph model that can be inserted into a larger factor graph. The backward message for S_a is a data point $b_{S_a}(s_a) = \delta(s_a - s_0)$. The backward message for S^1 in vector notation is $\mathbf{b}_{S^1} = (p_{S_a}(s_0|1), \dots, p_{S_a}(s_0|M))^T$ that after combination with entropic priors becomes $\mathbf{f}_{S^2} = (p_{S_a}(s_0|1)\pi_1, \dots, p_{S_a}(s_0|M)\pi_M)^T$. The soft quantization model satisfies a property of perfect reconstruction because if $\mathbf{b}_{S^2} = \mathbf{b}_{S^1}$, we have $\mathbf{f}_{S^1} = \mathbf{f}_{S^2}$ and $f_{S_a}(s_a) = \delta(s_a - \mathbf{f}_{S^1}^T(\xi_1, \dots, \xi_M)^T) = \delta(s_a - s_0)$ (lossless dequantization). More details about soft quantization with entropic priors will be reported in a longer paper elsewhere.

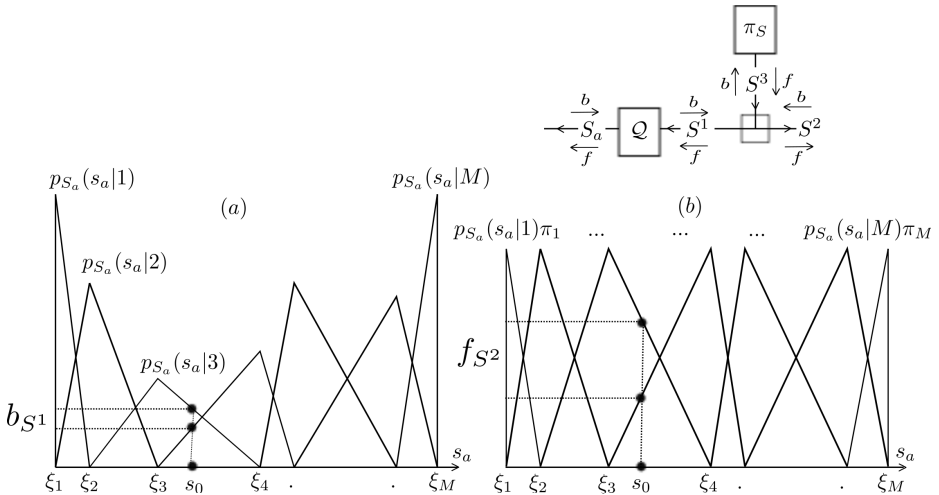


Fig. 2. Soft quantization on nodes $\{\xi_1, \dots, \xi_M\}$. (a) The triangular likelihoods; (b) The entropic priors-likelihoods products.

4 The Tipper Example

In this paper we report some experiments with the variational learning rules of Section 2 and with the soft quantization scheme of Section 3 on the well-known "Tipper" example. In this problems there are three continuous variables: S_a (Service), F_a (Food) and T_a (Tip). The Tipper example, often used as a teaching example in control classes (there is a Matlab demo available in the Fuzzy Control Toolbox), is a typical case of mapping between two input variables (Service and Food) and a final one (Tip). In the fuzzy framework is also very easy to include soft rules and various design constraints. Our objective is here to traslate this typical approach into a probabilistically-consistent Bayesian framework. The underlying factor graphs shown in Figure 4, in which messages travel back and forth, allows simultaneous inference and learning with inputs and outputs that become essentially indistinguishable.

Even though a priori soft-logic rules can be easily included as constraints in the prior block π_Θ , we have assumed here no prior knowledge about the variables S_a , F_a and T_a . We have simply presented 50 realizations of the triplet as f_{S_a} , f_{F_a} and f_{T_a} and let the system learn (simulations with combinations of soft rules and examples will be reported elsewhere). The triplets were obtained from a blind run of the Matlab demo. Forward and backward messages carry information in various parts of the system and inferences can also be made backward on Service and/or Food from Tip.

The graph structure assumes that variables Service and Food are mutually independent and that the $N = 50$ realizations are also statistically independent. The three analog variables S_a , F_a and T_a are soft quantized from ranges $[0 - 10][0 - 10][5 - 25]$ with $M = 6$ uniformly spaced nodes each into the three discrete variables S , F and T . Entropic priors are imposed in π_S , π_F and π_T . The 36×6 matrix of conditional probabilities $P(T|SF\Theta)$ is learned via message propagations with the variational algorithm described in Section 2. The simulations let the messages propagate 300 steps which is enough to cover the graph diameter. The graph is clearly a tree and convergence is guaranteed. Figure 4 shows the comparison of forward and backward information at each stage n . The thre plots show the comparison of the actual value of each variable, as carried by the backward input message, with the value provided by the rest of the system, as carried by the forward output message that uses all the other inputs after learning and propagation. Note that learning and inference is all done at the same time since information about the parameter θ are also carried by travelling messages. The simulation is self-contained and implements our best use of the data because the inference, say on $T_a[n]$, is based on all the examples except the one on $T_a[n]$. This is because $f_{\Theta[n]}$ does not contain information coming from $b_{\Theta[n]}$. Hence each stage n uses a slightly different estimate for $P(T|SF\Theta)$ because the n th examples is automatically excluded. Therefore in inferring $T_a[n]$, values of $S_a[n]$ and $F_a[n]$ are used for inference, but not for learning. The same considerations apply to inferences on $S_a[n]$ and $F_a[n]$.

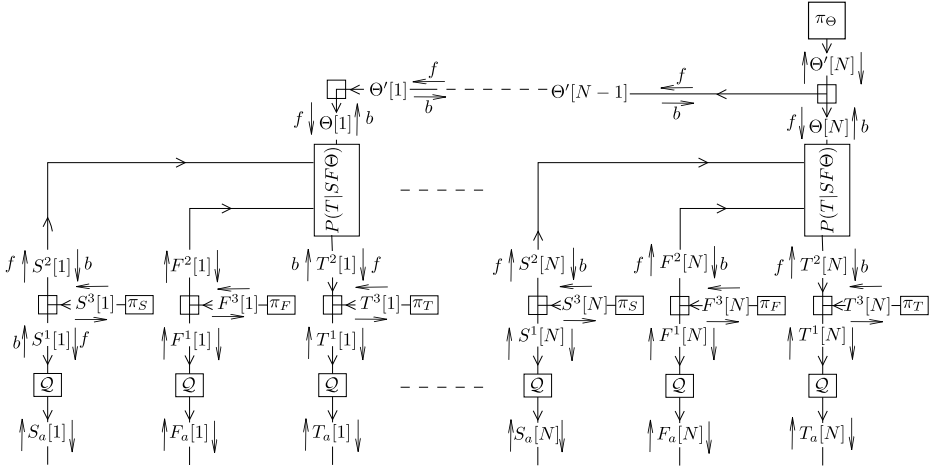


Fig. 3. The factor graph for the Tipper example




Fig. 4. Comparison of forward (inference) (*) and backward (true) (o) values for the three variables in the Tipper example

Conclusions

In this work we have reported partial results for a successful Bayesian paradigm that implements via message propagation on a factor graph simultaneous inference and learning. By means of an example, we have also proposed a quantization scheme, that via the introduction of entropic priors, allows us to build a probabilistically-consistent graph that can be adapted with belief propagation. More work will be devoted to further understanding of the adaptation rules and on the inclusion of soft-logic constraints.

References

1. Graphical models emerge. new connections between machine learning and signal processing. *Signal Processing Magazine*, 27(6) (2010)
2. Beal, M.J.: Variational algorithms for approximate bayesian inference. Ph.D. thesis, University of London (2003)
3. Beal, M.J., Ghahramani, Z.: Variational bayesian learning of directed graphical models with hidden variables. *Bayesian Analysis* 1, 1–44 (2004)
4. Cover, T.M., Thomas, J.A.: *Elements of Information Theory*. Wiley (2006)
5. Dauwels, J.: On variational message passing in factor graphs. In: ISIT2007, Nice, France, June 24–June 29 (2007)
6. Ghahramani, Z.: *Unsupervised Learning*. Springer (2004)
7. Heckerman, D.: A tutorial on learning with bayesian networks. Tech. Rep. MSR-TR-95-06, Microsoft Research (1996); March 1995 (Revised November 1996)
8. Dauwels, J., Eckford, A., Loeliger, S.K., Expectation, H.A.: Expectation maximization as message passing—part i: Principles and gaussian messages. arXiv:0910.1–14 (2009); Submitted to *IEEE Tr. on Information Theory*
9. Jantzen, J.: *Foundations of Fuzzy Control*. Wiley (2007)
10. Jaynes, E.T.: *Probability Theory: The Logic of Science*. Cambridge University Press (2003)
11. Loeliger, H.A.: An introduction to factor graphs. *IEEE Signal Processing Magazine* 21(1), 28–41 (2004)
12. Loeliger, H.A., Dauwels, J., Hu, J., Korl, S., Ping, L., Kschischang, F.: The factor graph approach to model-based signal processing. *Proceedings of the IEEE* 95(6), 1295–1322 (2007)
13. Choi, M.J., Tan, V.Y.F., Anandkumar, A., Willsky, A.S.: Learning latent tree graphical models. *Journal of Machine Learning Research* 12, 1771–1812 (2011)
14. Palmieri, F.A.N., Ciunzono, D.: Entropic priors for short-term stochastic process classification. In: 14th Int. Conf. on Information Fusion, Chicago, IL (2011)
15. Palmieri, F.A.N., Ciunzono, D.: Objective priors from maximum entropy in data classification. In: *Information Fusion* (2012), doi:10.1016/j.inffus.2012.01.012
16. Pearl, J.: *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers Inc., San Francisco (1988)
17. Novak, V., Perfilieva, I., Mockor, J.: *Mathematical Principles of Fuzzy Logic*. Kluwer Academic Press (1999)
18. Winn, J., Bishop, C.M.: Variational message passing. *Journal of Machine Learning Research* 6, 661–694 (2005)
19. Winn, J.M.: Variational message passing and its applications. Ph.D. thesis, University of Cambridge (2004)

Rival-Penalized Competitive Clustering: A Study and Comparison

Alberto Borghese and Wiliam Capraro

Applied Intelligent Systems Laboratory, Dept. of Computer Science, University of Milano
borghese@di.unimi.it, wiliam.capraro@studenti.unimi.it

Abstract. A major recurring problem in exploratory phases of data mining is the task of finding the number of clusters in a dataset. In this paper we illustrate a variant of the competitive clustering method which introduces a rival penalization mechanism, and show how it can be used to solve such problem. Additionally, we present some tests aimed at comparing the performance of our rival-penalized technique with other classical procedures.

1 Introduction

The term central clustering refers to a family of clustering algorithms that are based on moving a set of points, referred to as prototypes, inside the data space until their position minimizes a certain cost function, representing a measure of the goodness by which the prototypes represent the data points.

In the literature, approaches based on soft-clustering, like fuzzy c-means [14], Neural-Gas [13,4] or Self-Organizing Maps (SOM) [6,16], and competitive learning [5,15,12], have been proposed to partition a given dataset into a predefined number of clusters, each one represented by a prototype usually corresponding to the cluster centroid. All these algorithms suffer from several issues, most notably, the optimal value of the cost function is rarely reached. Only stochastic optimization [10], that is extremely costly, or a careful initialization of the prototypes allow escaping local minima. Although a few attempts have been proposed to derive a robust initialization (e.g. [11]), there seem to be no universal and reliable way to proceed, and some prototypes typically get stuck during the clustering process. These prototypes are referred to as “dead units” [13] and affect the proper operation of the algorithm and the quality of the result. As a consequence, the general approach is to repeat the clusterization process several times with different, random initializations of the prototypes, so as to allow the algorithm to escape from local minima from time to time.

A slightly different task is to find the number of clusters in a dataset (e.g. [3]). This is indeed a frequent problem in exploratory phases of data mining, and a straightforward approach is to adopt a parameterized version of a clustering algorithm using the desired number of data clusters K as parameter and to try a different clusterization for each possible value of the parameter. Subsequently, the best result would be chosen based on some validity measure or index.

However, it is impractical to run several random initializations of one algorithm for each possible values of its parameter, especially when the latter spans a wide interval of

values. Alternatively, in some cases it is possible to exploit the intrinsic characteristics of some algorithms to produce dead units (e.g. [6]), to facilitate the search for a good solution.

Recently, in the framework of central clustering, a different approach has been proposed for moving the prototypes. The idea is that, while the winning prototype is attracted by the closest data point, other prototypes are moved in the opposite direction. This mechanism, known as rival-penalization [5][15], is somehow similar to the BCM model proposed by Bienenstock, Cooper and Munro [9] in a typical Hebbian learning fashion.

Rival-penalization clustering has been overlooked in the past. In this paper we present the results of some tests we performed, aimed at comparing the rival-penalization approach with classical clustering techniques, namely standard competitive learning and SOM. The ability of rival-penalization of discovering the proper number of clusters in a given dataset is analysed and discussed. Specifically, we show that, by introducing the rival penalization mechanism into a competitive learning setting, results comparable with soft-clustering can be achieved. Moreover, the number of clusters can be discovered in a robust and reliable way.

2 Algorithms

We'll consider a collection of N d -dimensional observations, $\{\xi_j\}$. Goal of clustering algorithms is to assign each observation to one of K clusters Ψ_i , according to a similarity measure with the other elements in the same cluster. Each cluster is represented by its centroid, ψ_i , which is also a point in \mathbb{R}^d .

The following subsections give a quick coverage of the algorithms we employed.

2.1 Competitive Learning and Rival Penalization

Competitive learning (CL) is an effective tool for data clustering, widely applied in a variety of signal processing problems such as data compression, classification, adaptive noise cancelation, image retrieval and image processing [2].

For the purposes of this contribution, a feed-forward neural network with a single layer consisting of K output units is used to achieve a K -cluster data partitioning. Each unit represents a cluster centroid ψ_i .

The training of the network proceeds as follows. At each iteration, each data point ξ is presented in turn to the network and a winning unit, w , is elected. This is the prototype whose Euclidean distance from the point is minimum:

$$w = \underset{i}{\operatorname{arg\,min}} \|\xi - \psi_i\|. \quad (1)$$

Subsequently, the position of the winning unit is updated towards the data point using the following updating rule

$$\psi_{wj} = \psi_{wj} + \eta_{(t)}(\xi - \psi_w) \quad (2)$$

where j denotes a component of the prototype vector and $\eta_{(t)}$ is a learning rate parameter whose value decays as a function of the time t .

In a pure competitive learning setting, only the winning unit is updated. The procedure is repeated multiple times for each data point, until the prototypes converge to their final position—i.e. when the maximum difference in the position of any centroid in two successive iterations is smaller than a fixed tolerance ε , or when a maximum number of iterations is reached.

The prototypes are initialized using the "Forgy" approach [11]—i.e. K of the available data points are randomly chosen to serve as cluster prototypes. In this context, this is enough to guarantee that no dead unit will ever appear, as every prototype shall win the competition for at least one data point, that is, the prototype itself.

The rival-penalized competitive learning (RPCL) algorithm improves on the pure competitive learning approach by introducing a rival penalization mechanism, as proposed in [5] and [15]. With this approach, not only the position of the winning unit is updated towards the input vector, but additionally the position of its rival unit is updated in the opposite direction.

In order to find the winning unit and its rival, a relative winning frequency is introduced, which keeps track of how many times each unit happens to win a competition for some input vector. The relative winning frequency for unit i is defined as

$$\gamma_i = \frac{s_i}{\sum_{j=1}^K s_j} \quad (3)$$

where s_i is the number of times unit i was declared winner in the past. When $\sum_{j=1}^K s_j = 0$ —i.e. initially, then $\gamma_i = 1$ in order to give every prototype a fair chance to win.

The winning unit w for an input vector ξ is now given by

$$w = \arg \min_i \gamma_i \|\xi - \psi_i\|. \quad (4)$$

Notice how the parameter γ_i acts as a "conscience" for the unit—if the unit has won too often in the past, its chances to win the competition for the current data point are reduced accordingly. Moreover, for each input vector ξ , the rival penalized competitive learning algorithm computes not only the winning unit w , but also a second winning unit, referred to as the rival, defined by

$$r = \arg \min_i \gamma_i \|\xi - \psi_i\|, \quad i \neq w. \quad (5)$$

Equation 2 is used to update both the winner and its rival. The latter, however, moves away its centroid from the input point with a de-learning rate β , which is related to η by

$$\beta_{(t)} = -c\eta_{(t)}\gamma_r \quad (6)$$

where γ_r is the relative winning frequency of the rival and $c = 1/10$ is a predefined constant. Unlike the implementation of [5], here β depends on both the learning rate η and the winning frequency γ_r , so that the rival is dynamically penalized according to γ_r even for constant η (which is not the case anyway).

In contrast to the CL algorithm, here a "Forgy" initialization of the prototypes is not enough to guarantee dead unit avoidance. In fact, even if the prototypes are initialized using the input data points, depending on the de-learning rate β , a rival unit may incur

considerable modification in the value of its prototype, and thus it can fail to win the competition even for the input data point to which it had been initialized.

What is interesting with this approach is that, as reported in [5], if the learning rate η is chosen to be at least one order of magnitude larger than β , then the adequate number of output clusters will be automatically found. In other words, assuming that the actual number of clusters is unknown and that the number of units K is chosen greater than the cluster number, the prototype vectors will converge towards the centroids of the actual clusters with few of them overlapping in space. In our implementation, this condition holds in each iteration as $c = 1/10$. In each iteration, the RPCL algorithm pushes away the rival, thus allowing for faster convergence, and invalidates extra prototypes by eventually making their cluster empty. Hence, the RPCL algorithm is believed to be able to perform appropriate clustering without knowing the cluster number.

2.2 SOM

The limitation of considering only one data point at a time in competitive learning has been overcome by soft-clustering approaches [2] in which the position of all the prototypes is updated for each data point. Among these approaches, Self-Organizing Maps (SOMs) represent an excellent tool in exploratory phases of data mining. They project the input space onto prototypes in a low-dimensional regular grid that can be effectively used to visualize and explore properties of the data. The SOM consists of a regular, one- or two-dimensional grid of units, with each unit i represented by its prototype vector ψ_i . Additionally, each unit i is assigned a place in the output grid, represented by its coordinates $r_i = (x_i, y_i)$, and the units are logically linked to adjacent ones by a neighborhood relation. During training, data points lying near each other in the input space are mapped to nearby units in the output hyperplane. Thus, the SOM can be regarded as a topology-preserving tool for mapping the input space onto the output grid.

The SOM is trained iteratively. At each training step, a data point ξ is randomly chosen from the input data set, and the distance between ξ and all the prototype vectors is computed. Subsequently, all prototype vectors are updated, each proportionally to the distance of the corresponding unit from the winning unit in the output grid:

$$\psi_{ij} = \psi_{ij} + \eta_{(t)} \Lambda(i, w) (\xi - \psi_i). \quad (7)$$

In the hereabove equation $\Lambda(i, w)$ denotes the value of the neighborhood function between unit i and the winning unit w , as given by

$$\Lambda(i, w) = \exp\left(-\frac{\|r_i - r_w\|^2}{2\sigma^2}\right) \quad (8)$$

where the parameter σ defines the radius of the neighborhood. $\Lambda(i, w)$ therefore defines a region of influence for the prototype w . Notice that the value of Λ is exactly 1 when $i = w$, and decreases as the distance of the prototype from the other data points increases. Also, it is useful to adjust the radius as well as the learning rate at each iteration, so that the influence region of a prototype decays with time as a function of σ and η .

In this work, we are mainly concerned with the SOM’s ability to perform appropriate clustering of a given data set. Thus, only SOMs with one-dimensional output arrays are actually used. As stated in [16], this configuration is expected to produce better results as compared to the 2-dimensional grid configuration. This is due to the fact that the “tension” exerted in each unit by the neighboring units is much higher in the second configuration, and such a tension limits the plasticity of the SOM to adapt to the particular distributions of the dataset.

3 Experimental Setting and Test Results

In order to test the algorithms, we have generated a specific dataset containing 250 3- d data points distributed over 5 non-overlapping clusters. It has been generated by perturbing the centroid of each cluster with a Gaussian distribution with mean value 0 and variance 1.

Since the resulting clusterization depends strongly on the initialization of prototypes, it is essential that each algorithm be tested several times with different initializations. For our tests, 60 “Forgy” initializations (which we’ll refer to as trials) have been generated and evaluated for each algorithm. This should be enough to overcome random fluctuations.

As to the tests we conducted, they can be divided into two types. In a first type thereof—we call it type-A experiment—we focused on a specific algorithm and tried to partition our dataset varying the cluster number from $K = 2$ to $K = 10$. (And for each K , 60 trials have been performed as described before.) On the other hand, in type-B experiments we tested 60 trials of an algorithm with a fixed value of K but varying the parameters of the algorithm instead.

In our tests, the validity of the resulting clusterization is evaluated by means of quality indexes, and the number of iterations required by the algorithm to converge was also measured. The quality indexes used are the Davies-Bouldin (DB) index and the mean quadratic error (MQE). The mean quadratic error is simply the ratio of the sum of all the squared distances of each data point from its cluster prototype to the total number of data points:

$$MQE = \frac{\sum_{i=1}^K \sum_{\xi \in \Psi_i} \|\xi - \psi_i\|^2}{\sum_{i=1}^K |\Psi_i|}. \quad (9)$$

The Davies-Bouldin index is defined as

$$DB = \frac{1}{K} \sum_{i=1}^K \max_{j \neq i} \left\{ \frac{S_i + S_j}{\|\psi_i - \psi_j\|} \right\} \quad (10)$$

where S_i is the within i -th cluster scatter, as given by

$$S_i = \sqrt{\sum_{\xi \in \Psi_i} \frac{\|\xi - \psi_i\|^2}{|\Psi_i|}}. \quad (11)$$

For a detailed review of these and other cluster validity measures see [8]. Notice that it is geometrically plausible to seek clusters that have minimum within-cluster scatter

and maximum between-class separation, so the number of clusters \bar{K} that minimizes the Davies-Bouldin index can be reasonably taken as the optimal value of K . As reported in [8], for well-separated clusters, the Davies-Bouldin index is expected to decrease monotonically as K increases until the correct number of clusters is achieved.

In all the tests conducted, we fixed a tolerance of $\varepsilon = 0.001$ and the maximum number of iterations was set to 500. It should be enough for the algorithms to produce good clusterizations given the time-decay rule for η adopted, which is

$$\eta(t) = \exp\left(-\frac{t}{50}\right) \cdot \eta_0 \quad (12)$$

where $\eta_0 = 0.1$ is the initial value and $t = 0, 1, \dots$ is the iteration number.

In the remaining of this section we illustrate the results of our tests.

3.1 Competitive Learning

To begin with, we measured the effectiveness of the standard CL algorithm in partitioning our dataset by running a type-A test. The results can be used throughout the rest of this work as a reference for the other algorithms. For each value of K , Table 1 reports the Davies-Bouldin index and the quadratic error for the best outcome out of the 60 trials of the algorithm, along with the number of iterations performed.

As Table 1 shows, the algorithm succeeds in discovering the correct number of clusters: the DB index takes on its optimal value for $K = 5$.

As expected, the mean quadratic error is a decreasing function of the number of clusters (indeed one expects the within-cluster variance to decrease in this case), and hence it does not convey any useful information on the goodness of the result.

As a downside, the CL algorithm takes a considerable number of iterations to converge, as in each iteration only the winning unit is moved towards the current data point by a small, η -dependent, fraction of the distance. Moreover, in order to discover the optimal value of the parameter K , every possible value has to be investigated and the result evaluated. The average number of iterations is a decreasing function of K , which is rather obvious since, for small K , we expect the amount of modification in the position of each centroid as a function of the data points to be higher in each iteration as compared to when K is large.




Fig. 1. Scatterplot of the dataset

Table 1. Type-A test results for CL

K	DB	MQE	it
2	0.712	49.642	344
3	0.396	27.360	344
4	0.429	10.583	328
5	0.298	2.816	298
6	0.690	2.648	298
7	0.751	2.603	298
8	0.768	2.540	298
9	0.746	2.435	298
10	0.809	10.197	328

Table 2. Type-A test results for SOM

(a) $\sigma = 0.5$					(b) $\sigma = 1$					(c) $\sigma = 1.5$				
K	nD	DB	MQE	it	K	nD	DB	MQE	it	K	nD	DB	MQE	it
2	0	0.7123	49.6421	346	2	0	0.7123	49.6421	346	2	0	0.7123	49.6421	346
3	0	0.5779	32.8641	345	3	0	0.5779	32.8641	345	3	0	0.5779	32.8641	345
4	0	0.4291	10.5827	330	4	0	0.4291	10.5827	330	4	0	0.4291	10.5827	330
5	0	0.2983	2.8161	300	5	0	0.2983	2.8161	300	5	0	0.2983	2.8161	300
6	1	0.2983	2.8161	300	6	1	0.2983	2.8161	300	6	1	0.2983	2.8161	300
7	2	0.2983	2.8161	300	7	2	0.2983	2.8161	300	7	2	0.2983	2.8161	300
8	3	0.2983	2.8161	300	8	3	0.2983	2.8161	300	8	3	0.2983	2.8161	300
9	4	0.2983	2.8161	300	9	3	0.7244	2.6374	294	9	2	1.0505	2.5142	294
10	5	0.2983	2.8161	300	10	3	1.0228	2.4724	294	10	3	1.0710	2.5077	294

We have also investigated the role of the learning rate η in the learning process. To this end, a type-B test has been performed in which a value of $K = 5$ has been fixed and η takes on some values in the range (0.2-0.02). As before, 60 trials of the algorithm have been tested for each value of η , and the best outcome is considered. We do not report the results for space issues. We report, however, that the initial learning rate seems to play no crucial role in the learning process, since for every value of η the best value obtained for the DB index is the same as that of Table [1](#).

3.2 SOM

As our second test, we have investigated the performance of a SOM in achieving proper clusterizations of the dataset. As before we ran a type-A test using a 1-dimensional output array of units, as we are not interested in the spatial organization of the resulting cluster centroids. The distance of each prototype from its neighbor prototypes on the output array has been set arbitrarily to 1, which is also the initial value for the radius of the neighborhood σ . The general idea is that, by exploiting the SOM's inherent ability to produce dead units, it is possible to avoid testing every possible value of the parameter K provided it is chosen larger than the actual number of clusters. Results are reported in Table [2b](#), where nD represents the number of dead units and again each line represents the best outcome for all the 60 initializations, according to the Davies-Bouldin index.

As the table implies, the minimum of the DB index is obtained for values of K in the range between $K = 5$ and $K = 8$. The values reported confirm the ability of the SOM to produce good clusterizations of the dataset, with values comparable with that of the competitive learning approach for all values of the parameter K . The results also advocate the thesis that the SOM is able to invalidate extra clusters and discover the correct number of clusters if the parameter K is chosen in a neighborhood of its optimal value.

We did not expect the quality of the resulting clusterization to change considerably as a function of the initial learning rate η , so we did not conduct any test in this respect. It is interesting, however, to observe the behavior of the algorithm when the initial radius is enlarged or restricted. Tables [2a](#) and [2c](#) also show the result of type-A tests

Table 3. Type-B test results for RPCL

(a) $\eta = 0.1$					(b) $\eta = 0.3$					(c) $\eta = 0.5$				
#	K	DB	MQE	it	#	K	DB	MQE	it	#	K	DB	MQE	it
1	5	0.2973	2.8166	298	1	5	0.3312	3.6632	366	1	6	0.9961	1743.7830	499
2	8	14.1375	436.7141	499	2	5	0.2955	2.8196	353	2	7	1.1205	529.9871	499
3	5	0.2983	2.8161	298	3	5	0.2978	2.8162	353	3	1	n.a.	n.a.	499
4	7	0.8392	3.4301	337	4	6	1.0019	594.3351	499	4	9	1.0335	8.1133	440
5	5	0.2977	2.8162	298	5	10	3.2478	12092.0839	499	5	8	2.0086	1834.3202	499
6	5	0.2972	2.8165	298	6	9	2.2782	27149.0575	499	6	6	0.6528	2.9607	411
7	5	0.2977	2.8162	298	7	7	2.2396	407.3565	499	7	7	1.7469	27172.1130	499
8	5	0.2972	2.8168	298	8	1	n.a.	n.a.	499	8	1	n.a.	n.a.	499
9	8	1.2113	127.1700	485	9	9	2.0219	14926.2644	499	9	1	n.a.	n.a.	499
10	5	0.2980	2.8161	298	10	8	1.0686	111.9457	499	10	1	n.a.	n.a.	499
11	9	1.9186	283.6699	486	11	1	n.a.	n.a.	494	11	3	1.3013	84.3499	454
12	5	0.2981	2.8161	298	12	1	n.a.	n.a.	496	12	7	1.0529	659.0174	499
13	9	1.5104	2.2541	285	13	8	3.3781	736.2310	499	13	8	1.5026	353.7500	499
14	10	1.4525	32.4955	438	14	6	0.6007	2.6212	353	14	5	0.2983	2.8161	379
15	8	1.8814	119.3174	478	15	9	1.1331	2.4473	363	15	7	1.3130	663.7360	499
16	9	3.2711	1289.8495	499	16	8	1.1660	254.0439	499	16	1	n.a.	n.a.	499
17	5	0.2981	2.8161	298	17	9	4.7230	68607.6355	499	17	4	0.9470	45.7815	432
18	10	1.3259	2.0547	283	18	1	n.a.	n.a.	499	18	9	8.7524	200.3968	499
19	5	0.2964	2.8175	298	19	1	n.a.	n.a.	486	19	8	1.3188	35309.9092	499
20	7	0.8711	2.4896	298	20	1	n.a.	n.a.	459	20	8	2.2503	712.0748	499
21	5	0.2961	2.8180	298	21	5	0.2948	2.8272	353	21	5	0.2983	2.8161	379
22	8	1.2890	187.3231	478	22	1	n.a.	n.a.	499	22	1	n.a.	n.a.	499
23	5	0.2983	2.8161	298	23	8	1.4005	22.1895	463	23	1	n.a.	n.a.	499
24	10	1.4116	2.0938	287	24	7	0.8376	2.4683	345	24	1	n.a.	n.a.	499
25	8	1.7381	87.3881	429	25	9	1.1667	2.7230	374	25	7	0.7805	3.0149	401
26	5	0.2959	2.8205	298	26	1	n.a.	n.a.	495	26	6	1.0100	270.6201	499
27	8	1.4844	786.6143	499	27	9	1.1134	6.0580	410	27	1	n.a.	n.a.	454
28	9	2.2929	931.2280	499	28	8	1.1568	355.1591	499	28	7	2.1072	12438.7022	499
29	5	0.2983	2.8161	298	29	6	1.4096	91.4646	443	29	1	n.a.	n.a.	499
30	5	0.2967	2.8174	298	30	7	1.1727	7076.8766	499	30	7	1.0933	3285294.0383	499
31	5	0.2970	2.8167	298	31	9	1.2331	22.9518	479	31	8	0.9186	3.4217	410
32	5	0.2975	2.8164	298	32	9	1.4621	599.8796	499	32	9	2.0005	4802.1818	499
33	5	0.2983	2.8161	298	33	9	1.2037	2.8699	384	33	7	0.9568	2.8659	402
34	5	0.2959	2.8183	298	34	8	0.9067	2.8137	374	34	8	2.0441	15709.4822	499
35	7	0.8889	2.5653	292	35	1	n.a.	n.a.	470	35	9	8.6214	631.4390	499
36	5	0.2983	2.8161	298	36	10	3.6183	406.4338	499	36	7	1.0028	502.4863	499
37	5	0.2973	2.8165	298	37	7	2.0509	729.5399	499	37	1	n.a.	n.a.	499
38	5	0.2979	2.8161	298	38	8	0.9922	2.6372	352	38	1	n.a.	n.a.	499
39	5	0.2973	2.8164	298	39	9	1.2467	32.5807	487	39	7	2.0152	147.8561	482
40	7	1.4684	206.7596	491	40	1	n.a.	n.a.	467	40	2	0.7452	106.3408	445
41	5	0.2975	2.8164	298	41	1	n.a.	n.a.	499	41	8	2.7086	503.6955	499
42	5	0.2959	2.8188	298	42	8	1.8052	5702.3187	499	42	6	2.6311	41958.7468	499
43	9	1.0727	2.5370	297	43	9	1.0551	2.6401	363	43	8	2.3926	889.5130	499
44	5	0.2975	2.8163	298	44	1	n.a.	n.a.	499	44	6	0.9915	904.8396	499
45	9	1.2583	69.7734	465	45	9	1.6941	22.7070	461	45	5	1.6410	564.5119	499
46	9	1.2287	150.8042	494	46	1	n.a.	n.a.	499	46	1	n.a.	n.a.	499
47	8	1.3926	246.4092	499	47	7	1.7933	317.9108	499	47	1	n.a.	n.a.	499
48	5	0.2959	2.8183	298	48	9	1.1090	3.3827	386	48	7	1.1248	208.1925	499
49	7	2.1666	222.3385	487	49	9	1.1030	3.8625	396	49	1	n.a.	n.a.	499
50	5	0.2983	2.8161	298	50	7	1.0535	1192.2263	499	50	1	n.a.	n.a.	499
51	5	0.2958	2.8192	298	51	10	1.9694	1978.7778	499	51	2	0.8228	102.4013	450
52	9	1.4211	51.5298	445	52	5	0.2982	2.8161	353	52	7	0.8945	2.7744	378
53	5	0.2976	2.8162	298	53	8	5.0827	783.8008	499	53	7	1.0473	1132.3572	499
54	5	0.2970	2.8168	298	54	7	1.9763	3376261.9184	499	54	1	n.a.	n.a.	499
55	9	1.3163	988.7258	499	55	9	1.2069	39.0237	497	55	6	1.0058	582.4626	499
56	5	0.2982	2.8161	298	56	5	0.2983	2.8161	353	56	1	n.a.	n.a.	499
57	9	1.2233	2.2424	292	57	5	0.2894	2.9133	354	57	10	6.9798	282791.2331	499
58	8	1.7596	428.0865	499	58	7	1.0936	701.7895	499	58	1	n.a.	n.a.	469
59	5	0.2978	2.8161	298	59	1	n.a.	n.a.	499	59	6	0.9993	464.6136	499
60	5	0.2982	2.8161	298	60	8	2.7499	133.5437	499	60	1	n.a.	n.a.	499

for $\sigma = 1.5$ and $\sigma = 0.5$: results show a general tendency of the radius to influence the ability of the SOM to kill extra units—this ability seems to increase as the radius of the neighborhood narrows.

3.3 Competitive Clustering with Rival Penalization

Lastly, we have analysed the performance of our RPCL implementation in discovering the correct number of clusters. As suggested in [5] we have chosen a number of clusters, $K = 10$, larger than the true number of clusters. Recall that the de-learning rate β is always at least one order of magnitude smaller than η . Once again we considered 60 “Forgy” initialization of the algorithm, and ran three type-B tests using different learning rates, namely $\eta = 0.1$, $\eta = 0.3$ and $\eta = 0.5$. Results are reported in Tables 3a through 3c, where we have indicated with k the number of partitions in the resulting clusterization, and with $\#$ the trial number. For each value of η , we have highlighted the best result according to the Davies-Bouldin index.

Results reveal the following aspects. As expected, the algorithm exhibits a strong ability to invalidate extra units. Such ability appears to be stronger compared to the SOM, as suggested by the fact that the algorithm has always been able to obtain correct clusterizations of the dataset—i.e. five clusters, associated with extremely good values for the Davies-Bouldin index.

Moreover, this ability is only partially affected by the choice of the initial learning rate η —as Table 3 implies, the RPCL algorithm has been able to obtain correct partitionings of the dataset in the 56.67% of the trials for $\eta = 0.1$, 11.67% for $\eta = 0.3$ and 5% for $\eta = 0.5$. In this respect, a higher learning rate does augment the ability of the rival units to move in the data space, and hence the ability of the algorithm to invalidate extra clusters¹, but the success or failure of the algorithm is ultimately due to the goodness of the initialization of the prototypes. As a consequence, we expect the algorithm to succeed independently of the learning rate as long as the number of initializations tested is large enough.

4 Discussion and Conclusion

In conclusion, all the algorithms tested work reasonably well and produce good clusterizations of the dataset. However, if the main task is to make use of one such methods to discover the number of clusters in a given dataset, the rival-penalized competitive learning approach appears to be more robust and practical, since it exhibits a remarkable ability to invalidate extra units—i.e. clusters—depending on the prototype initialization, provided the number of initializations tested is large enough. Hence, this method comes in handy when the number of clusters of a dataset is unknown.

If the number of clusters is not known exactly but it is known to belong to a range of a few possible values, then the self-organizing map can also guess the correct number of clusters and yield a good clusterization, providing multiple, random initializations are tested and the prototypes are drawn from the input dataset. However, the performance

¹ Note that, in some cases, this ability has reached a point in which the algorithm produced a 1-cluster partitioning, for which the Davies-Bouldin index is structurally not defined and the quadratic error loses its significance.

of the SOM exhibits a strong dependency on the value of the parameters, and finding the optimal values for the radius and the step-length can be a challenging task. In this respect, the rival-penalized competitive learning approach is to be preferred over the SOM. Moreover, the SOM involves greater workload compared to the rival-penalized method or the standard competitive learning method, and hence its use in a context where the spatial organization of the output units is of little or no interest appears to be questionable.

Lastly, if the number of clusters is known in advance and the goal is simply to produce a clusterization of the dataset, the basic competitive learning algorithm works fairly well and is less susceptible to the initialization of the prototypes and does not suffer from the dead unit problem. Additionally, it has the highest performance-to-cost ratio, although the number of clusters and the learning rate can play a crucial role in this respect.

References

1. Forgy, E.W.: Cluster analysis of multivariate data: efficiency vs interpretability of classifications. *Biometrics* 21, 768–769 (1965)
2. Xu, R., Wunsch, D.: Survey of Clustering Algorithms. *IEEE Transactions on Neural Networks* 16, 645–678 (2005)
3. Sugar, C.A., James, G.M.: Finding the number of clusters in a dataset: an information theoretic approach. *Journal of the American Statistical Association* 98, 750–763 (2003)
4. Martinetz, T.M., Berkovich, S.G., Schulten, K.J.: “Neural Gas” network for vector quantization and its application to time-series prediction. *IEEE Transactions on Neural Networks* 4, 558–569 (1993)
5. Budura, G., Botoca, C., Miclău, N.: Competitive Learning Algorithms for Data Clustering. *Electronics and Energetics* 19, 261–269 (2006)
6. Kohonen, T.: Self-organized formation of topologically correct feature maps. *Biological Cybernetics* 43, 59–69 (1982)
7. Erwin, E., Obermayer, A., Schulten, K.: Self-organizing maps: ordering, convergence, properties and energy functions. *Biological Cybernetics* 67, 47–55 (1992)
8. Bezdek, J.C., Pal, N.R.: Some new indexes of cluster validity. *IEEE Transactions on Systems, Man, and Cybernetics* 28, 301–315 (1998)
9. Bienenstock, E.L., Cooper, L.N., Munro, P.W.: *Neuroscience* 2, 32–48 (1982)
10. Kirkpatrick, S., Gelatt, C., Vecchi, M.: Optimization by Simulated Annealing. *Science* 220, 671–680 (1983)
11. Ferrari, S., Ferrigno, G., Piuri, V., Borghese, N.A.: Reducing and Filtering Point Clouds with Enhanced Vector Quantization. *IEEE Transactions on Neural Networks* 18, 161–177 (2007)
12. Uchiyama, T., Arbib, M.A.: An algorithm for competitive learning in clustering problems. *Pattern Recognition* 27, 1415–1421 (1994)
13. Fritzke, B.: A growing Neural Gas network learns topologies. *Advances in Neural Information Processing Systems* 7, 625–632 (1995)
14. Bezdek, J.: *Pattern recognition with fuzzy objective function algorithms*. Plenum Press, New York (1981)
15. King, I., Lau, T.-K.: Non-hierarchical Clustering with Rival Penalized Competitive Learning for Information Retrieval. In: Perner, P., Petrou, M. (eds.) *MLDM 1999*. LNCS (LNAI), vol. 1715, pp. 116–130. Springer, Heidelberg (1999)
16. Bação, F., Lobo, V., Painho, M.: Self-organizing Maps as Substitutes for K-Means Clustering. In: Sunderam, V.S., van Albada, G.D., Sloot, P.M.A., Dongarra, J. (eds.) *ICCS 2005*. LNCS, vol. 3516, pp. 476–483. Springer, Heidelberg (2005)

An Interpretation of the Boundary Movement Method for Imbalanced Dataset Classification Based on Data Quality

Dario Malchiodi

Dipartimento di Informatica, Università degli Studi di Milano, Italy
malchiodi@di.unimi.it

Abstract. This paper describes how the classification of imbalanced datasets through support vector machines using the boundary movement method can be easily explained in terms of a cost-sensitive learning algorithm characterized by giving each example a cost in function of its class. Moreover, it is shown that under this interpretation the boundary movement is measured in terms of the squared norm of the separator's slopes in feature space, thus providing practical insights in order to properly choose the boundary surface shift.

1 Introduction

Many real-world problems cannot be solved directly because no formulation attempting to process them, even in an inefficient way according to the computational complexity theory [1], is known. Despite of this fact, such problem *instances* can be tackled *indirectly* because given a set of candidate solutions it is relatively easy to assert which ones will actually solve the problem and which won't. Refer to these two types of candidate solutions as *positive* and *negative examples*, respectively, and consider for instance the problem of face recognition [2]: while there is no agreement on the mechanisms underlying the way our brain recognizes a given face, it is fairly easy to state with good confidence if the face of known person, say Mr. White, appears on an image. Thus it is possible to build a *dataset* gathering several (positive and negative) examples. Such dataset can be processed by a *machine learning* algorithm having the aim of inferring an automatic classifier able to answer future queries related to the recognition of Mr. White.

Datasets such as those involved in face recognition often share the property of being *imbalanced*, as finding images not representing the face of a given person is far easier than finding images of that person. Thus examples from one class are represented in a sensibly lower quantity than those of the remaining classes. Many real-world instances of the classification problem fall into this special category, such as for instance happens in the fields of fraud detection or fingerprinting recognition. Indeed, it is expectable that when considering, e.g., payment requests issued by a credit card, only a small fraction of these will actually describe a fraud. Likewise, when comparing a latent fingerprint

found on a crime scene against a forensic database, most of the available records will not refer to the found fingerprint.

The problem of imbalanced classification has been considered by various perspectives, including those:

- rebalancing positive and negative examples’ representatives, either undersampling the over-represented class [3], oversampling the under-represented one [4], or performing both operations [5];
- post-processing of the classifiers in output of general-purpose learning algorithms [6];
- focusing only on the under-represented class, to be learnt through unsupervised techniques [7], so that the over-represented class is indirectly inferred through complementation of the under-represented one;
- tailoring learning algorithms for the special case of imbalanced data [8].

This paper focus on the relations between the second and the fourth approach when they are applied to the widely used classification methodology based on support vector machines [9]. When using these tools, the classifier inference is reformulated in terms of the solution of a constrained optimization problem depending on the original examples, and this solution is in turn translated into a hyperplane separating the images of examples in a suitable feature space. In particular, it is shown that the post-processing technique proposed in [6], consisting in shifting the threshold of previously mentioned hyperplane, constitutes the special case of a cost-sensitive algorithm for support vector classification (that is, an algorithm assigning distinct costs to each example and using these costs in order to build the optimization problem to be solved [10]). This equivalence will suggest a practical rule in order to set the shift for the hyperplane learnt through support vector classification in order to account for the amount of class imbalance.

The paper is organized as follows: Sect. 2 briefly reviews the employed cost-sensitive methodology tailored for support vector classification, while Sect. 3 describes the application of this methodology to the problem of imbalanced classification. Finally, Sect. 4 is devoted to concluding remarks.

2 Cost-Sensitive Classification through Support Vector Machines

The proposed approach uses as a starting point the cost-sensitive classification algorithm presented in [11]. This algorithm receives as input a sample

$$\{(x_i, y_i, r_i), i = 1, \dots, m\}, \quad (1)$$

for some $m \in \mathbb{N}$, where the pair $(x_i, y_i) \in X \times \{-1, 1\}$ gathers a *pattern* and a *label*, thus corresponding to a labeled example for the classification problem while r_i (henceforth referred to as *quality*) quantifies the confidence that the association of label y_i to

object x_i in previous example is actually correct. The algorithm is formalized, as usual, through a constrained optimization problem that reads as follows:

$$\begin{aligned}
 \min_{w,b} & \frac{1}{2} w \cdot w + C \sum_{i=0}^m \xi_i \\
 & w \cdot \left(\Phi(x_i) - \frac{r_i}{2} w \right) + b \geq 1 - \xi_i & \forall i : y_i = +1 & \quad (2) \\
 & w \cdot \left(\Phi(x_i) + \frac{r_i}{2} w \right) + b \leq -1 + \xi_i & \forall i : y_i = -1, \\
 & \xi_i \geq 0 & \forall i = 1, \dots, m.
 \end{aligned}$$

In the above formulation Φ is a mapping from X onto a space H , within which the optimization process aims at finding a hyperplane (having w and b respectively as slopes and threshold) separating the images through Φ of patterns associated to positive labels from those associated to negative labels. Slack variables ξ_1, \dots, ξ_m allow this hyperplane to perform mistakes in the separation process (precisely, pattern x_i is misclassified when $\xi_i > 1$), and the value of parameter $C > 0$ balances the two optimization components, namely the number of mistakes and the separator margin (the latter being related to the generalization capability of the inferred classifier [12]). The difference between (2) and the problem at the basis of classical support vector classifier [9] lies in the presence of r_i , whose effect is that of performing a *virtual shift* on the patterns images according to the related quality value. More precisely:

- when $r_i > 0$ point $\Phi(x_i)$ is shifted along the direction normal w.r.t. the separating surface, moving towards the latter with the effect of increasing the distance between the actual classifier and the original pattern position (see Fig. 1(a));
- when $r_i < 0$ the shift occurs in the opposite direction, so that if r_i is sufficiently large pattern x_i will be misclassified (see Fig. 1(b));
- finally, when r_i has a null value nothing changes in the problem formulation.

Applying standard results of duality theory [13] the solution of (2) is linked to that of the following problem:

$$\begin{aligned}
 \max_{\alpha_1, \dots, \alpha_m} & \sum_i \alpha_i - \frac{1}{2(1 + \sum_l \alpha_l r_l)} \sum_{i,j} \alpha_i \alpha_j y_i y_j k(x_i, x_j) \\
 & \sum_i \alpha_i y_i = 0 & (3) \\
 & 0 \leq \alpha_i \leq C & \forall i = 1, \dots, m,
 \end{aligned}$$

where k is the kernel function associated to Φ , defined by $k(x_i, x_j) = \Phi(x_i) \cdot \Phi(x_j)$. More precisely, denoted by $\alpha_1^*, \dots, \alpha_m^*$ the optimal values for (3), the solution of (2) occurs in correspondence of:

$$w^* = \frac{1}{1 + \sum_l \alpha_l^* r_l} \sum_i \alpha_i^* y_i \Phi(x_i), \quad (4)$$

$$b^* = y_i - w^* \cdot \Phi(x_i) + \frac{y_i r_i}{2} w^* \cdot w^* \quad \text{for } i \text{ such that } 0 < \alpha_i^* < C, \quad (5)$$




Fig. 1. Virtually shifting patterns in order to take into account the quality of their corresponding examples: (a) the image of a pattern such that $r_i > 0$ is shifted toward the separating surface of the original problem (plain line), with the effect of increasing the distance between the classifier (dashed line) and the original point position; (b) when $r_i < 0$ the image is shifted in the opposite direction, so that its distance w.r.t. the classifier is reduced.

so that the label of a new point x^{new} can be inferred as $y^{\text{new}} = \text{sign}(w^* \cdot \Phi(x^{\text{new}}) + b^*)$, which translates into:

$$y^{\text{new}} = \text{sign} \left(\frac{1}{1 + \sum_l \alpha_l^* r_l} \sum_i \alpha_i^* y_i k(x_i, x^{\text{new}}) + y_j + \right. \\ \left. - \frac{1}{1 + \sum_l \alpha_l^* r_l} \sum_i \alpha_i^* y_i k(x_i, x_j) + \frac{y_j r_j}{2} \frac{1}{(1 + \sum_l \alpha_l^* r_l)^2} \sum_{i,h} \alpha_i^* \alpha_h^* y_i y_h k(x_i, x_h) \right),$$

where j in the second sum has been chosen so that $0 < \alpha_j^* < C$.

It has been shown that this approach promotes the correct classification of examples when r_i is positive and conversely tends to misclassify examples characterized by a negative r_i [11]. The evident drawback consists in the high nonlinearity of the objective function in (3), while in standard support vector classification this function is quadratic. This fact essentially precludes an efficient processing of the related optimization problem in order to numerically approximate its solution.

3 Classification of Imbalanced Data

An interesting application of the procedure described in Sect. 2 having the additional benefit of being formulated as a quadratic optimization concerns the classification of imbalanced data. Indeed, this can be easily accomplished through association of a fixed positive quality value to each example from the under-represented class and a fixed

negative quality value to the remaining examples. As a result, the classification algorithm will be pushed to correctly classify the under-represented class, possibly at the expense of an erroneous classification of the over-represented one. The simplest way to implement this strategy is that of setting $r_i = y_i$, so that (2) becomes

$$\begin{aligned} \min_{w,b} \frac{1}{2} w \cdot w + C \sum_{i=0}^m \xi_i \\ y_i \left(w \cdot \left(\Phi(x_i) - \frac{1}{2} \right) + b \right) \geq 1 - \xi_i \quad \forall i = 1, \dots, m \quad (6) \\ \xi_i \geq 0 \quad \forall i = 1, \dots, m. \end{aligned}$$

Note how this formulation uses the ‘‘costs’’ r_i in order to directly shape the problem constraint instead of modifying its objective function as in other cost-sensitive approaches to support vector classification. The dual form of this problem simplifies to

$$\begin{aligned} \max_{\alpha_1, \dots, \alpha_m} \sum_i \alpha_i - \frac{1}{2} \sum_{i,j} \alpha_i \alpha_j y_i y_j k(x_i, x_j) \\ \sum_i \alpha_i y_i = 0 \quad (7) \\ 0 \leq \alpha_i \leq C \quad \forall i = 1, \dots, m. \end{aligned}$$

The latter problem coincides with the original support vector classification algorithm [9] whose solution $(w^{\text{svc}}, b^{\text{svc}})$ in the primal space can be efficiently found using any of the available standard tools and techniques (refer for instance to the SMO algorithm [14]). Substituting $r_i = y_i$ in (4-5) shows that the the optimal solution of (6) occurs in correspondence of $w^* = w^{\text{svc}}$ and $b^* = b^{\text{svc}} + \frac{1}{2} \|w^{\text{svc}}\|^2$.

It is also easy to prove that the same analysis can be applied to the general case $r_i = r y_i$ for any $r > 0$, obtaining the optimization problem

$$\begin{aligned} \min_{w,b} \frac{1}{2} w \cdot w + C \sum_{i=0}^m \xi_i \\ y_i \left(w \cdot \left(\Phi(x_i) - \frac{r}{2} \right) + b \right) \geq 1 - \xi_i \quad \forall i = 1, \dots, m \quad (8) \\ \xi_i \geq 0 \quad \forall i = 1, \dots, m, \end{aligned}$$

whose optimal values are:

$$w^* = w^{\text{svc}}, \quad (9)$$

$$b^* = b^{\text{svc}} + r \frac{1}{2} \|w^{\text{svc}}\|^2. \quad (10)$$

The solution in (9-10) corresponds to rediscovering the so-called *boundary movement method* for the classification of imbalanced datasets with support vector machines [6], which captures the intuitive idea of suitably shifting the threshold value for the separating hyperplane in order to account for the under-representation of one class w.r.t. the remaining one. Equation (10) states that this shift can be measured in terms of the

squared norm of the separating hyperplane slopes. The latter quantity represents the optimal value of the first addend in the objective function of (8), which in turn can be expressed as a function of the maximized *margin* of the classifier [9] through an inverse relation property: precisely

$$\mu^* = \frac{2}{\|w^*\|}, \quad (11)$$

where μ^* denotes the maximal margin.

Note also that the shift is always positive: indeed it is expressed as a positive multiple of a squared norm in (10), meaning that the region of X associated to the under-represented class is being enlarged.

The obtained result implicitly sets the following rule in order to preliminary fix the shift amount when applying the boundary movement method:

1. solve the original formulation support vector classification problem (i.e. that not taking into account the data imbalance), obtaining w^{svc} and b^{svc} through standard tools;
2. compute the value $\frac{1}{2}\|w^{\text{svc}}\|^2$;
3. fix the initial shift as r times the quantity computed on previous point, with $r > 0$.

Elementary algebraic considerations involving equations (10) and (11) rediscover the natural assertion that when the optimal margin assumes a big value, a unit move of b^{svc} will correspond to shifting the hyperplane by a small (possibly fractional) number of margins and *vice versa*.

4 Conclusions

This work showed that the boundary movement method used as a post-processing technique for classifiers learnt through the support vector method in order to account for class imbalance can be viewed as a special case of a cost-sensitive modification of the original support vector learning algorithm. The formulation of the latter, where costs are used to shape the optimal problem's constraints rather than its objective function, suggests a rule for setting the boundary movement involving the classifier margin.

References

1. Papadimitriou, C.H.: Computational complexity. Addison-Wesley, Reading (1994)
2. Zhao, W., Chellappa, R., Phillips, P.J., Rosenfeld, A.: Face Recognition: A Literature Survey. *ACM Computing Surveys* 35(4), 399–458 (2003)
3. Tang, Y., Zhang, Y., Chawla, N.V., Krasser, S.: SVMs Modeling for Highly Imbalanced Classification. *IEEE Transactions on Systems, Man, and Cybernetics, Part B* 39(1), 281–288 (2009)
4. Chawla, N., Bowyer, K., Hall, L., Kegelmeyer, W.: SMOTE: Syntetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Reserarch* 16, 321–357 (2002)
5. Stefanowski, J., Wilk, S.: Improving Rule-Based Classifiers Induced by MODLEM by Selective Preprocessing of Imbalanced Data. In: *Proceedings of the Workshop RSKD at European Conference on Machine Learning and Principles of Knowledge Discovery in Databases, Warszawa, September 17-21, pp. 54–65 (2007)*

6. Wu, G., Chang, E.: Class-Boundary Alignment for Imbalanced Dataset Learning. In: ICML 2003 Workshop on Learning from Imbalanced Data Sets (II), pp. 49–56 (2003)
7. Ben-Hur, A., Horn, D., Siegelmann, H.T., Vapnik, V.: Support Vector Clustering. *Journal of Machine Learning Research* 2, 125–137 (2001)
8. Wu, G., Chang, E.Y.: Adaptive Feature-Space Conformal Transformation for Imbalanced-Data Learning. In: Proceedings of the Twentieth International Conference on Machine Learning (ICML 2003), pp. 816–823 (2003)
9. Cortes, C., Vapnik, V.: Support-Vector Networks. *Machine Learning* 20, 121–167 (1995)
10. Elkan, C.: The Foundations of Cost-Sensitive Learning. In: Proceedings of the 17th International Joint Conference on Artificial Intelligence, pp. 973–978 (2001)
11. Apolloni, B., Malchiodi, D., Natali, L.: A Modified SVM Classification Algorithm for Data of Variable Quality. In: Apolloni, B., Howlett, R.J., Jain, L. (eds.) KES 2007, Part III. LNCS (LNAI), vol. 4694, pp. 131–139. Springer, Heidelberg (2007)
12. Schölkopf, B., Smola, A.J.: *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge (2002)
13. Fletcher, R.: *Practical Methods of Optimisations*, 2nd edn. John Wiley & Sons, Chichester (1987)
14. Platt, J.: Fast Training of Support Vector Machines Using Sequential Minimal Optimization. In: Schölkopf, B., Burges, C.J.C., Smola, A.J. (eds.) *Advances in Kernel Methods – Support Vector Learning*, pp. 185–208. MIT Press, Cambridge (1999)

Genetic Algorithm Modeling with GPU Parallel Computing Technology

Stefano Cavuoti¹, Mauro Garofalo², Massimo Brescia^{1,3,*}, Antonio Pescape²,
Giuseppe Longo^{1,4}, and Giorgio Ventre²

¹ Department of Physics, University Federico II, Via Cinthia 6, I-80126 Napoli, Italy
`brescia@oacn.inaf.it`

² Department of Computer Engineering and Systems, University Federico II,
Via Claudio 21, I-80125 Napoli, Italy

³ INAF, Astronomical Observatory of Capodimonte, Via Moiariello 16, I-80131
Napoli, Italy

⁴ Visiting Associate, California Institute of Technology, Pasadena, CA 91125, USA

Abstract. We present a multi-purpose genetic algorithm, designed and implemented with GPGPU / CUDA parallel computing technology. The model was derived from a multi-core CPU serial implementation, named GAME, already scientifically successfully tested and validated on astrophysical massive data classification problems, through a web application resource (DAMEWARE), specialized in data mining based on Machine Learning paradigms. Since genetic algorithms are inherently parallel, the GPGPU computing paradigm has provided an exploit of the internal training features of the model, permitting a strong optimization in terms of processing performances and scalability.

Keywords: genetic algorithms, GPU programming, data mining.

1 Introduction

Computing has started to change how science is done, enabling new scientific advances through enabling new kinds of experiments. They are also generating new kinds of data of increasingly exponential complexity and volume. Achieving the goal of being able to use, exploit and share most effectively these data is a huge challenge. The harder problem for the future is heterogeneity, of platforms, data and applications, rather than simply the scale of the deployed resources. Current platforms require the scientists to overcome computing barriers between them and the data [1].

The present paper concerns the design and development of a multi-purpose genetic algorithm implemented with the GPGPU/CUDA parallel computing technology. The model comes out from the machine learning supervised paradigm, dealing with both regression and classification scientific problems applied on

* Corresponding author.

massive data sets. The model was derived from the original serial implementation, named GAME (Genetic Algorithm Model Experiment) deployed on the DAME [8] Program hybrid distributed infrastructure and made available through the DAMEWARE [9] data mining (DM) web application. In such environment the GAME model has been scientifically tested and validated on astrophysical massive data sets problems with successful results [2]. As known, genetic algorithms are derived from Darwin's evolution law and are intrinsically parallel in its learning evolution rule and processing data patterns. The parallel computing paradigm can indeed provide an optimal exploit of the internal training features of the model, permitting a strong optimization in terms of processing performances.

2 Data Mining Based on Machine Learning and Parallel Computing

Let's start from a real and fundamental assumption: we live in a contemporary world submerged by a tsunami of data. Many kinds of data, tables, images, graphs, observed, simulated, calculated by statistics or acquired by different types of monitoring systems. The recent explosion of World Wide Web and other high performance resources of Information and Communication Technology (ICT) are rapidly contributing to the proliferation of such enormous information repositories. Machine learning (ML) is a scientific discipline concerned with the design and development of algorithms that allow computers to evolve behaviors based on empirical data. A *learner* can take advantage of examples (data) to capture characteristics of interest of their unknown underlying probability distribution. These data form the so called Knowledge Base (KB): a sufficiently large set of examples to be used for training of the ML implementation, and to test its performance. The DM methods, however, are also very useful to capture the complexity of small data sets and, therefore, can be effectively used to tackle problems of much smaller scale [2].

DM on Massive Data Sets (MDS) poses two important challenges for the computational infrastructure: asynchronous access and scalability. With synchronous operations, all the entities in the chain of command (client, workflow engine, broker, processing services) must remain up for the duration of the activity: if any component stops, the context of the activity is lost.

Regarding scalability, whenever there is a large quantity of data, the more affordable approach to making learning feasible relies in splitting the problem in smaller parts (parallelization) sending them to different CPUs and finally combine the results together. So far, the parallel computing technology chosen for this purpose was the GPGPU.

GPGPU is an acronym standing for General Purpose Computing on Graphics Processing Units. It was invented by Mark Harris in 2002, [3], by recognizing the trend to employ GPU technology for not graphic applications. With such term we mean all techniques able to develop algorithms extending computer graphics but running on graphic chips. In general the graphic chips, due to

their intrinsic nature of multi-core processors (many-core) and being based on hundreds of floating-point specialized processing units, make many algorithms able to obtain higher (one or two orders of magnitude) performances than usual CPUs (Central Processing Units). They are also cheaper, due to the relatively low price of graphic chip components.

The choice of graphic device manufacturers, like NVIDIA Corp., was the many-core technology (usually many-core is intended for multi-core systems over 32 cores). The many-core paradigm is based on the growth of execution speed for parallel applications. Began with tens of cores smaller than CPU ones, such kind of architectures reached hundreds of core per chip in a few years. Since 2009 the throughput peak ratio between GPU (many-core) and CPU (multi-core) was about 10:1. Such a large difference has pushed many developers to shift more computing-expensive parts of their programs on the GPUs.

3 The GAME Model

An important category of supervised ML models and techniques, in some way related with the Darwin's evolution law, is known as evolutionary (or genetic) algorithms, sometimes also defined as based on genetic programming [4]. The slight conceptual difference between evolutionary and genetic algorithms is that the formers are problem-dependent, while the latter are very generic.

GAME is a pure genetic algorithm specially designed to solve supervised optimizations problems related with regression and classification functionalities, scalable to efficiently manage MDS and based on the usual genetic evolution methods (crossover, genetic mutation, roulette/ranking, elitism). In order to give a level of abstraction able to make simple to adapt the algorithm to the specific problem, a family of polynomial developments was chosen for GAME model. This methodology makes the algorithm itself easily expandable, but this abstraction requires a set of parameters that allows fitting the algorithm to the specific problem.

From an analytic point of view, a pattern, composed of N features contains an amount of information correlated between the features corresponding to the target value. Usually in a real scientific problem that correlation is *masked* from the noise (both intrinsic to the phenomenon, and due to the acquisition system); but the unknown correlation function can ever be approximated with a polynomial sequence, in which the degree and non-linearity of the chosen function determine the approximation level. The generic function of a polynomial sequence is based on these simple considerations:

Given a generic dataset with N features and a target t , pat a generic input pattern of the dataset, $pat = (f_1, \dots, f_N, t)$ and $g(x)$ a generic real function, the representation of a generic feature f_i of a generic pattern, with a polynomial sequence of degree d is:

$$G(f_i) \cong a_0 + a_1g(f_i) + \dots + a_dg^d(f_i) \quad (1)$$

Hence, the k -th pattern (pat_k) with N features may be represented by:

$$Out(pat_k) \cong \sum_{i=1}^N G(f_i) \cong a_0 + \sum_{i=1}^N \sum_{j=1}^d a_j g^j(f_i) \quad (2)$$

Then target t_k , concerning to pattern pat_k , can be used to evaluate the approximation error of the input pattern to the expected value:

$$E_k = (t_k - Out(pat_k))^2 \quad (3)$$

If we generalize the expression (2) to an entire dataset, with NP number of patterns ($k = 1, \dots, NP$), at the end of the *forward* phase (batch) of the GA, we obtain NP expressions (2) which represent the polynomial approximation of the dataset.

In order to evaluate the fitness of the patterns as extension of (3), the Mean Square Error (MSE) or Root Mean Square Error (RMSE) may be used.

Then we define a GA with the following characteristics:

- The expression (2) is the fitness function;
- The array (a_0, \dots, a_M) defines M genes of the generic chromosome (initially they are generated random and normalized between -1 and +1);
- All the chromosomes have the same size (constrain from a classic GA);
- The expression (3) gives the standard error to evaluate the fitness level of the chromosomes;
- The population (genome) is composed by a number of chromosomes imposed from the choice of the function $g(x)$ of the polynomial sequence.

About the last item, this number is determined by the following expression:

$$NUM_{chromosomes} = (B \cdot N) + 1 \quad (4)$$

where N is the number of features of the patterns and B is a multiplicative factor that depends from the $g(x)$ function, which in the simplest case is just 1, but can arise to 3 or 4 in more complex cases. The parameter B also influences the dimension of each chromosome (number of genes):

$$NUM_{genes} = (B \cdot d) + 1 \quad (5)$$

where d is the degree of the polynomial. For example if we use the trigonometric polynomial expansion, given by the following expression (hereinafter polytrigo),

$$g(x) = a_0 + \sum_{m=1}^d a_m \cos(mx) + \sum_{m=1}^d b_m \sin(mx) \quad (6)$$

in order to have 200 patterns composed by 11 features, the expression using (2) with degree 3, will become:

$$Out(pat_{k=1\dots 200}) \cong \sum_{i=1}^{11} G(f_i) \cong a_0 + \sum_{i=1}^{11} \sum_{j=1}^3 a_j \cos(jf_i) + \sum_{i=1}^{11} \sum_{j=1}^3 b_j \sin(jf_i) \quad (7)$$

In the last expression we have two groups of coefficients (sin and cosine), so B will assume the value 2. Hence the generic genome (population at a generic evolution stage), will be composed by 23 chromosomes, given by equation (4), each one with 7 genes $[a_0, a_1, a_2, a_3, b_1, b_2, b_3]$, given by equation (5), with each single gene (coefficient of the polynomial) in the range $[-1, +1]$.

In the present project, the idea is to build a GA able to solve supervised crispy classification and regression problems, typically related to an high-complexity parameter space where the background analytic function is not known, except for a limited number of couples of input-target values, representing valid solutions to a physical category of phenomena. A typical case is to classify astronomical objects based on some solution samples (the KB) or to predict new values extracted by further observations. To accomplish such behavior we designed a function (a polynomial expansion) to combine input patterns. The coefficients of such polynomials are the chromosome genes. The goal is indeed to find the best chromosome so that the related polynomial expansion is able to approximate the right solutions to input pattern classification/regression. So far, the fitness function for such representation consists of the training error, obtained as absolute difference between the polynomial output and the target value for each pattern. Due to the fact that we are interested to find the minimum value of the error, the fitness is calculated as the complement of the error (i.e. 1-error) and the problem is reduced to find the chromosome achieving the maximum value of fitness.

4 The GPU-Based GAME Implementation

In all execution modes (use case), GAME exploits the polytrigo function (6), consisting in a polynomial expansion in terms of sum of sins and cosines. Specifically in the training use case, corresponding to the GA building and consolidation phase, the polytrigo is used at each iteration as the transformation function applied to each chromosome to obtain the output on the problem input dataset, and indirectly also to evaluate the fitness of each chromosome. It is indeed one of the critical aspects of the serial algorithm to be investigated during the parallelization design process.

Moreover, after having calculated the fitness function for all genetic population chromosomes, this information must be back-propagated to evolve the genetic population. This back and forth procedure must be replicated as many times as it is the training iteration number or the learning error threshold, both decided and imposed by the user at setup time of any experiment. The direct consequence of the above issues is that the training use case takes much more execution time than the others (such as test and validation), and therefore is the one we are going to optimize.

Main design aspect approaching the software architecture analysis for the GPU is the partition of work: i.e. which work should be done on the CPU vs. the GPU. We have identified the time consuming critical parts to be parallelized by executing them on the GPU. They are the generation of random chromosomes and the calculation of the fitness function of chromosomes. The key principle is that we need to perform the same instruction simultaneously on as much data as possible. By adding the number of chromosomes to be randomly generated in the initial population as well as during each generation, the total number of involved elements is never extremely large but it may occur with a high frequency. This is because also during the population evolution loop a variable number of chromosomes are randomly generated to replace older individuals. To overcome this problem we may generate a large number of chromosomes randomly *una tantum*, by using them whenever required. On the contrary, the evaluation of fitness functions involves all the input data, which is assumed to be massive datasets, so it already has an intrinsic data-parallelism. Since CUDA programming involves code running concurrently on a host with one or more CPUs and one or more CUDA-enabled GPU, it is important to keep in mind that the differences between these two architectures may affect application performance to use CUDA effectively. The function `polytrigo` takes about three-quarters of the total execution time of the application, while the total including child functions amounts to about 7/8 of total time execution. This indeed has been our first candidate for parallelization. In order to give a practical example, for the interested reader, we report the source code portions related to the different implementation of the `polytrigo` function, of the serial and parallelized cases.

C++ serial code for polytrigo function (equation 6):

```
for (int i = 0; i < num_features; i++) {
    for (int j = 1; j <= poly_degree; j++) {
        ret += v[j] * cos(j * input[i]) + v[j + poly_degree] *
            * sin(j * input[i]); } }
```

CUDA C (Thrust) parallelized code for polytrigo function (equation 6):

```
struct sinFunctor { __host__ __device__
    double operator()(tuple <double, double> t) {
        return sin(get < 0 > (t) * get < 1 > (t)); };
};
struct cosFunctor { __host__ __device__
    double operator()(tuple <double, double> t) {
        return cos(get < 0 > (t) * get < 1 > (t)); };
};
thrust::transform(thrust::make_zip_iterator(
    thrust::make_tuple(j.begin(), input.begin())),
    thrust::make_zip_iterator(
        thrust::make_tuple(j.end(), input.end())),
    ret.begin(), sinFunctor(), cosFunctor());
```

Noting that, while the vector $v[]$ is continuously evolving, *input[]* (i.e. the elements of the input dataset) are being used in calculation of *ret* at each iteration but they are never altered. We rewrite the function by calculating in advance the sums of sines and cosines, storing the results in two vectors and then use them in the function `polyTrigo()` at each iteration. This brings huge benefits because we calculate trigonometric functions, which are those time consuming, only once instead of at every iteration and exploit the parallelism on large amount of data because it assumes that we have large input datasets.

From the time complexity point of view, by assuming to have as many GPU cores as population chromosomes, the above CUDA C code portion would take constant time, instead of polynomial time required by the corresponding C++ serial code.

5 The Experiment

In terms of experiments, the two CPU versions of GAME, the original and an optimized version of the serial algorithm (hereinafter *serial* and *Opt* respectively), together with the final version for GPU (hereinafter *ELGA*), have been compared basically by measuring their performance in terms of execution speed, by also performing an intrinsic evaluation of the overall scientific performances. The optimized algorithm is the serial version adapted by modifying the code portions which are candidate to be parallelized in the final GPU release.

Initially, the tests have been organized by distinguishing between classification and regression functional modes. By analyzing early trials, however, it resulted that the performance growth was virtually achieved in both cases. So far, we limit here the discussion details to a classification experiment, done in the astrophysical context.

The scientific problem used here as a test bed for data mining application of the GAME model is the search (classification) of Globular Cluster (GC) populations in external galaxies [2]. This topic is of interest to many astrophysical fields: from cosmology, to the evolution of stellar systems, to the formation and evolution of binary systems.

The dataset used in this experiment consists in wide field HST observations of the giant elliptical NGC1399 in the Fornax cluster, [5]. The subsample of sources used to build our Base of Knowledge, to train the GAME model is composed by 2100 sources with all photometric and morphological information, [2]. Finally, our classification dataset consisted of 2100 patterns, each composed by 11 features (including the two targets, corresponding to the classes GC and not GC used during the supervised training phase).

The performance was evaluated on several hardware platforms. We compared our production GPU code with a CPU implementation of the same algorithm. The benchmarks were run on a 2.0 GHz Intel Core i7 2630QM quad core CPU running 64-bit Windows 7 Home Premium SP1. The CPU code was compiled using the Microsoft C/C++ Optimizing Compiler version 16.00 and GPU benchmarks were performed using the NVIDIA CUDA programming toolkit version

4.1 running on several generations of NVIDIA GPUs GeForce GT540M. As execution parameters were chosen combinations of:

- Max number of iterations: 1000, 2000, 4000, 10000, 20000 and 40000;
- Order (max degree) of polynomial expansion: 1, 2, 4 and 8;

The other parameters remain unchanged for all tests:

- error function (fitness): MSE with threshold = 0.001;
- Selection criterion: RANKING and ROULETTE;
- Crossover probability rate 0.9 and mutation probability rate 0.2;
- Elitism chromosomes at each evolution: 2.

For the scope of the present experiment, we have preliminarily verified the perfect correspondence between CPU- and GPU-based implementations in terms of classification performances. In fact, the scientific results for the CPU-based algorithm have been already evaluated and documented in a recent paper [2], where the CPU version of GAME were also compared with other ML models, provided by our team.

Referring to the best results as described in [2], for the serial code version we obtained the following percentages on a dataset consisting of 2100 patterns each composed by 7 column features:

- Classification accuracy = 86.4%;
- Completeness = 78.9%;
- Contamination = 13.9%;

In both optimized serial and parallelized version, we obtained, as expected, the same values, slightly varying in terms of least significant digit, trivially motivated by the intrinsic randomness of the genetic algorithms.

Here we investigated the analysis of performances in terms of execution speed. By using the defined metrics we compared the three versions of GAME implementation, under the same setup conditions. As expected, while the two CPU-based versions, serial and Opt, appear comparable, there is a quite shocking difference with the GPU-based version ELGA. The diagram of Fig. 1 reports the direct comparisons among the three GAME versions, by setting a relatively high degree of the polynomial expansion which represents the evaluation function for chromosomes.

We performed also other tests, by varying the polynomial degree. The trends show that the execution time increases always in a linear way with the number of iterations, once fixed the polynomial degree. This is what we expected because the algorithm repeats the same operations at each iteration. The GPU-based version speed is always at least one order of magnitude less than the other two implementations. We remark also that the classification performances of the GAME model increases by growing the polynomial degree, starting to reach good results from a value equal to 4. Exactly when the difference between CPU and GPU versions starts to be 2 orders of magnitude.




Fig. 1. Comparison among the GAME implementations with the polynomial degree = 8




Fig. 2. Speedup comparison among GAME CPU implementations against the GPU version

In the diagram of Fig. 2, the GPU version is compared against the CPU implementations. As shown, the speedup increases proportionally with the increasing of the polynomial degree. The diagram shows that for the average speed in a range of iterations from 1000 to 40000, the ELGA algorithm exploits the data parallelism as much data are simultaneously processed. As previously mentioned, an increase of maximum degree in the polynomial expansion leads to an increase in the number of genes and consequently to a larger population matrix. The GPU algorithm outperforms the CPU performance by a factor ranging from $8x$

to $200x$ in the not optimized (serial) case and in a range from $6x$ to $125x$ in the optimized case (Opt), enabling an intensive and highly scalable use of the algorithm that were previously impossible to be achieved with a CPU.

6 Conclusions

We investigated the state of the art computing technologies, by choosing the one best suited to deal with a wide range of real physical problems. A multi-purpose genetic algorithm (GA) implemented with GPGPU/CUDA parallel computing technology has been designed and developed. The model comes from the paradigm of supervised machine learning, addressing both the problems of classification and regression applied on massive data sets.

The model was derived from a serial implementation named GAME, deployed on the DAME Program, [6] and [7], hybrid distributed infrastructure and already scientifically tested and validated on astrophysics massive data sets problems with successful results. Since GAs are inherently parallel, the parallel computing paradigm has provided an exploit of the internal training features of the model, permitting a strong optimization in terms of processing performances. We described our effort to adapt our genetic algorithm for general purpose on GPU. We discussed the efficiency and computational costs of various components involved that are present in the algorithm. Several benchmark results were shown. The use of CUDA translates into a $75x$ average speedup, by successfully eliminating the largest bottleneck in the multi-core CPU code. Although a speedup of up to $200x$ over a modern CPU is impressive, it ignores the larger picture of use a Genetic Algorithm as a whole. In any real-world the dataset can be very large (those we have previously called Massive Data Sets) and this requires greater attention to GPU memory management, in terms of scheduling and data transfers host-to-device and *vice versa*. Moreover, the identical results for classification functional cases demonstrate the consistency of the implementation for the three different computing architectures, enhancing the scalability of the proposed GAME model when approaching massive data sets problems.

Finally, the very encouraging results suggest to investigate further optimizations, like: (i) moving the formation of the population matrix and its evolution in place on the GPU. This approach has the potential to significantly reduce the number of operations in the core computation, but at the cost of higher memory usage; (ii) exploring more improvements by mixing Thrust and CUDA C code, that should allow a modest speedup justifying development efforts at a lower level; (iii) use of new features now available on NVIDIA Fermi architecture, such as faster atomics and more robust thread synchronization and multi GPUs capability.

Acknowledgments. The experiment has been done by means a M.Sc. degree in Informatics Engineering arisen from a Collaboration among several Italian academic institutions. The hardware resources have been provided by Dept. of

Computing Engineering and Systems and S.Co.P.E. GRID Project infrastructure of the University Federico II of Naples. The data mining model has been designed and developed by DAME Program Collaboration. MB wishes to thank the financial support of PRIN-INAF 2010, "Architecture and Tomography of Galaxy Clusters". This work has been partially funded by LINCE project of the F.A.R.O. programme jointly financed by the Compagnia di San Paolo and by the Polo delle Scienze e delle Tecnologie of the University of Napoli "Federico II" and it has been carried out also thanks to a hardware donation in the context of the NVIDIA Academic Partnership program. The authors also wish to thank the financial support of Project F.A.R.O. III Tornata, from University Federico II of Naples.

References

1. Fabbiano, G., Calzetti, D., Carilli, C., Djorgovski, S.G.: Recommendations of the VAO Science Council. E-print arXiv:1006.2168v1 [astro-ph.IM] (2010)
2. Brescia, M., Cavuoti, S., Paolillo, M., Longo, G., Puzia, T.: The Detection of Globular Clusters in Galaxies as a data mining problem. MNRAS MN-11-2315-MJ.R1 (2012)
3. Harris, M.J.: Real-Time Cloud Simulation and Rendering. University of North Carolina Technical Report TR03-040 (2003)
4. Mitchell, M.: An Introduction to Genetic Algorithms. The MIT Press, Cambridge (1998)
5. Paolillo, M., Puzia, T.H., Goudfrooij, P., Zepf, S.E., Maccarone, T.J., Kundu, A., Fabbiano, G., Angelini, L.: Probing the GC-LMXB Connection in NGC 1399: A Wide-field Study with the Hubble Space Telescope and Chandra. ApJ 736, 90 (2011)
6. Brescia, M., Longo, G., Djorgovski, G.S., Cavuoti, S., D'Abrusco, R., Donalek, C., et al.: DAME: A Web Oriented Infrastructure for Scientific Data Mining & Exploration. E-print arXiv:1010.4843v2 [astro-ph.IM] (2010)
7. Brescia, M., Cavuoti, S., Djorgovski, G.S., Donalek, C., Longo, G., Paolillo, M.: Extracting knowledge from massive astronomical data sets. In: Barrosaro, L.M., et al. (eds.) Astrostatistics and Data Mining in Large Astronomical Databases, E-print arXiv:1109.2840v1. Springer Series on Astrostatistics, 15 pages (2011)
8. DAME Program official website, <http://dame.dsf.unina.it>
9. DAMEWARE Web Application entry page, http://dame.dsf.unina.it/beta_info.html

An Experimental Evaluation of Reservoir Computation for Ambient Assisted Living

Davide Bacciu¹, Stefano Chessa¹, Claudio Gallicchio¹,
Alessio Micheli¹, and Paolo Barsocchi²

¹ Dipartimento di Informatica, Università di Pisa, Largo B. Pontecorvo,
3 - 56127 Pisa, Italy

{bacciu,ste,gallicch,micheli}@di.unipi.it

² ISTI-CNR, Pisa Research Area, Via Morruzzi - 56124 Pisa, Italy
paolo.barsocchi@isti.cnr.it

Abstract. In this paper we investigate the introduction of Reservoir Computing (RC) neural network models in the context of AAL (Ambient Assisted Living) and self-learning robot ecologies, with a focus on the computational constraints related to the implementation over a network of sensors. Specifically, we experimentally study the relationship between architectural parameters influencing the computational cost of the models and the performance on a task of user movements prediction from sensors signal streams. The RC shows favorable scaling properties results for the analyzed AAL task.

Keywords: Reservoir Computing, Echo State Networks, Wireless Sensor Networks, Ambient Assisted Living.

1 Introduction

The aim of Ambient Assisted Living (AAL) [8] applications is to integrate different technologies to improve the quality of life of elders and disable people, by assisting them in the environments where they live and work. Recently, the EU FP7 RUBICON [1] (Robotic UBiquitous COgnitive Network) project [1,3], has proposed the use of self-adaptive robotic ecologies to approach tasks in the field of AAL. A robotic ecology consists in a network of heterogeneous devices including mobile robots, wireless sensor networks (WSNs) [5] and actuators. The objective of RUBICON is to integrate self-sustaining learning solutions to provide cheap, adaptive and efficient coordination of the robotic ecology. RUBICON is based on the interplay among different layers implementing state-of-the-art techniques in machine learning, WSNs, cognitive robotics and agent control systems. Learning in a network of distributed devices with very limited computational capabilities (such are WSNs), raises novel challenges related to the effectiveness and the efficiency of the learning models used to handle the temporal data sensed by

¹ <http://www.fp7rubicon.eu/>

the ecology. Reservoir Computing (RC) [14], and in particular the Echo State Network (ESN) [12,11] model, represent an emerging paradigm for modeling Recurrent Neural Networks (RNNs), and offer in principle an interesting trade-off between computational efficiency and the ability of learning in domains of temporal sequences. RC is therefore identified as a potentially suitable approach for the implementation of learning models on-board the nodes of robotic ecologies, although the parameters of the RC networks should be tailored to this specific applicative context. Moreover, in [10,4] it was already experimentally shown that ESNs are particularly suitable for treating information gathered by WSNs for tasks related to the prediction of indoor user movements. The problem of reducing the computational requirements of RC implementations (e.g. [15,9]), in particular for their embedding into WSN nodes [6] is currently a topic of active research.

In this paper we present an experimental study of the relation between the performance and the implementation cost of an RC system in a real-world AAL task consisting in predicting user movements in indoor environments. Such study extends the work in [10,4] to consider the effect of different components of the RC architecture such as the number of reservoir units and the weight encoding. The proposed investigation is useful to understand the potentiality of the RC approach for practical implementations on the nodes of robotic ecologies.

2 Learning in RUBICON Robotic Ecology

In this Section, we characterize the RC based Learning Layer in the RUBICON ecology system. By embedding learning functionalities on-board the nodes of the ecology, the role of the Learning Layer in RUBICON is to supply the general purpose learning infrastructure necessary for achieving self-sustaining learning capabilities.

The purpose of the Learning Layer (LL) is to respond to the need of adaptivity and analysis of temporal context (including the sensor data dynamics) of the RUBICON environments, by providing learning mechanisms that are used by higher (control and cognitive) layers of RUBICON architecture. Specifically, its role is to:

- recognize and detect relevant sensed information by providing predictions which depend on the temporal history of the input signals;
- analyze and process sensed information to extract refined goal-significant information (e.g. information fusion, event recognition).

To this purpose, a networked learning infrastructure, named Learning Network (LN), is built on top of the robotic ecology in order to provide the core learning services to the higher levels of the RUBICON architecture. The LN is a *flexible environmental memory* that serves as a task driven model of the environment that can readily be shared by new nodes connecting to RUBICON. This allows a straightforward sharing of the learned experience.

In the LN design, each node hosts a RC network composed by a number of artificial neurons. Neurons are connected by remote synapses (implemented by

means of communication channels) with neurons residing on other nodes, thus creating a distributed, artificial recurrent neural network. The same synapses are also used to interconnect the output of the LN with the control and cognitive layers of RUBICON. With this approach, each neuron in the local learning model may be considered as a node of a distributed reservoir; the instantaneous state of such reservoir thus incorporates the combined knowledge attained by each sensor node, while reflecting the dynamics of the overall RUBICON. The design and development of the LN is strongly influenced by the issues of scalability and efficiency, especially with regards to the limitations of the nodes (most of which are wireless sensors) in terms of computation, communication and energy constraints. In particular, the design of the LN makes use of RNNs, and specifically of RC models, due to their modular, networked structure which naturally adapts to the distributed nature of the RUBICON ecology, while their recurrent nature allows to effectively capture the dynamics of the system processes. At the same time, the neural paradigm offers a natural approach to robustly cope with noisy sensed data.

Preliminary studies of RC used in combination with sensors for localization applications aimed at forecasting users movements can be found in [10] and [4]. Differently from previous works, in the following we extend the study to consider different components of the RC architecture.

3 Echo State Network Architecture

An ESN [12],[11] is made up of an input layer, a reservoir and a readout with N_U , N_R and N_Y units, respectively (see Fig. 1). The reservoir is a large, sparsely connected, non linear, *untrained* recurrent hidden layer, used to compute an input-driven *fixed* contractive encoding of the input sequences into a state space. The readout is a linear, feed-forward output tool, responsible for the output computation and representing the only trained component of the ESN architecture. In standard ESNs, sigmoid reservoir units are used, e.g. implementing *tanh* activation function. In this paper, we take into consideration leaky integrator ESNs (LI-ESNs) [13], in which leaky integrator reservoir units are used, applying an exponential moving average to the reservoir state values. The use of leaky




Fig. 1. The ESN architecture

integration of reservoir states in LI-ESNs, implies a better handling of input sequences changing slowly with respect to the sampling frequency [13,14,2], and results in an RC model particularly suitable for treating real-world RSS input signals, as experimentally shown in [10].

Given an input sequence $\mathbf{s} = [\mathbf{u}(1), \dots, \mathbf{u}(n)]$ over the real input space \mathbb{R}^{N_U} , for each time step t the reservoir of a LI-ESN computes the state transition function: $\mathbf{x}(t) = (1 - a)\mathbf{x}(t-1) + af(\mathbf{W}_{in}\mathbf{u}(t) + \hat{\mathbf{W}}\mathbf{x}(t-1))$, where $\mathbf{x}(t) \in \mathbb{R}^{N_R}$ is the state of the network at time t , $\mathbf{W}_{in} \in \mathbb{R}^{N_R \times N_U}$ is the input-to-reservoir weight matrix (possibly including also a bias term), $\hat{\mathbf{W}} \in \mathbb{R}^{N_R \times N_R}$ is the recurrent reservoir weight matrix, f is a component-wise applied activation function (we use $f \equiv \tanh$) and $a \in [0, 1]$ is a *leaking rate* parameter, controlling the speed of reservoir dynamics (smaller values of a imply slower reservoir dynamics). Note that for $a = 1$, LI-ESN state transition function reduces to the case of standard ESN. The output is computed by the readout tool, by linearly combining the activation of the reservoir units. In the following, we restrict our consideration to the case of binary classification tasks on sequences, which is of a particular interest for this paper. In this case, the readout is applied only when the input sequence has been completely seen and the reservoir encoding process has terminated. For each input sequence \mathbf{s} of length n , the readout computes $y(\mathbf{s}) = \text{sgn}(\mathbf{W}_{out}\mathbf{x}(n))$, where $\mathbf{W}_{out} \in \mathbb{R}^{N_Y \times N_R}$ is the reservoir-to-readout weight matrix (possibly including also a bias term), $y(\mathbf{s}) = \{-1, +1\}$ is the output classification computed for the input sequence \mathbf{s} , and sgn is a sign threshold function.

The reservoir of a LI-ESN is initialized in order to satisfy the *Echo State Property* (ESP) [11,12], i.e. the network state should asymptotically depend only on the driving input sequence [9]. A sufficient and a necessary condition for the ESP are provided in literature [11]. A necessary condition for the ESP is typically considered in RC applications [13,14,11], i.e. $\rho(\hat{\mathbf{W}}) < 1$, where $\hat{\mathbf{W}} = (1 - a)\mathbf{I} + a\hat{\mathbf{W}}$, and $\rho(\hat{\mathbf{W}})$ is the *spectral radius* of $\hat{\mathbf{W}}$. Accordingly, $\hat{\mathbf{W}}$ is randomly initialized and then rescaled to meet the condition on ρ , where values of ρ close to 1 are generally used [14,12,16]. Weight values in \mathbf{W}_{in} are chosen from a random (uniform) distribution over $[-scale_{in}, scale_{in}]$, where $scale_{in}$ is an *input scaling* parameter.

The readout of a LI-ESN is typically trained off-line by using Moore-Penrose pseudo-inversion or ridge regression (see e.g. [10,14]).

Beside the choice of the spectral radius, one of the most relevant parameters in RC applications is the number of reservoir units. Indeed, larger reservoirs often lead to better network performances (e.g. [16,9]). Notice that the cost of ESN application (both in time and space) increases with the reservoir dimension². Another interesting aspect is related to the weight dimension, i.e. the number of bits used to represent each weight in \mathbf{W}_{in} and $\hat{\mathbf{W}}$. Typically, such weights are randomly initialized to assume values in an infinite real interval, and in practice floating-points with double precision are used, requiring 64 bits per

² The cost is quadratically bounded in general and, due to the sparsity of reservoir connectivity, it can be linear under the assumption of a fixed maximum number of connections for each reservoir unit.

weight memorization. In the following, this setting is referred to as LI-ESN *full weight encoding*. In this paper we also consider the alternative strategy of assuming a small finite set of possible non-zero weights in the untrained RC matrices \mathbf{W}_{in} and $\hat{\mathbf{W}}$. The number of different non-zero values is denoted by N_w . In this case, referred in the following as LI-ESN with *reduced weight encoding*, each weight value can still be a double precision floating-point, but a small number of bits (i.e. $\lceil \log_2 N_w \rceil$) are sufficient to encode each weight in the network matrices.

The issue of minimizing the cost for the network memorization is of a fundamental relevance in view of the embedding of the RC networks directly on the motes. Indeed, it is worth to note that the motes of a WSN usually host a very limited *total* amount of RAM memory (usually in the order of 8-10 Kbytes). We therefore experimentally assess the impact on the predictive performance of RC networks due to the variation of the number of reservoir units and the type of weight encoding used (i.e. full or reduced), specifically in a task of user movement prediction type.

4 Experiments

Real WSN data for our experiments was collected during a measurement campaign at the first floor of the the ISTI institute of CNR in the Pisa Research Area, in Italy. The environment comprises 4 rooms organized into pairs of coupled rooms (denoted as Room 1 and Room 2 in each couple), with front doors separated by a hallway. Rooms contained typical office furniture including chairs, cabinets, monitors and desks (see Fig. 2), representing harsh conditions for indoor wireless communications. In each couple of rooms, we set up a WSN composed of 5 IRIS motes 7: 4 fixed sensors, or *anchors*, and 1 *mobile* sensor worn by the user. Sensors embedded a Chipcon AT86RF230 radio subsystem implementing the IEEE 802.15.4 standard. The user moved in the environment according to the 6 possible paths illustrated in Fig. 2. Following a curved trajectory the user remained in the same room, while straight trajectories led to a room change. Each measurement gathered the received signal strength (RSS) information between the anchors and the mobile at a constant frequency of 8 Hz, from the starting point of each movement until the user reached a *marker* point (denoted as M in Fig. 2) located at 60 cm from the door. The RSS data was used to set up a dataset for a binary classification task in which RSS sequences from the different anchors are used as input data and the target consists in +1 if the corresponding user movement led to a room change, and -1 otherwise. In practice, the classification tasks consists in predicting if the user is about to change room or not based on the history of the RSS information until the marker. Note that the marker M is the same for all the movements, hence it not possible to distinguish the different paths based only on the RSS values collected at M. The dataset contains 210 input sequences, where 104 of them where collected in the first couple of rooms, and the remaining 106 in the second couple. Using a stratified sampling over the different movement paths, the dataset was divided into a training and a test set comprising 168 and 42 sequences, respectively.

We run experiments for all the 15 possible configurations of the number of anchors considered, i.e. for $N_{anchors} = 1, 2, 3, 4$. We used reservoirs with $N_R \in \{10, 20, 50, 100, 300, 500\}$ units, 10% of connectivity, input scaling $scale_{in} = 1$ and spectral radius $\rho = 0.99$. In the reduced weight encoding setting, each non-zero weight value in \mathbf{W}_{in} and $\hat{\mathbf{W}}$ was randomly chosen in a small weight alphabet with $N_w = 8$ possible values, uniformly sampled in the interval $[-0.4, 0.4]$, thus leading to just a 3 bit encoding per weight memorization. For every choice of the reservoir hyper-parametrization, we considered 10 independent (random guessed) reservoirs (and the results were averaged over the 10 guesses). The readout was trained using pseudo-inversion and ridge-regression, with regularization parameter $\lambda_r \in \{10^{-1}, 10^{-3}, 10^{-5}\}$. The readout regularization was chosen by model selection with stratified holdout validation, using $\approx 30\%$ of the training set as validation set.

Fig. 3 shows the averaged test accuracy of LI-ESN with full weight encoding, for increasing reservoir dimension N_R and varying the number of anchors used $N_{anchors}$. For any value of $N_{anchors}$ considered, results in Fig. 3 are averaged (and standard deviations are computed) over the number of possible configurations of the anchors. It can be noticed that the RC networks can reach excellent test performances, and the test accuracy scales very well with both the reservoir dimension and the number of anchors. As already pointed out in 10.4, the test accuracy of the LI-ESN networks is improved when more anchors are progressively considered in the WSN setting. The best performance (test accuracy up to 97.1%) is obtained for $N_{anchors} = 4$, although very good performances are achieved for simpler WSN settings with a fewer number of anchors (up to 88.6%, 91.5% and 95.4%, for 1, 2 and 3 anchors, respectively). Table 1 details the training and test accuracies (and standard deviations over the reservoir guesses) in the case $N_{anchors} = 4$, for increasing reservoir dimensionality. The best test accuracy, i.e. 97.1%, is obtained for $N_R = 500$, corresponding to test sensitivity and specificity of 99.5% and 95.0%, respectively. Moreover, from Fig. 3 it is clear that for every choice of $N_{anchors}$, the LI-ESN performance is improved as larger reservoirs are considered (adopting regularization in readout training in order




Fig. 2. Prototypical environment where the RSS measurements have been collected, showing the positions of the anchors of the WSN in one couple of rooms and the possible user movement paths




Fig. 3. Mean test accuracy (in %) and standard deviation of LI-ESNs with *full weight encoding*, varying the reservoir dimension and the number of anchors considered

Table 1. Mean accuracy (in %) and standard deviation of LI-ESNs with *full weight encoding* for $N_{anchors} = 4$

	Reservoir Dimension (N_R)					
	10	20	50	100	300	500
Training	76.1(±3.1)	86.7(±2.5)	96.4(±2.1)	99.8(±0.3)	99.0(±0.5)	99.9(±0.2)
Test	74.0(±5.6)	85.2(±4.6)	90.2(±3.9)	92.1(±3.2)	96.9(±1.5)	97.1(±2.3)

to avoid overfitting). At the same time, it can also be observed that the test performance tends to be saturated as the number of reservoir units is increased. Indeed, when a sufficiently large reservoir is considered, a further increase of the reservoir dimension can only slightly improve the test accuracy. This means that small enough reservoirs can be sufficient in practice and thus a small amount of memory can be sufficient for storing the network parameters on the nodes of the WSN. For instance (see Table 1), for $N_{anchors} = 4$, the test performance with $N_R = 500$ units is 97.1%, while with smaller reservoirs e.g. with $N_R = 100$ and $N_R = 50$, the test performance is 92.1% and 90.2%, respectively, which are still very satisfactory results for this kind of tasks (from noisy input data).

The averaged test accuracy for LI-ESNs with reduced weight encoding is shown in Fig. 4, varying the reservoir dimensionality and the number of anchors. Analogous considerations can be done as for the case of full weight encoding. RC networks achieve very good test accuracy, which scales well with the number of anchors in the environment and the reservoir dimension. Also in the case of reduced weight encoding, a saturating effect of the networks performance can be observed for increasing the number of reservoir units. Table 2 reports the training




Fig. 4. Mean test accuracy (in %) and standard deviation of LI-ESNs with *reduced weight encoding*, varying the reservoir dimension and the number of anchors considered

and test accuracies achieved by LI-ESNs in correspondence of $N_{anchors} = 4$ and varying the number of reservoir units. The best test performance is obtained for $N_R = 500$, corresponding to a test accuracy of 97.6%, test sensitivity of 99.0% and test specificity of 96.4%. Very interestingly, comparing the performances obtained by LI-ESNs with full and reduced weight encodings (see Figures 3 and 4, and Tables 1 and 2), we can observe that in the two cases, the accuracies are substantially the same (except for small, statistical fluctuations) for correspondent experimental settings. As a result, the use of a reduced weight encoding scheme does not have a relevant practical effect on the LI-ESN model performances, and in particular it does not strictly depend on the reservoir dimension, due to the prevailing role of the number of reservoir units on the accuracy results³. This means that in practice, using a small finite set of reservoir weights, thus a few bits per weight encoding, is sufficient for tackling the task at hand, and the small amount of memory typically available on the motes of a WSN can be enough for ESN embedding. To give an idea of the possible reduction in memory requirements, for storing the weights in $\hat{\mathbf{W}}$ of a 100-dimensional reservoir, roughly 8 Kbytes of memory are needed in case of full weight encoding, while roughly 800 bytes are sufficient in case of reduced weight encoding. For 50-dimensional reservoirs, the memory requirement reduces from roughly 2 Kbytes (for full weight encoding) to roughly 250 bytes (for reduced weight encoding).

³ Note that the number of bits used in our experiments for the reduced weight encoding scheme, i.e. $\log_2(N_w) = 3$ bits, is even smaller than what could be necessary for embedded implementations of our RC system on-board the motes of a WSN. Considering a further reduction of the number of bits used for the weight values encoding could lead to investigations of architectural variants of the ESN model, e.g. see [15].

Table 2. Mean accuracy (in %) and standard deviation of LI-ESNs with *reduced weight encoding* for $N_{anchors} = 4$

	Reservoir Dimension (N_R)					
	10	20	50	100	300	500
Training	76.8(± 2.4)	85.5(± 3.9)	96.6(± 1.4)	100.0(± 0.0)	100.0(± 0.0)	100.0(± 0.0)
Test	73.6(± 4.3)	84.8(± 3.9)	93.8(± 3.4)	94.3(± 2.7)	96.4(± 2.7)	97.6(± 1.8)

Concerning the computational power constraints, we observed that on motes with an 8 MHz processor, like the IRIS motes used in our experiments, the LI-ESN computations required for each input-output step approximatively 0.7 and 2 milliseconds, respectively for reservoirs of dimension 50 and 100, which is therefore plausible for a real-time processing on-board the nodes of a WSN.

5 Conclusions

We have presented an experimental investigation of the relation between the performance and the implementation cost of RC networks in an AAL task consisting in anticipating user movements in indoor environments. Results have shown that the proposed RC system achieves a very good predictive performance on the considered task. Such performance scales well with the number of anchors used and with the cost of the network memorization. In particular, a decrease in the number of reservoir units does not lead to a dramatic degeneration of the performance (up to 50 units), and the use of a reduced encoding scheme (requiring only 3 bits of memory for each weight value) does not affect significantly the accuracy of the model. A small reservoir with each weight value encoded in a few bits is sufficient in practice for the analyzed task. Such promising results allow us to envisage practical solutions for the embedding of the RC networks on-board the motes of WSNs.

The experimental analysis presented in this paper represents a ground for further investigations on the development of local and distributed learning capabilities based on RC systems distributed across the nodes of the RUBICON ecology.

Acknowledgement. This work is partially supported by the EU FP7 RUBICON project (contract n. 269914).

References

1. Amato, G., Broxvall, M., Chessa, S., Dragone, M., Gennaro, C., Lopez, R., Maguire, L., Mcginnity, M., Micheli, A., Renteria, A., O'Hare, G., Pecora, F.: Robotic UBIquitous COgnitive Network. In: Proceedings of the 3rd International Symposium on Ambient Intelligence (ISAMI), Salamanca, March 28-30 (To Appear, 2012)

2. Antonelo, E.A., Schrauwen, B., Stroobandt, D.: Event detection and localization for small mobile robots using reservoir computing. *Neural Networks* 21(6), 862–871 (2008)
3. Bacciu, D., Broxvall, M., Coleman, S., Dragone, M., Gallicchio, C., Gennaro, C., Guzman, R., Lopez, R., Lozano-Peiteado, H., Ray, A., Renteria, A., Saffiotti, A., Vairo, C.: Self-Sustaining Learning for Robotic Ecologies. In: *Proceedings of the 1st International Conference on Sensor Networks (SensorNets)*, Rome, February 24–26 (To Appear, 2012)
4. Bacciu, D., Gallicchio, C., Micheli, A., Chessa, S., Barsocchi, P.: Predicting user movements in heterogeneous indoor environments by reservoir computing. In: Bhatt, M., Guesgen, H.W., Augusto, J.C. (eds.) *Proceedings of the IJCAI Workshop on Space, Time and Ambient Intelligence STAMI 2011*, pp. 1–6 (2011)
5. Baronti, P., Pillai, P., Chook, V.W.C., Chessa, S., Gotta, A., Hu, Y.F.: Wireless sensor networks: A survey on the state of the art and the 802.15.4 and ZigBee standards. *Computer Communications* 30(7), 1655–1695 (2007)
6. Chang, M., Terzis, A., Bonnet, P.: Mote-Based Online Anomaly Detection Using Echo State Networks. In: Krishnamachari, B., Suri, S., Heinzelman, W., Mitra, U. (eds.) *DCOSS 2009. LNCS*, vol. 5516, pp. 72–86. Springer, Heidelberg (2009)
7. Crossbow Technology Inc.: <http://www.xbow.com>
8. Ducatel, K., Bogdanowicz, M., Scapolo, F., Leijten, L., Burgelman, J.: Scenarios for Ambient Intelligence in 2010. Tech. rep., IST Advisory Group (February 2001)
9. Gallicchio, C., Micheli, A.: Architectural and markovian factors of echo state networks. *Neural Networks* 24(5), 440–456 (2011)
10. Gallicchio, C., Micheli, A., Barsocchi, P., Chessa, S.: User movements forecasting by reservoir computing using signal streams produced by mote-class sensors. In: Del Ser, J., Jorswieck, E.A., Miguez, J., Matinmikko, M., Palomar, D.P., Salcedo-Sanz, S., Gil-Lopez, S. (eds.) *Mobilight 2011. LNICST*, vol. 81, pp. 151–168. Springer, Heidelberg (2012)
11. Jaeger, H.: The “echo state” approach to analysing and training recurrent neural networks. Tech. rep., GMD - German National Research Institute for Computer Science (2001)
12. Jaeger, H., Haas, H.: Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science* 304(5667), 78–80 (2004)
13. Jaeger, H., Lukosevicius, M., Popovici, D., Siewert, U.: Optimization and applications of echo state networks with leaky-integrator neurons. *Neural Networks* 20(3), 335–352 (2007)
14. Lukosevicius, M., Jaeger, H.: Reservoir computing approaches to recurrent neural network training. *Computer Science Review* 3(3), 127–149 (2009)
15. Rodan, A., Tiño, P.: Minimum complexity echo state network. *IEEE Transactions on Neural Networks* 22(1), 131–144 (2011)
16. Verstraeten, D., Schrauwen, B., D’Haene, M., Stroobandt, D.: An experimental unification of reservoir computing methods. *Neural Networks* 20(3), 391–403 (2007)

Balancing Recall and Precision in Stock Market Predictors Using Support Vector Machines

Marco Lippi, Lorenzo Menconi, and Marco Gori

Dipartimento di Ingegneria dell'Informazione, Università degli Studi di Siena
{lippi,menconi,marco}@dii.unisi.it

Abstract. Computational finance is one of the fields where machine learning and data mining have found in recent years a large application. Nevertheless, there are still many open issues regarding the predictability of the stock market, and the possibility to build an automatic intelligent trader able to make forecasts on stock prices, and to develop a profitable trading strategy. In this paper, we propose an automatic trading strategy based on support vector machines, which employs recall-precision curves in order to allow a buying action for the trader only when the confidence of the prediction is high. We present an extensive experimental evaluation which compares our trader with several classic competitors.

1 Introduction

How far is the time when most of the trading volume in financial markets will come from intelligent and automated trading strategies rather than by human interaction? And how is actually complex to build a financial strategy that can operate autonomously in the market? Research now provides several tools to approach the problem. In particular, in this paper, we use the basic concepts underlying computational finance and algorithmic trading, with the purpose of predicting the behavior of the market and of developing an intelligent trading strategy, really capable of directly operating on the market, and completely independent of human intervention.

In recent years, there has been a tremendous growth in the field of computational finance [9], with the rise of a wide variety of different automatic intelligent strategies, based on data mining and machine learning techniques. This spread of interest has been mainly due to the growing availability of data coming from the World Wide Web: nowadays, huge data flows have become accessible every day from all the stock markets around the world, and computational intelligence techniques have been increasingly employed to process and analyze stock prices and trading patterns.

One of the main goals of computational finance is to predict the evolution of the stock market, given observations of its past behavior. An example is given by forecasting the trend of a stock price: given a price time series $X_t = \{x_1, \dots, x_t\}$ the aim is to predict some property of the series at $t + \Delta$, and use such prediction to build an automatic trading strategy. When Δ is within the order of minutes or hours, we talk of intra-day predictions, which will be the main object of investigation of this paper. Short-term prediction methods are the ones which in the last years have received most of the

attention [7], and many different algorithms have been proposed, using several machine learning or statistical techniques, including artificial neural networks, support vector machines, volatility models and many others. Typical applications of such predictors include regression (e.g., predict the future price of a stock) or classification tasks (e.g., predict the trend of a stock in the forthcoming minutes/hours). Nevertheless, most of the literature lacks an accurate experimental evaluation of the proposed algorithms: sometimes only the accuracy of the predictor is measured, but no concrete analysis of the trading strategy is performed, too often commissions are not taken into account, and finally only few works directly compare *intelligent* traders with trivial strategies such as *Buy and Hold*. For these reasons, it is still an open question to infer at what extent is the stock market predictable: such question boils down to facing what is called the *random walk hypothesis* for the stock market, which states that no better estimate than the current price can be given for the future price of a given stock [4].

In this paper, we present an automatic intelligent trading system which uses support vector machines to identify the stocks to buy, while employing a sophisticated technique based on recall-precision curves in order to filter predictions, maintaining a higher level of confidence. In the following sections of the paper, we first tackle some questions regarding the general implementation of an automatic trading agent, then we describe our model based on support vector machines, and finally give an extensive experimental evaluation of our trader, in comparison with other classic approaches.

2 Limitations of Automatic Intelligent Traders

Differently from many other contexts, in financial time series forecasting some characteristics of the domain should be taken into account when building a machine learning predictor and therefore also an intelligent trader. In this work, we use support vector machines (SVMs) to predict whether a certain stock will rise more than a certain amount within a short horizon (e.g., a few hours). The observed past time series is used to build the input features for the classifier. In our preliminary experiments we aimed at answering several questions, which are fundamental in order to create an accurate model.

1. Do larger training sets always improve SVM performance ?
2. Do performance degrade over days/weeks/months of prediction, if the predictive model is kept fixed ?
3. At what extent is the stock market really predictable ?

2.1 Training Set Dimension

Machine learning algorithms typically take advantage of large training sets to reduce their approximation error. Yet, in the case of time series forecasting, if the training data is too distant in the past from the current prediction time, the effect of a larger training set does not necessarily reflect in an improvement in the performance of the predictor. Figure 1(a) shows the precision of an SVM classifier – which employs the past time series as input features – as a function of the training set dimension, keeping the test set




Fig. 1. Left: learning curve as a function of training set dimension. Right: learning curve as a function of test set distribution drift.

fixed. In this case, the binary classification task consists in predicting (positive class) whether the stock AIG (American International Group, NYSE) will gain at least 0.5% in the forthcoming 60 minutes: the test set consists of one whole week of data, with samples taken every 30 seconds, and the predictor is asked to make a prediction for each new sample. The time series made of the previous 60 samples is used as feature vector. It is clear from the graph that, after a certain period during which the precision of the predictor increases while growing the training set, then the performance starts degrading. It should also be noticed that we report the precision of the predictor (false positives ratio), which is always under 50 %: as it will be discussed in the following subsections, this somehow confirms the great difficulty of the task. The key idea of our SVM-based trader will be that of filtering the predictions, with the goal of maintaining a higher level of precision.

2.2 Test Set Horizon

Similarly to the previous case, in this second experiment we want to observe how the performance of a predictor behaves over a long test period, this time maintaining a fixed training set: this should give the idea of how often predictors should be updated in order to keep always a fresh and ready-to-use model. Figure 1 confirms that the performance degrades very quickly, suggesting that the models should be re-trained as often as possible. This phenomenon is well known in statistics and machine learning, and is called *distribution drift* or *covariate shift* [8]. When dealing with data coming from a distribution which changes over time, the performance of a predictor strongly depends on the temporal distance between training and test examples: for this reason, periodically re-training the models is a key strategy in order to get an accurate and updated forecasting system.

2.3 Stock Market Predictability

One of the main criticisms which are moved with regards to computational finance is that the stock market follows the random walk theory (RWT) and is actually not predictable [4]. The RWT comes as an interpretation of the efficient market hypothesis,

and states that in an efficient market, on average, the information is instantly absorbed by the prices, and consequently the time-series of prices has no memory. According to our experiments, this is partially true. In fact, if an SVM is trained to predict the future price of a stock (either with a regression or a classification task) at any time, such predictor will not behave better than a random walk classifier, on average. On the other hand, this does not exclude the possibility that there might be rare predictable features and arbitrage opportunities, which could be exploited by a human or artificial trader. In particular, an artificial trader has the ability of processing continuously huge amounts of data, analyzing cross-correlations between different stocks, observing micro-trends across thousands of stocks in different markets, and can therefore take into consideration many trading strategies in parallel, which is unfeasible for humans. This is the fundamental reason for which computational finance has recently conquered a large slice of the trading market [15].

In the following section we will present a machine learning predictor based on support vector machines, which will drive the trading strategy of an automatic agent, with the goal of making few transactions, but with high level of precision.

3 Trading Strategy

Following the results obtained in our preliminary experiments, we designed an SVM-based trader which merges two fundamental ideas: (1) the models used at prediction time should be updated at least daily, because predictions quickly degrade if keeping the models fixed; (2) the trader should buy only if the predictor is somehow *confident*, which means that some parameters have to be designed in order to score and guide the decisions.

A different SVM is trained for each different stock in our basket. A binary classification task is structured as follows: a feature vector describes the past time series of the prices of a given stock at time t , and the goal is to predict whether the price will have a certain gain in the forthcoming future. More precisely, a positive example is given if the stock will have a percentage gain greater than K within Δ minutes, being K and Δ two tunable parameters. By training an SVM in these conditions, it is possible to build a recall-precision curve on a validation set. Recall (R) — sometimes called True Positive Rate or Sensitivity — is defined as the ratio between correctly identified positive examples and total number of positive examples: $R = \frac{TP}{TP+FN}$, while precision (P) — sometimes called Positive Predictive Value — is the ratio between correctly identified positive examples and total number of examples predicted as positives: $P = \frac{TP}{TP+FP}$. By using different SVM margins as the values for discriminating between positive and negative classes, we obtain different values for recall and precision, which can be plotted to draw a curve: we select a working point on this curve, such that precision is above a certain threshold (chosen as $\hat{P} = 90\%$) and the $F_{0.5}$ value is as maximum as possible¹.

¹ In general, F_β combines precision and recall measures in order to obtain a single comprehensive score function: $F_\beta = (1 + \beta^2) \frac{P \cdot R}{\beta^2 \cdot P + R}$. Typically, F_1 is employed, which corresponds to the harmonic mean between recall and precision. The choice of $\beta = 0.5$ gives a higher weight to the precision.

The idea is to use the SVM margin corresponding to this working point on the RP curve as a triggering threshold τ_S for the positive class at prediction time: more precisely, if at time t the classifier outputs a margin $M > \tau_S$ for stock S , then the trader will buy S .

The feature vector is built as follows. Given the time series $X_t = \{x_{t-W \cdot A}, \dots, x_t\}$ containing the last $W \cdot A$ samples, we defined the vector of W features f_1, \dots, f_W , where feature f_j is computed as:

$$f_j = \frac{\hat{x}_j - \hat{x}_{j-1}}{\hat{x}_{j-1}} \quad (1)$$

where \hat{x}_j are aggregations over A samples:

$$\hat{x}_j = \frac{x_{t-j \cdot A} + x_{t-j \cdot A + 1} + \dots + x_{t-j \cdot A + A}}{A} \quad (2)$$

Variations expressed in percentage are preferable features than plain prices, as they are independent from the absolute value of the stock. In our preliminary experiments for the design of the classifier, we also tried logarithmic returns as input features, but they headed to slightly worse performance.

Several parameters have to be tuned for each stock S during model selection: the percentage gain predicted as positive class K (in our experiments, chosen as $K = 0.1\%$ for all stocks), the training set dimension T , the prediction horizon Δ , the aggregation time for input features A , the input window W , the SVM parameter C which controls generalization. As a selling criterion, we employed a stoploss s , one of the most common tools used by investors to exit the market and limit losses. The idea is simple: when the price of a purchased stock drops below a predefined threshold, the stock is sold. This stoploss threshold can be easily made dynamic, if computed as a function of the highest price reached by the stock after its purchase:

$$\text{lastHighestPrice}(t) = \max(\text{lastHighestPrice}(t-1), \text{price}(t)) \quad (3)$$

$$\text{stoplossThreshold}(t) = \text{lastHighestPrice}(t) \cdot (1-s) \quad (4)$$

In our experiments, we use a stoploss value $s = 0.005$.

Finally, the trading strategy was designed so as to close all trading days without any shares in the portfolio: 15 minutes before the market closes, all the remaining stocks in the portfolio are automatically sold².

4 Experimental Evaluation

We evaluate our SVM-based trading strategy on several test periods between 2010 and 2011. The data set consists in the stock price series of 86 assets of NYSE, from January 1st, 2010, to December 31st, 2011, with prices taken every 30 seconds. We use Interactive Brokers³ as trading platform. The experiments were conducted over 12 different test periods of 6 months, starting from July 1st, 2010 up to November 31st, 2011, shifting the semester by one month for each test period (e.g., the first period runs from July 1st, 2010 to January 1st, 2011). In this way, we will observe the behavior of the traders in different market scenarios, which include both bull and bear periods.

² This typically corresponds to propitious trading operations, because as a matter of fact the stoploss threshold had not yet ordered a sell operation.

³ <http://www.interactivebrokers.com>

4.1 Competitors

In order to have an extensive experimental evaluation, we compared the performance achieved by our SVM-based trader with several algorithms implementing different baseline trading strategies.

Random Trader (RT). The comparison with a random trader allows to measure how much an intelligent strategy is better than a purely random and uninformed one. Several different implementations of a random trader can be employed: in our version, every day the trader randomly chooses a buy and a sell moment, in the range from 30 minutes after the opening to 30 minutes before closing, always performing the buying action first.

Random Trader with Stoploss (RST). This is a variant of RT, which uses a stoploss $s = 0.005$ to decide when to sell.

Daily Open-Close Trader (DOCT). By using this strategy, every day a trader places a buy order 30 minutes after the opening and a sell order 30 minutes before closing. This is a simple strategy, that gives a measure of how much the market has grown in a day.

Buy and Hold Trader (BHT). The Buy and Hold trading strategy offers a point of view which is similar to that of the daily open-close trader, but on a larger time horizon. The BHT buys at the beginning of the test period, holding the purchased shares until the last day, when he sells just before the closing.

Technical Analysis Trader (TAT). A smarter trading strategy uses the combination of four classic technical analysis tools [2]: the Bollinger's bands (BBs), the Money flow index (MFI), the On balance volume index (OBV), and the Stoploss policy (SL). The TAT buys a certain stock when the majority of its three indicators votes for a purchase. The stoploss indicator is used to decide when to sell.

4.2 Results

To evaluate the performance of each trading strategy, we consider the gain obtained in each semester, including commissions fees as applied by Interactive Brokers. The orders are assumed to be executed with no time lag. At each period a portfolio of 10,000 USD is initialized.

In Table I we report for each test period the returns of each trade strategy, averaged over all the examined stocks: the RT, RST, DOCT, BHT, TAT strategies were in fact independently run for each stock, and their performance averaged over the whole set of stocks, while our trader can decide which stock to buy. As expected, the BHT, which keeps the position open for a whole semester, has good returns in bull periods, and bad returns in bear periods: obviously, this strategy is also influenced by inter-day price gaps. Like all the trading strategies which maintain an amount of stocks for several months, its behavior is almost independent of commission fees, and has the advantages and disadvantages of all long investments. Anyhow, it should be remarked that

Table 1. Trading strategies comparison for 12 different test semesters between 2010 and 2011. We report the percentage gain of each trader, starting at each period with a portfolio of 10,000 USD. The SVMT has no standard deviation, as it decides on which stocks it operates every day, while for all the other competitors results are averaged on all the considered stocks: the very large standard deviations indicate that all the other traders have very heterogeneous behaviors on the different stocks.

Semester	DOCT	BHT	RT	RST	TAT	SVMT
07/10 – 12/10	3.57 ± 13.37	29.94 ± 28.02	-4.34 ± 9.23	0.33 ± 9.42	1.52 ± 7.62	4.20
08/10 – 01/11	1.91 ± 13.80	20.40 ± 22.98	-3.26 ± 10.26	-2.57 ± 7.56	-0.12 ± 7.67	9.02
09/10 – 02/11	-0.83 ± 10.68	22.88 ± 21.94	-4.77 ± 8.75	-2.77 ± 8.27	-0.47 ± 7.90	6.19
10/10 – 03/11	-1.53 ± 10.82	13.61 ± 18.53	-5.88 ± 7.99	-4.14 ± 7.63	-0.11 ± 7.32	-2.30
11/10 – 04/11	-2.92 ± 10.49	12.06 ± 18.59	-4.65 ± 8.28	-4.17 ± 7.44	-0.15 ± 6.01	1.66
12/10 – 05/11	-5.27 ± 12.69	6.36 ± 19.69	-5.72 ± 7.95	-2.65 ± 7.43	-1.24 ± 6.55	0.03
01/11 – 06/11	-7.07 ± 13.83	-1.11 ± 20.48	-5.94 ± 8.37	-3.70 ± 6.60	-2.32 ± 6.66	-8.92
02/11 – 07/11	-8.16 ± 13.01	-5.30 ± 15.55	-6.60 ± 8.28	-4.96 ± 5.92	-3.13 ± 6.61	-4.47
03/11 – 08/11	-14.29 ± 14.12	-13.65 ± 17.70	-8.32 ± 10.35	-3.74 ± 6.56	-4.48 ± 8.01	-12.26
04/11 – 09/11	-18.59 ± 17.39	-23.46 ± 19.09	-8.61 ± 10.64	-6.21 ± 6.91	-6.59 ± 8.94	-9.26
05/11 – 10/11	-12.63 ± 19.29	-18.71 ± 16.70	-7.65 ± 12.21	-7.83 ± 6.69	-5.97 ± 8.79	0.23
06/11 – 11/11	-9.10 ± 19.49	-15.20 ± 18.19	-6.72 ± 12.67	-8.94 ± 7.39	-4.89 ± 9.48	-3.33

the performance of BHT are extremely dependent on the stock which is chosen at the beginning of the strategy: this is confirmed by the very large standard deviations (results are averaged on the set of stocks) in the first column of Table 1, which indicate that such a strategy can bring very large gains as well as very large losses.

In the early (bull) semesters, there are considerable inter-day positive variations in price, while in the last few (bear) semesters the price changes are mostly intraday: for this reason, the DOCT earns much less than the BHT during bull periods, while having similar behavior during bear periods. The TAT, which uses BBs, MFI and OBV, performs generally better than a random trader and the DOCT, but is a very conservative trader and does not seem to offer real gain opportunities.

The SVMT, on the other hand, presents encouraging results, especially containing the losses during many bear periods. Further improvements can be applied to this trader, including a smarter analysis of model parameters, with the aim of buying stocks with high share values, so that commission fees have a lower impact on the trading performance. Similarly, also information about volumes and volatility should be included within the model. It should be remarked that, in our preliminary experiments, a plain SVM trader which does not use the recall-precision method for threshold selection headed to very poor results, with large losses in all semesters.

5 Future Works

The goal of this work was to compare our SVM-based trading technique with common baselines and technical analysis strategies, which are typically employed in computational finance. Future directions of research will include a more complete experimental analysis, in order to compare the proposed approach also with other machine learning algorithms, such as neural networks, decision trees and classical time-series forecasting algorithms like SARIMA [10] or GARCH [3]. In order to improve the performance of

our trading strategy, we are currently investigating techniques based on statistical relational learning, which might be used to model interdependencies between different time series: this kind of approach has been recently applied to other forecasting domains, like traffic forecasting [6]. Finally, the use of textual information obtained by several sources across the World Wide Web might greatly improve the opportunities of the automatic trader, for example to predict opening prices.

6 Conclusions

Computational finance has become nowadays a constantly evolving research area. As the amount of data becoming available with the world wide web has enormously grown in the last years, huge money streams are daily moved by computers all around the world. Machine learning and data mining have therefore found in finance a challenging and profitable domain.

In this paper we presented an automatic trading strategy based on support vector machines, which uses recall-precision curves in order to suggest a buying action only when the confidence level of the trader is high. We compared our trader on several test periods between 2010 and 2011, showing strongly encouraging results with respect to classic baseline traders, even based on technical analysis.

Many future directions of research can be developed: relational learning techniques can be used to capture relations among stocks; volumes and volatility information can be added to our model; text mining techniques can be applied in order to acquire knowledge from news, blogs and forums around the web, to guide the choices of the trader.

References

1. Chaboud, A., Chiquoine, B., Hjalmarsson, E., Vega, C.: Rise of the machines: algorithmic trading in the foreign exchange market. International Finance Discussion Papers 980, Board of Governors of the Federal Reserve System, U.S. (2009)
2. Edwards, R.D., Magee, J., Bassetti, W.H.C.: Technical analysis of stock trends. John Magee Investment Series. St. Lucie Press (2001)
3. Engle, R.: GARCH 101: The Use of ARCH/GARCH Models in Applied Econometrics. The Journal of Economic Perspectives 15(4), 157–168 (2001)
4. Fama, E.F.: Random Walks in Stock Market Prices. Financial Analysts Journal 21, 55–60 (1965)
5. Kirilenko, A.A., Kyle, A.P., Samadi, M., Tuzun, T.: The flash crash: The impact of high frequency trading on an electronic market. Social Science Research Network Working Paper Series (October 2010)
6. Lippi, M., Bertini, M., Frasconi, P.: Collective traffic forecasting. In: ECML/PKDD 2010, pp. 259–273 (2010)
7. Menconi, L., Gori, M., Lippi, M.: Computational models for short-term prediction of the stock market. Intelligenza Artificiale 5(2), 217–227 (2011)
8. Quionero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N.D.: Dataset Shift in Machine Learning. MIT Press (2009)
9. Seydel, R.U.: Tools for Computational Finance. Springer, Heidelberg (2009)
10. Williams, B.M., Hoel, L.A.: Modeling and forecasting vehicular traffic flow as a seasonal arima process: Theoretical basis and empirical results. Journal of Transportation Engineering 129(6), 664–672 (2003)

Measures of Brain Connectivity through Permutation Entropy in Epileptic Disorders

Domenico Labate*, Giuseppina Inuso, Gianluigi Occhiuto,
Fabio La Foresta, and Francesco C. Morabito

MECMAT - Mediterranea University of Reggio Calabria,
via Graziella Feo di Vito, 89122 Reggio Calabria, Italy
{domenico.labate,giuseppina.inuso,gianluigi.occhiuto,fabio.laforesta,
morabito}@unirc.it

Abstract. Most of the scientist assume that epileptic seizures are triggered by an abnormal electrical activity of groups of neural populations that yields to dynamic changes in the properties of Electroencephalography (EEG) signals. To understand the pathogenesis of the epileptic seizures, it is useful detect them by using a tool able to identify the dynamic changes in EEG recordings. In the last years, many measures in the complex network theory have been developed. The aim of this paper is the use of Permutation Entropy (PE) with the addition of a threshold method to create links between the different electrodes placed over the scalp, in order to simulate the network phenomena that occur in the brain. This technique was tested over two EEG recordings: a healthy subject and an epileptic subject affected by absence seizures.

Keywords: Permutation Entropy, EEG, Complex Network.

1 Introduction

EEG is a measure of neurophysiological brain electrical activity recorded through electrodes placed over the scalp. Neurons, in fact, communicate by means of electrical impulses and generate a bio-electromagnetic field that propagates through the tissues of the brain, skull and scalp [1]. EEG signal is produced by the extracellular current flow generated by the sum of the activities of a large number of neurons, including cortical pyramidal neurons, arranged in correspondence of the cortex underlying the electrode [2][3].

Epilepsy is one of the most common neurological disorders. Epileptic seizures are triggered by neurons that exhibit an abnormal electrical activity [4], as if a critical area is trying to recruit other cerebral areas until the brain can't longer bear this hyper-synchronization and triggers the seizure in order to repress this abnormal state. Then, epileptic phenomena seem a network rather than local phenomena [5].

* Corresponding author.

In order to understand the development of epileptic seizures, it is important to understand how this abnormal synchronization influences the randomness of the EEG. The possible solution to this issue can be entropy. In particular, Permutation Entropy is chosen [6][7][8] because, being a symbolic procedure, pre-normalization step can be discarded and because it was applied to epileptic EEG analysis obtaining encouraging results.

In this paper, a study of links between different cerebral areas for two types of subjects is presented: an healthy subject (called "control") and an epileptic subject (called "patient"). The aim of this paper is to use PE in order to underline different network phenomena in epileptic subject during EEG recordings. The paper is organized as follows: section 2 will address Permutation Entropy; section 3 will explain the applied method used to analyze the EEG recordings; section 4 will comment the obtained results and, ultimately, section 5 will draw the conclusions.

2 Permutation Entropy

PE has been introduced [9] as a robust method to extract information from a time series. The time series is analyzed from a pure ordinal point of view. PE is based by counting of the ordinal patterns (called "motifs"), that describe the up-and-down in a dynamical signal. Since just ordinal patterns are considered, the amplitude of the signal has no relevance, thus yielding a structural robustness to noise. In the study of biological signals, like EEG, this also implies the independence of the choice of the reference electrode [10]. As a consequence, there is no need of any normalization in the pre-processing step.

PE combines the concept of Shannon Entropy to the ordinal pattern analysis by analyzing the frequencies of ordinal patterns within a time series. If most of the ordinal patterns are the same, we have a regular time series; on the contrary, the presence of different patterns with similar frequency is indicative of high complexity.

PE is dependent on two-parameters, an embedding dimension, m , and a time-lag, l ; m is the number of samples belonging to the segment, while l represents the distance between the samples points spanned by each section of the motif. The variation of the time-lag basically implies working on different scales of the time series. In this work, the embedding parameter and time-lag are fixed to 3 and 1 respectively.

The calculation of PE is simple and an example is depicted in Figure 1. First of all, the EEG signal is fragmented into a sequence of motifs; then, each motif is identified as belonging to one of the possible types according to their shape; it is counted the number of motifs to each of the six categories in order to obtain the probability of occurrence of each motif in the signal, and, finally, is calculated the PE of the resulting normalized probability distribution of the motifs using the Shannon uncertainty formula:

$$PE = -\frac{\sum p_i * \ln(p_i)}{\ln(\text{number of motifs})} \quad (1)$$




Fig. 1. (a) EEG recording. (b) PE time evolution. (c) Motifs distribution ($m=3$) of the two highlighted segments. (d) The six possible motifs.

The normalized entropy is maximal ($PE=1$) if there is an equal distribution of motifs between each of the six patterns; on the contrary, when the signal is dominated by slow waves, there will be a greater occurrence of a few motifs, and the entropy decreases. The PE of a single motif is zero [6].

PE shows a different dependence with frequency with respect to the selected time-lag. To understand how, a sequence of pure sine waves time series is simulated and then the PE for each of them is performed. The results are reported in Figure 2. It is clear that, with $l=1$, the PE is monotonically decreasing with decreasing frequency. The frequency dependence of PE with $l=2$ is markedly different and more difficult to interpret. In Figure 2, it is possible to observe that the PE function has a minimum for a frequency that is a fraction of the sampling frequency, F_s . Some authors decided to define a suitably averaged composite PE index that incorporates both the cases, this way exploiting the effect [6].

The effects of noise on PE have been originally discussed in [9]. If a time series is perturbed by noise, the PE value strongly increases. Thus, if the added white noise is of sufficient power, the value of PE tends to its maximum at very low frequencies. Figure 3 shows the effect of adding white noise to the simulated sine-wave time series. It is clear that, the added noise leaves undisturbed the PE curve after a frequency threshold. This threshold level depends on the power of the noise [11].

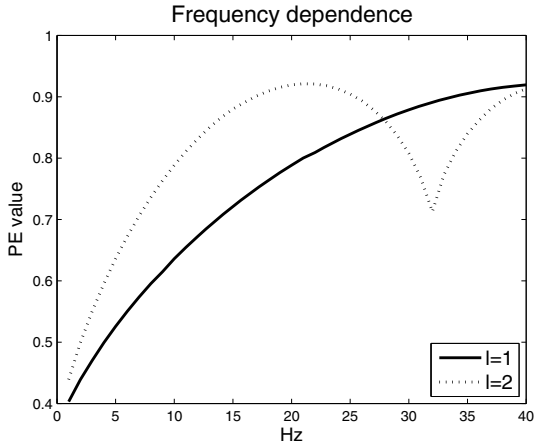


Fig. 2. Frequency dependence of PE. The two distributions for $l=1$ and $l=2$, $m=3$, and $F_s=128$ Hz are reported.

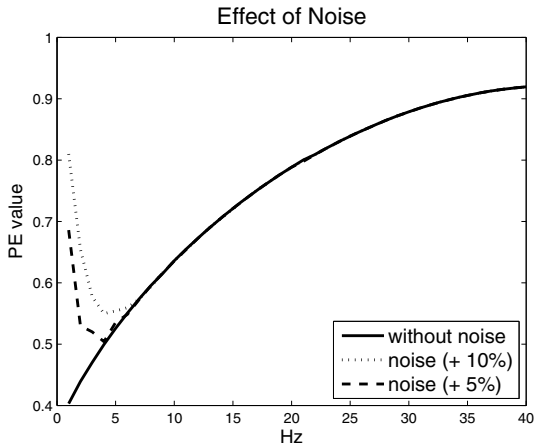


Fig. 3. The effect of adding white noise on PE calculation ($m=3$, $l=1$ and $F_s=128$ Hz)

3 Method

The EEG recording is processed by a sliding overlapping windows of 10 seconds with 1 second of overlap, i.e. each window includes 2560 time samples. The PE is performed, channel by channel, for each temporal window using $m=3$ and $l=1$. Once PE values are obtained, they are used in order to build a picture in which a scalp is depicted with all its 19 electrodes. If two electrodes have both a PE value greater than a fixed threshold, then the two electrodes are linked. This procedure is repeated for all the time windows obtained in the previous step. At the end of the analysis a movie with all the obtained pictures is built.

The purpose of the movie is to better clarify the link time evolution between the electrodes during the different stages of the recording. To better compare the obtained results, the two EEG recordings are cut in order to have the same size. This method is implemented in Matlab code.

4 Results

4.1 Data Description

The analyzed dataset is composed of two EEGs of nineteen channels recorded simultaneously, the sampling rate is 256 Hz. During the acquisition, Notch filter was set to ON and the EEGs was band-pass filtered (0.5 - 70 Hz). The analyzed EEGs are relative to a healthy subject (control) and a patient affected by absence epilepsy (patient). During EEG recording, the patient has undergone four seizures.

4.2 Discussion

The EEG recordings was analyzed as described in section 3. PE was performed for each channel and window by window.

In Figure 4 are reported PE values for the two EEG recordings. In order to consider all the 19 channels, the average of PE values is reported. It is clear as for the patient there are four peaks down due to the four absence seizures. Looking at a typical PE temporal trend, it can probably hypothesized that the sudden drop during the ictal stage is due to a mechanism of the brain to reset the high randomness that develops during the inter-ictal stage [5][8]. For the control, instead, the PE values are mostly the same for the entire duration of the EEG

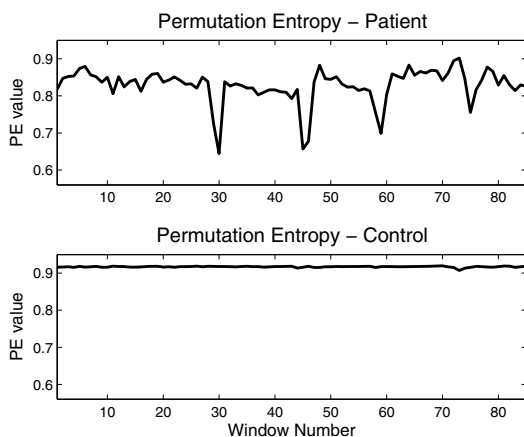


Fig. 4. (on the top) PE value for the patient; (on the bottom) PE value for the control

Comparison between Patient and Control

Threshold = 0.85, Window Number: 27

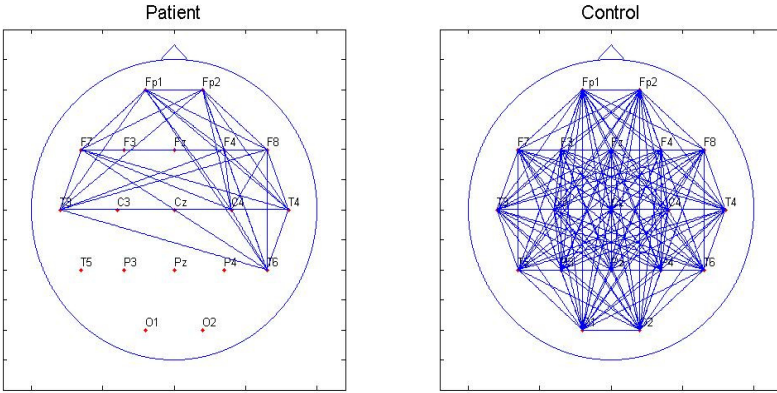


Fig. 5. The comparison of electrodes linked between the patient and the control for the window number 27 (pre-ictal stage)

Comparison between Patient and Control

Threshold = 0.85, Window Number: 30

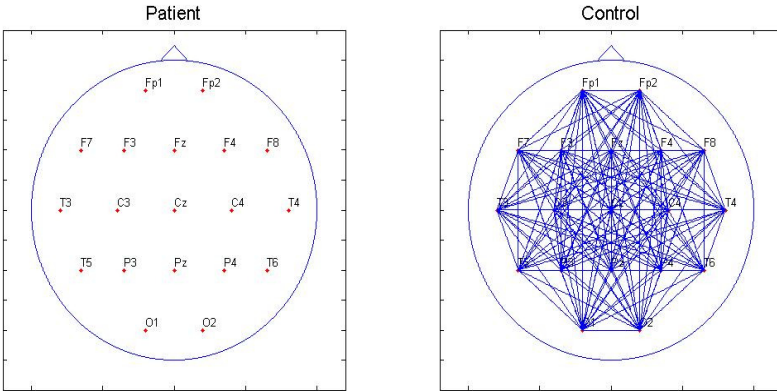


Fig. 6. The comparison of electrodes linked between the patient and the control for the window number 30 (ictal stage)

recording. Then PE can be used as a useful tool to identify the dynamic changes in the EEG.

Figure 5 and Figure 6 show how PE values change during the EEG recording. In these pictures are reported the scalps for patient and control both. Figure 5 is relative to a pre-ictal stage for the patient; it is clear that some electrodes are linked. Figure 6, instead, is relative to a ictal stage for the patient. In this case no electrodes are linked due to fact that all the PE values are lower than the threshold value; this change is due to absence seizure occurring in Figure 6. Observing the scalp relative to the control, there are no changes between the two figure because in both the PE values are always greater than the threshold value. Moreover, analyzing Figure 4, Figure 5 and Figure 6, it is clear that the PE values relative to the control are greater than the PE values relative to the patient; this is due to a most important cerebral activity in the healthy subject.

Table 1 and Table 2 report PE values for each electrode and for several EEG areas for three different thresholds: 0.85, 0.82 and 0.79. This is due to the fact that there isn't a fixed threshold value able to distinguish a patient by a control or a pre-ictal stage by an ictal stage. However the behaviour is similar for all the thresholds reported.

Observing Table 1 and Table 2, it is clear that the electrodes belonging to the frontal/temporal area have PE values greater than the other electrodes throughout the entire EEG recording and not only during the ictal stage. On the

Table 1. In the first column there are the electrodes; in the second, third and fourth column, there are the number of times in which the Patient's PE value exceeds the threshold value set to 0.85, 0.82 and 0.79 respectively

Electrode	th=0.85	th=0.82	th=0.79
Fp1	76/85	77/85	79/85
Fp2	76/85	78/85	80/85
F7	68/85	79/85	80/85
F3	24/85	73/85	78/85
Fz	24/85	71/85	76/85
F4	50/85	74/85	77/85
F8	73/85	80/85	80/85
T3	80/85	81/85	84/85
C3	15/85	38/85	66/85
Cz	5/85	27/85	53/85
C4	5/85	53/85	73/85
T4	79/85	80/85	82/85
T5	12/85	39/85	64/85
P3	7/85	20/85	48/85
Pz	2/85	15/85	38/85
P4	12/85	29/85	53/85
T6	35/85	53/85	71/85
O1	2/85	13/85	37/85
O2	3/85	23/85	46/85

Table 2. In the first column there are the EEG areas; in the second, third and fourth column, there are the number of times in which the Patient's PE value exceeds the threshold value set to 0.85, 0.82 and 0.79 respectively

Area	th=0.85	th=0.82	th=0.79
Fp1-Fp2	76/85	77/85	79/85
F7-F3	24/85	73/85	78/85
F4-F8	50/85	74/85	77/85
T3-C3	15/85	38/85	66/85
T4-C4	22/85	53/85	73/85
P3-T5	7/85	20/85	48/85
P4-T6	12/85	29/85	53/85
O1-O2	2/85	13/85	36/85

contrary, there are some electrodes belonging to the parietal/occipital area, that have PE values smaller than the average of PE values. However, the analysis of a unique EEG recording affected by absence seizures is not enough in order to suggest a clinical reason for this behaviour.

PE values reported in Table 2 are meaningful because they allow detection of network phenomena relative to different EEG areas by means of a local measure like PE.

5 Conclusions

In this paper a EEG analysis based on PE was performed in order to study the evolution of links between the different electrodes placed over the scalp. PE was proposed to measure the randomness/synchrony of the brain. It has been shown the different features between an healthy subject and an epileptic subject, and how there are no links during an epileptic seizure. Observing PE values, it is clear how the frontal/temporal area has PE values greater than PE values relative to other cerebral areas. PE has been a good tool to identify dynamical changes in epileptic EEG.

References

1. Mammone, N., La Foresta, F., Morabito, F.C.: Automatic Artifact Rejection from Multichannel Scalp EEG by Wavelet ICA. *IEEE Sensors Journal* 12(3), 533–542 (2012)
2. Labate, D., La Foresta, F., Inuso, G., Morabito, F.C.: Remarks about Wavelet Analysis in the EEG Artifacts Detection. In: *Frontiers in Artificial Intelligence and Application: Neural Nets WIRN10*, vol. 226, pp. 99–106 (2011)
3. Labate, D., La Foresta, F., Inuso, G., Morabito, F.C.: Multiscale Entropy Analysis of Artifactual EEG Recordings. In: *Frontiers in Artificial Intelligence and Application: Neural Nets WIRN11*, vol. 234, pp. 170–177 (2011)

4. Mammone, N., Inuso, G., La Foresta, F., Versaci, M., Morabito, F.C.: Clustering of entropy topography in epileptic electroencephalography. *Neural Computing and Applications* 20, 1–9 (2011)
5. Mammone, N., La Foresta, F., Morabito, F.C.: Discovering Network Phenomena in the Epileptic Electroencephalography through Permutation Entropy Mapping. In: *Frontiers in Artificial Intelligence and Application: Neural Nets WIRN10*, vol. 226, pp. 260–269 (2011)
6. Olofsen, E., Sleigh, J.W., Dahan, A.: Permutation Entropy of the Electroencephalogram: a Measure of Anaesthetic Drug Effect. *British Journal of Anaesthesia* 101, 810–821 (2008)
7. Bruzzo, A.A., Gesierich, B., Santi, M., Tassinari, C.A., Birbaumer, N., Rubboli, G.: Permutation entropy to detect vigilance changes and preictal states from scalp EEG in epileptic patients. A preliminary study. *Neurological Science* 29, 3–9 (2008)
8. Li, X., Ouyang, G., Richards, D.A.: Predictability analysis of absence seizures with permutation entropy. *Epilepsy Research* 77, 70–74 (2007)
9. Bandt, C., Pompe, B.: Permutation entropy - a natural complexity measure for time series. *Physical Review Lett.* 88, 174102 (2002)
10. Mammone, N., Lay-Ekuakille, A., Morabito, F.C., Massaro, A., Casciaro, S., Trabacca, A.: Analysis of absence seizure EEG via Permutation Entropy spatio-temporal clustering. In: *Proceedings of IEEE International Symposium on Medical Measurements and Applications (MEMEA 2011)*, Bari, Italy, May 30-31, pp. 532–535 (2011)
11. Morabito, F.C., Labate, D., La Foresta, F., Bramanti, A., Morabito, G., Palamara, I.: Multivariate Multi-Scale Permutation Entropy for Complexity Analysis of Alzheimer's Disease EEG. *Entropy* 14(7), 1186–1202 (2012)

A New System for Automatic Recognition of Italian Sign Language

Marco Fagiani, Emanuele Principi, Stefano Squartini, and Francesco Piazza

Dipartimento di Ingegneria dell'Informazione
Università Politecnica delle Marche, Ancona, Italy
{m.fagiani,e.principi,s.squartini,f.piazza}@univpm.it

Abstract. This work proposes a preliminary study of an automatic recognition system for the Italian Sign Language (Lingua Italiana dei Segni - LIS). Several other attempts have been made in the literature, but they are typically oriented to international languages. The system is composed of a feature extraction stage, and a sign recognition stage. Each sign is represented by a single Hidden Markov Model, with parameters estimated through the resubstitution method. Then, starting from a set of features related to the position and the shape of head and hands, the Sequential Forward Selection technique has been applied to obtain feature vectors with the minimum dimension and the best recognition performance. Experiments have been performed using the cross-validation method on the Italian Sign Language Database A3LIS-147, maintaining the orthogonality between training and test sets. The obtained recognition accuracy averaged across all signers is 47.24%, which represents an encouraging result and demonstrates the effectiveness of the idea.

1 Introduction

Sign languages [15] fulfil the same social and mental functions as spoken languages, allowing deaf and hearing impaired people to communicate in a comfortable way. These languages arise spontaneously, evolve rapidly, and are so geographical specific to present differences also in regions where the oral language is unique.

Interaction and communication between deaf and hearing people entails the same misunderstanding problem found in oral languages, given that is widely accepted that sign languages are independent on oral ones. The number of problems increase taking into account that the interpreters, with their indispensable work, cannot be always present and that less of the 0.1% of total population belong to hearing impaired community. Therefore, “signers” are faced with a considerable difficulty and a social inclusion problem. That is why the development of sign language recognition and synthesis tools have been gaining an increasing interest among the scientific community.

In [1], algorithms have been developed to robustly extract features from hands and face, with experiments conducted in uncontrolled environments. A multiple

hypotheses tracking approach is used for the extraction of hands features, while the extraction of face ones is based on the active appearance model. The recognition stage is based on Hidden Markov Models (HMM) and is designed for the detection of both isolated and continuous signs. Theodorakis *et al.* [17] propose a so-called product HMMs for efficient multi-stream fusion, while Maebatake *et al.* [14] discussed the optimal number of the states and mixtures to obtain the best accuracy. SignSpeak [9] is one of the first European funded projects that tackles the problem of automatic recognition and translation of continuous sign language. Among the works of the project, Dreuw and colleagues in [8,7,6] introduced important innovations: they used a tracking adaptation method to obtain an hand tracking path with optimized tracking position, and they trained robust models through visual speaker alignments and virtual training samples. A first framework for continuous sign language recognition applied to Italian has been proposed in [12]. The system uses features based on centroid coordinates of the hands and a self-organizing map neural network to classify the single sign and to provide a list of probable meanings. Experiments are performed on a corpus of 160 videos with a vocabulary of 40 signs (each sign is repeated four times).

This paper proposes an automatic sign language recognition system based on HMM. The structure and the parameters of the sign models are investigated by means of the resubstitution method, and a study on the feature set is presented using the Sequential Forward Selection approach. In addition, the system uses only features obtained from the evaluation of position and contours of face and hands. Facial features related to mouth, lips, nose and eyes are not used. Being a preliminary study, the simpler task of isolated sign recognition is addressed, while continuous sign recognition is reserved for future works. Experiments have been carried out on a recently proposed Italian Sign Language Database (A3LIS-147) [10], and the evaluation has been performed using the cross-validation approach, this way guaranteeing the independence between training and test sets. This represents a notable difference from [12], where the same signers appear in both sets.

The outline of the paper is as follows: Section 2 briefly illustrates the used video corpus. Section 3 proposes an overview of the developed system. Section 4 presents and discusses recognition results. Finally, Section 5 draws conclusions and proposes future developments.

2 Database A3LIS-147

This section describes the principal characteristics of the video corpus that has been used for the evaluation of the developed system. The database has been presented in [10], and is composed of 147 distinct isolated signs executed by 10 different signers: 7 males and 3 females. Each video presents a single sign which is preceded and succeeded by the occurrence of the “silence” sign.

Video sequences have been acquired by using a commercial camera located in front of the subject. Behind the signer, a green chroma-key background is placed and two diffuse lights (400 W each) are used to ensure an uniform lighting. The




Fig. 1. Components of the sign language recognition system

video streaming has been acquired at 25 fps with a resolution of 720x576 pixel. The video sequences have been stored in Digital Video (DV) format in order to keep the maximum image quality.

3 System Overview

Given a sign language video input, the goal of the developed recognition system is to produce the corresponding sign as text. An overview of the whole system architecture is depicted in Fig. 1 in the first stage, the system processes the input video and extracts meaningful features. As a preliminary step, head and hands regions are located using a skin detection algorithm, then the actual feature matrix of size $N \times L$ is calculated. Here, N is the total feature vector dimension and L is equal to the video frames number.

The HMM decoder takes as input the feature matrix, calculates the output probabilities for each model, and selects the sign which scores best. The HMM models of each sign are not known and have to be estimated from the training data.

3.1 Feature Extraction

As aforementioned, the first step of the feature extraction stage is the discrimination of skin regions. In this work, skin recognition is based on the evaluation of each pixel colour in the YC_bC_r colour space, which produces better performance with respect to RGB since it is luminance independent. The corresponding skin cluster is given as [13]:

$$\begin{cases} Y > 80 \\ 85 < C_b < 135 \\ 135 < C_r < 180 \end{cases} \quad (1)$$

where $Y, C_b, C_r \in [0, 255]$.

The pixels satisfying the above conditions are included in the skin mask, with head and hands determined from the regions with largest areas (Fig. 2(b)). A morphological transformation [3,16], closing operation, has been introduced to merge the nearby regions on the skin mask that belong to the same zones of




Fig. 2. Frame from the sign “abbonamento” (a). Skin regions detected (b).

interest. The closing is obtained applying first the dilation operator, then the erosion one. Fig. 3(a) shows the contours of the selected skin regions.

Table 1 shows the entire feature set calculated from the interest regions. The set can be divided into two different subsets: the first comprises features 1–9 and is obtained starting from the relative position of head and hands obtained by the mass center of the detected regions. It includes head and hands positions, position derivatives (movements velocity) and the hands distance. The second subset comprises features 10–18 and represents the general characteristics of head and hands, such as the area and area derivative, the shape type and the 2-D spatial orientation of the hands.

In the first subset of features, two different ways to compute the spatial moments for the mass center are used: features 1–3 use the Canny algorithm [4], while features 4–9 directly exploit the contours segmented from the skin mask (Fig. 3(a)). The Canny algorithm is applied within the bounding rectangle of the hands and obtains their contours as illustrated in Fig. 3(c). For the features 7–9 the head position is used as the center of the coordinate system. Features 10 and 11 are moments invariant to translation, changes in scale, and rotation [11]. Both for the head and hands regions, 7 different values are computed. As for the first features group, the Hu-moments are obtained in two different ways: using the Canny algorithm for 10 and using directly the contours for 11. The compactness 14 and eccentricity 15 of head and hands have been proposed in [17]. The first one is computed as the ratio of the minor and major axis lengths of the detected region. The second is obtained as the ratio of area and perimeter squared of the region. Features 16–18 have been introduced in [1], where $\mu_{p,q}$ indicates the corresponding central moments. The hands orientation $\alpha \in [-90^\circ, 90^\circ]$ is split into α_1 and α_2 to ensure stability at the interval borders. Note that all measurements are normalized respect the head dimension, so that they are independent of the distance from the camera and the signer height.




Fig. 3. Contours detection from the skin regions (a). Bounding rectangle of hands skin zones (b). The hands edge obtained by the Canny algorithm (c).

The overlap between head and hands, and between hands complicates the estimation of the mass centers of these regions. An overlap is detected when the skin mask contains two regions instead of three. When this occurs, the system evaluates whether the head area increased from the previous frame. An increase is considered as an head-hand overlap. Otherwise, if the area is unchanged, an hand-hand overlap is reported and the same region is assigned to both hands. With regard to the head, its position during the execution of a sign is considered as almost stationary. In order to limit incongruous deviations due to overlap with hands, a simple smoothing of the $(x, y)_t$ coordinates of the mass center at time t is computed:

$$(x, y)_t = \alpha(x, y)_t + (1 - \alpha)(x, y)_{t-1}, \quad (2)$$

where $\alpha \in [0, 1]$. With regard to the hands, the overlap with the head is managed assigning the position of the whole region at time t to the overlapped hand.

3.2 Sign Recognition

Hidden Markov Models are widely used in speech recognition applications, and they have been chosen for the sign recognition stage due to the similarity of the task. For each sign, a different left-to-right model with S states and G mixtures of Gaussians is created. The choice of the values of S and G is discussed in the following section. Recognition is performed calculating the likelihood for each sign model and selecting the sign whose model likelihood is highest.

Table 1. Feature set: A and P represent respectively the area of the region and its perimeter

Features	Parameters
1 Hands centroid center Canny-Filter normalized respect head position	$(x, y)_{right} (x, y)_{left}$
2 Hands centroid center derivatives Canny-Filter normalized respect head position	$(\dot{x}, \dot{y})_{right} (\dot{x}, \dot{y})_{left}$
3 Hands distance Canny-Filter normalized respect head dimension	d_{norm}
4 Hands centroid center contours (segmented by skin detection)	$(x, y)_{right} (x, y)_{left}$
5 Hands centroid center derivatives contours	$(\dot{x}, \dot{y})_{right} (\dot{x}, \dot{y})_{left}$
6 Hands distance contours	d_{norm}
7 Hands centroid center contours referring at head coordinates	$(x, y)_{right} (x, y)_{left}$
8 Hands centroid center derivatives contours referring at head coordinates	$(\dot{x}, \dot{y})_{right} (\dot{x}, \dot{y})_{left}$
9 Hands distance contours referring at head coordinates	d_{norm}
10 Head and hands Hu-moments Canny-Filter	$I_1 I_2 I_3 I_4 I_5 I_6 I_7$
11 Head and hands Hu-moments contours	$I_1 I_2 I_3 I_4 I_5 I_6 I_7$
12 Hands area normalized with head area	$A_{right} A_{left}$
13 Hands area derivative	$\dot{A}_{right} \dot{A}_{left}$
14 Head and hands compactness 1	$AXIS_{min}/AXIS_{max}$
15 Head and hands eccentricity 1	A_{norm}/P^2
16 Head and hands compactness 2	$(4\pi A_{norm})/P^2$
17 Head and hands eccentricity 2	$((\mu_{2,0} - \mu_{0,2})^2 + 4\mu_{1,1})/A_{norm}$
18 Hands orientation	$o_1 = \sin(2\alpha) o_2 = \cos(\alpha)$

4 Experimental Results

All the video elaboration have been executed with the OpenCV, an open source computer vision library [2]. The aforementioned A3LIS-147 database has been used for all the recognition test and training, and the sign recognition stage has been implemented using the Hidden Markov Model toolkit [19].

The first experiment studies the number of states and mixtures in left-to-right HMM representing each sign. The experiments has been performed evaluating the recognition results for different number of states (from 1 to 10) and mixtures (from 1 to 8). The second experiments analyses the performance of the system for different combinations of features.




Fig. 4. Resubstitution accuracy for different states and mixtures number

The performance have been evaluated using the sign recognition accuracy. It is defined on the basis of the number of substitution (S), deletion (D) and insertion errors (I) as in speech recognition tasks [19]:

$$\text{Accuracy} = \frac{(N - D - S - I)}{N}. \quad (3)$$

The term N is the total number of labels in the reference transcriptions. It is worth pointing out that in the isolated sign recognition task addressed in this work only the substitution errors are present.

4.1 Hidden Markov Models Parameters Estimation

In order to find an estimate of the number of states and of the number of mixtures per state of the sign models, the resubstitution method [18] has been used. The method entails the utilization of the same data set both for training and for testing with the aim of providing an optimistic estimation of the true error probability. The method has been applied recursively with different states and mixtures number using features 1, 3, 10, 12, 14 and 15.

Fig. 4 shows the obtained results: as expected, increasing the number of states and the number of mixtures per state produces better recognition accuracies. A *plateau* is reached when the number of states is above 7 and the number of mixtures is more than 5, suggesting that a good trade-off between computational cost and performance can be obtained using such values. Using 7 states and 5 mixtures gives a recognition accuracy of 96.69%.

4.2 Features Selection

The best feature combination has been obtained by means of the Sequential Forward Selection (SFS) [18] technique. SFS is sub-optimal, and consists of the following steps:

1. compute the criterion value for each of the features;
2. select the feature with the best value;
3. form all possible vectors that contain the winner from the previous step;
4. compute the criterion value for each of the new vectors;
5. select the best one;
6. repeat until the end of the features.

The results for each combination of features have been evaluated by means of the cross-validation method [18], a variation of the leave-one-out method, where $K > 1$ samples are used for testing and the remaining $N - K$ samples are used for training. The procedure is repeated for N iterations, each time excluding a different set of K samples. This way, training is performed using all samples and the independence between the training and test sets is maintained, differently from the approach followed in [12]. In the experiments, $N = 1470$ and $K = 147$, where the K samples have been chosen so that they all belong to the same signer.

Preliminary recognition tests have been executed and, at the fourth step, the selection reached a maximum accuracy of 38.76% with the combination of features 5, 13, 7 and 8: hands centroid center derivatives contours, hands area derivative, hands centroid center contours and derivatives referring at head coordinates. The analysis of these results has led to the introduction of crucial improvements (e.g., smoothing of the head position), already discussed in Section 3.1

Table 2. Sequential Forward Selection experimental results (accuracy of first four best selection/selected step). The “Total” row values are obtained as an average of the recognition rate performed for each cross-validation step, that is for each signer.

Signer	1 SFS	2 SFS	3 SFS	4 SFS
	Feature 5	Feat. 5-13	Feat. 5-13-8	Feat. 5-13-8-7
fal	51.03%	55.86%	57.24%	47.59%
fef	17.93%	20.69%	15.86%	22.07%
fsf	28.28%	37.24%	30.34%	35.86%
mdp	39.31%	42.76%	46.21%	44.83%
mdq	47.59%	49.66%	46.21%	53.10%
mic	48.97%	51.03%	53.10%	45.52%
mmr	49.66%	57.93%	52.41%	55.17%
mrla	49.66%	53.10%	51.72%	61.38%
mrlb	50.34%	55.17%	56.55%	60.00%
msf	48.97%	48.97%	46.90%	46.21%
Total	43.17%	47.24%	45.65%	47.17%

The best results of the subsequent feature selection are reported in Table 2. Feature 5 obtains the best performance (43.17%) for the first evaluation of the feature selection. In the second step the combination with feature 13 achieves an average recognition rate of 47.24%. The combinations with the feature 8, at the third step, and then with the feature 7, fourth step, produce respectively an overall accuracy of 45.65% and 47.17%. At the end of the sequential forward feature selection, an overall best accuracy of 47.24% is obtained at the second step of the selection by the combination of the hands centroid center derivatives and hands area derivative.

The main causes of errors have been analyzed by means of the confusion matrix among all signers. The analysis gives better insights of the system weak points, and indicates which signs limit the overall recognition accuracy. One of the most frequent causes of errors is the overlap between hands, and between head and hands, which causes the wrong revelation of their position and boundaries. Another source of errors, is the inability to discriminate in a proper way the 2-D projection respect the 3-D true form of a hand. This problem is emphasized if there are signs that have the same movements for the arms. In addition, some movements, such as those parallel to the line of sight of the camera, are not detectable in a 2-D field.

5 Conclusions

In this paper, a preliminary study of an automatic recognition system of Italian Sign Language based on Hidden Markov Models has been proposed. Using features obtained from the position of hands and face, a study on the parameters of the model representing each sign has been proposed, and the feature set has been selected by means of the Sequential Forward Selection approach. Experiments have been performed on the A3LIS-147 database of isolated Italian signs through the cross-validation approach, this way guaranteeing the orthogonality between training and test sets. In the final system, each sign is represented with an HMM composed of 7 states and 5 mixtures and 4 of the original 18 features are used. The obtained recognition accuracy averaged across all signers is 47.24%, with a maximum of 57.93%. Noteworthy the recognition accuracy of 47.17% reached by the fourth step of the selection with a remarkable maximum signer accuracy of 61.38%.

In order to have better insights on the performance of the system, the confusion matrix among all signers has been analysed: the main causes of errors identified are the overlap between head and hands and between the two hands, and the inability to discriminate the 3-D position of the subject. These problems are more evident for the signs related to the days of the week, they have similar arm movements and the information is entirely contained in the hands shapes. Therefore, an overlap impedes to detect the correct shape of the hands resulting in loss of valuable information. The computational burden of the recognition process has been evaluated in terms of real-time factor, defined as the ratio between the processing time and the length of the video sequence. On average, the whole

process, requires about 8 seconds, and video segments are 3 seconds long. This results in a real-time factor of about 2.6, without taking into account solutions to lower the computational cost.

As future works, the overlap problem will be primarily addressed with the introduction of tracking techniques, as Mean-Shift or CAM-Shift. In addition, the inclusion of novel features containing lip movements and Histogram of Oriented Optical Flow (HOOF) [5] will be investigated. HOOF features, in particular, are suitable for real-time computation, and could be obtained directly from the flow information encoded in a compressed video. Finally, the authors are considering the development of a new sign database including audio and spatial information (e.g., using multiple cameras, or Microsoft KinectTM) to augment the feature set and reduce errors due to overlapping and wrong spatial detection of the subject.

References

1. von Agris, U., Zieren, J., Canzler, U., Bauer, B., Kraiss, K.F.: Recent Developments in Visual Sign Language Recognition. *Univers. Access Inf. Soc.* 6(4), 323–362 (2008)
2. Bradski, G.: The OpenCV library. *j-DDJ* 25(11), 122–125 (2000)
3. Bradski, G., Kaehler, A.: *Learning OpenCV: Computer Vision with the OpenCV Library*. O'Reilly Media, Inc., Sebastopol (2008)
4. Canny, J.: A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 8(6), 679–698 (1986)
5. Chaudhry, R., Ravichandran, A., Hager, G., Vidal, R.: Histograms of Oriented Optical Flow and Binet-Cauchy Kernels on Nonlinear Dynamical Systems for the Recognition of Human Actions. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1932–1939. IEEE Computer Society, Los Alamitos (2009)
6. Dreuw, P., Forster, J., Deselaers, T., Ney, H.: Efficient Approximations to Model-Based Joint Tracking and Recognition of Continuous Sign Language. In: *8th IEEE International Conference on Automatic Face Gesture Recognition*, pp. 1–6. IEEE Computer Society, Los Alamitos (2008)
7. Dreuw, P., Neidle, C., Sclaroff, V.A.S., Ney, H.: Benchmark Databases for Video-Based Automatic Sign Language Recognition. In: *Proc. of the Sixth International Conference on Language Resources and Evaluation (LREC 2008)*. European Language Resources Association (ELRA), Marrakech (2008)
8. Dreuw, P., Ney, H.: Visual Modeling and Feature Adaptation in Sign Language Recognition. In: *ITG Conference on Voice Communication (SprachKommunikation)*, pp. 1–4 (2008)
9. Dreuw, P., Ney, H., Martinez, G., Crasborn, O., Piater, J., Moya, J.M., Wheatley, M.: The SignSpeak Project - Bridging the Gap Between Signers and Speakers. In: *Proc. of the Seventh International Conference on Language Resources and Evaluation (LREC 2010)*. European Language Resources Association (ELRA), Valletta (2010)
10. Fagiani, M., Principi, E., Squartini, S., Piazza, F.: A New Italian Sign Language Database. In: Zhang, H., Hussain, A., Liu, D., Wang, Z. (eds.) *BICS 2012. LNCS*, vol. 7366, pp. 164–173. Springer, Heidelberg (2012)
11. Hu, M.-K.: Visual Pattern Recognition by Moment Invariants. *IRE Trans. on Information Theory* 8(2), 179–187 (1962)

12. Infantino, I., Rizzo, R., Gaglio, S.: A Framework for Sign Language Sentence Recognition by Commonsense Context. *IEEE Trans. on Systems, Man, and Cybernetics, Part C: Applications and Reviews* 37(5), 1034–1039 (2007)
13. Kukharev, G., Nowosielski, A.: Visitor Identification - Elaborating Real Time Face Recognition System. In: *Proc. of WSCG (Short Papers)*, pp. 157–164. UNION Agency, Plzen (2004)
14. Maebatake, M., Suzuki, I., Nishida, M., Horiuchi, Y., Kuroiwa, S.: Sign Language Recognition Based on Position and Movement Using Multi-Stream HMM. In: *Proc. of the 2nd Int. Symposium on Universal Communication*, pp. 478–481. IEEE Computer Society, Los Alamitos (2008)
15. Sandler, W., Lillo-Martin, D.: *Sign Language and Linguistic Universals*. Cambridge University Press, Cambridge (2006)
16. Serra, J.: *Image Analysis and Mathematical Morphology*. Academic Press, Inc., Orlando (1983)
17. Theodorakis, S., Katsamanis, A., Maragos, P.: Product-HMMs for Automatic Sign Language Recognition. In: *Proceedings of the 2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 1601–1604. IEEE Computer Society, Washington (2009)
18. Theodoridis, S., Koutroumbas, K.: *Pattern Recognition, Fourth Edition*. Academic Press, Burlington (2008)
19. Young, S.J., Evermann, G., Gales, M.J.F., Hain, T., Kershaw, D., Moore, G., Odell, J., Ollason, D., Povey, D., Valtchev, V., Woodland, P.C.: *The HTK Book*, version 3.4. Cambridge University Press, Cambridge (2006)

Fall Detection Using an Ensemble of Learning Machines

Simon Bulotta¹, Hassan Mahmoud¹, Francesco Masulli^{1,2},
Ernesto Palummeri³, and Stefano Rovetta³

¹ DIBRIS - Università di Genova, Genova, Italy

² Center for Biotechnology, Temple University, Philadelphia, USA

³ Dipartimento di Gerontologia e Scienze Motorie, E. O. Ospedali Galliera, Genova, Italy
ernesto.palummeri@galliera.it

Abstract. A random ensemble of random perceptrons is studied and applied in fall detection and categorization, an important and growing problem in Ambient Assisted Living and other fields related to the care of elder and in general of “fragile” people. The classifier ensemble is designed around an ECOC aggregator and compensates for the lack of an accurate training with the number of base learners, which increases accuracy and strengthens the error-correcting capabilities of class codewords. The approach is suitable when some memory is available, but computational power is limited: this is the standard situation in mobile computing, and to an even larger extent in wearable computing. Performances on the two applicative tasks of fall recognition (dichotomic) and categorization (multi-class) are compared with those of support vector machines.

Keywords: Ensemble classifiers, Fall detection, Ambient-assisted living.

1 Introduction

As technology progresses, advances in different fields can mutually compensate for each other, so that some technical problems remain in spite of the evolution. Such is the case in resource-constrained computing. Programming and system design techniques that were once necessary to exploit primitive processors and architectures are obviously no longer needed for mainstream computers; however, advances in microelectronics have made it possible to provide computing capabilities to many devices in the range size from small to tiny, such as mobile phones, wearable computers, or even textile computers (integrated in fabric). These devices have inherent and obvious limitations, so that special software design techniques are still necessary.

The research presented in this work is motivated by the need to target this type of resource-constrained computers for an application that aims at integrating transparently into the live of a growing class of citizens, the elder and more generally the so-called “fragile” subjects. The application, which is part of a collection of related tasks for which similar techniques are in order, is the detection and categorization of falls, an event that occurs with a surprisingly high frequency and may have short- and long-term consequences.

Specifically, we present a neural classifier ensemble designed around an ECOC aggregator. The model compensates for the lack of an accurate training with the number of base learners, which increases accuracy and strengthens the error-correcting capabilities of class codewords. The approach is suitable when some memory is available, but computational power is limited.

2 Fall Detection

2.1 Significance of Falls

The treatment of the injuries and complications associated with falls of older adults is a notable social expense. For instance, every year one-third to one-half of the U.S. population aged 65+ experience fall injuries, generating a public expense of over 20 billion dollars [14]. Falls are the leading cause of moderate to severe injuries in older adults (e.g., hip fractures), and the leading cause of accidental death for subjects over 75. These injuries can make it difficult for the subject to get around and live independently, and increase his/her risk of an early death. Moreover, regardless of the actual injuries, experiencing a fall may generate the fear of falling again, in a feedback cycle that ultimately increases the actual risk of falling.

In case of a fall injury it is vital to have an alert system to signal for help. For this purpose, portable or wearable sensors appear to be the best technical choice.

2.2 Sensors for Fall Detection

In the last ten years many studies focused on fall detection systems based on various sensor types, including piezoelectric, infrared, inertial, microphones, and video-cameras. However, most ambient-installed sensors have a limited range of operation and can only be used in indoor environments, in most cases with the need for a permanent or semi-permanent installation. A clear example of this issue is represented by “intelligent floors”, but this is also true for other modalities.

By contrast, accelerometer-based sensors can be manufactured in very small sizes and are suitable for portable/wearable devices, even for outdoor use. Accelerometers are inertial sensors based on electro-mechanical devices that can measure dynamic accelerations resulting from movements or shocks, and/or static acceleration, such as gravity (tilt sensors). Micro Electro-Mechanical Systems (MEMS) accelerometers are very popular for their accuracy, compactness, low power consumption, low cost. As a consequence, they are widely used in consumer electronics, especially in mobile devices. Due to their compactness, MEMS sensors make it possible to integrate more sensors for multiple axes in a single chip.

2.3 Fall Detection Algorithms

The recent literature proposes many methods for fall detection and classification. Threshold algorithms are widely applied to fall detection using MEMS accelerometers [11], including several sensors with hardware built-in fall detection. However, research indicates that those approaches can have limited sensitivity for some types of

falls. Current fall-detection approaches can also result in poor specificity, in part because impacts can be experienced as part of ordinary daily living activities. The accuracy obtained with these methods is over 80 %.

Many machine learning and intelligent computation methods have also been applied to fall detection using MEMS data. The best results from the literature report accuracies ranging from 93% to 97% [15][17], but comparative performance assessment is not easy because of the different experimental and data acquisition conditions of each individual study. To the best of our knowledge, no standardized benchmark data collection exists in this very specific field. The actual data acquisition and preprocessing procedure followed in the present study is described in subsection 5.1.

3 The Method

3.1 Ensemble Classifiers

An *ensemble classifier* or *multiple classifier system* or *committee machine* [8] uses more than one classifier to increase classification performance with respect to an individual classifier.

Given a classification learning task represented by a set of labeled examples $[X|Y] = \{(x_1, y_1), \dots, (x_P, y_P)\}$, where the x_l are D -dimensional input patterns and the y_l are corresponding class labels from the set of classes $\Omega = \{\omega_1, \dots, \omega_C\}$, an individual learning algorithm (such as a neural network) outputs a hypothesis $y = H(x)$, a *classifier*. An ensemble classification method is formed by training a set of N *base classifiers* H_1, \dots, H_N . A subsequent aggregation step is performed by a combiner, which takes all individual classifications y_1, \dots, y_N as inputs and outputs a collective decision y . The simplest combiner is probably the majority voting criterion, $y = \arg \max_{\omega_k} \{\text{freq}(\omega_1), \dots, \text{freq}(\omega_C)\}$, where $\text{freq}(\omega_k)$ is the fraction of classifiers that output the class ω_k .

To provide an advantage with respect to a single classifier, multiple classifiers co-operating in an ensemble should be *diverse*. What constitutes diversity is subject of discussion [8], and there are several ways to differentiate base classifiers [5][3] (operating on the data, on the training procedure, on the base classifier structure, . . .). However, it is quite obvious that a set of identical base classifiers has no reason to perform better than any one of them taken individually.

There are several reasons why an ensemble classifier should be able to outperform an individual base classifier [5]. The first is statistical: diverse hypotheses can provide a more thorough exploration of the hypothesis space, providing better evidence on their relative distribution. Another reason is computational: if diversity is obtained by training several classifiers with different random starting points, the need for accurate minimization of the objective function (often a difficult task) is reduced. Yet another reason is that the representational power of individual classifiers should be kept under control, to avoid over-fitting, so that having multiple learners provides a set of “alternative opinions”, any one of which is not sufficiently powerful to solve the classification problem: collectively, they can have the required representational ability.

An interesting fact in ensemble learning stems from the observation that perfect classifiers cannot be diverse, since they are all identical to the perfect hypothesis. In fact, in some methods [6], classifiers are only required to have just the minimum bias toward

being correct, i.e., in two-class problems they only need to be correct just a bit more than 50% of the time. Such base classifiers are termed *weak*.

3.2 ECOC Ensembles

One way to aggregate individual classifiers into an ensemble is to encode the required output into binary words, acting as output codes, and then have the individual base classifiers learn one bit of each code. In this way, a multi-class ($C > 2$) problem can be stated in terms of a set of dichotomic problems. Note that this is the standard procedure as employed in neural network classifier to translate a multi-class problem into a set of values representable by output neurons, which only have a binary output.

There are many kinds of coding schemes [12]; one classical choice is 1-out-of- C encoding, where the output codeword is of length C bits and the bit corresponding to the encoded class is 1 while all others are 0. This encoding scheme is popular because it has a soft version directly related to logistic regression, but other schemes (e.g. thermometer encoding) are possible as well.

Each encoding has features that make it desirable in a specific framework. A particular choice of a family of encoding schemes was proposed by Dietterich and Bakiri [4]. They employed error-correcting codes [7] to provide diversity between the tasks learned by base classifiers.

In the ECOC ensemble approach, the codes are generated so as to optimize the discernability of classes. Each class is attributed a bit string, and each base classifier is assigned the task of classifying one of these bits. As a consequence, this encoding splits the multi-class problem into a set of dichotomic problems, with the property that they are as diverse as possible. The most well-known example of error-correcting encoding is Hamming codes, a block coding method that is commonly presented in any digital electronics and computer architecture text book, since it is adopted (for instance) in RAM modules.

To recover the output class, the Hamming distance between the output bit string and each code is computed; the class corresponding to the nearest code (minimum Hamming distance) is selected as the overall output. This strategy has been found to be particularly effective [10]. However, its application is usually focused on small or moderately large ensembles, for which the best possible encoding is sought; diversity is due to the different tasks assigned to individual classifiers. In the next section we propose an alternative approach where a huge ensemble with sub-optimal encoding and very weak learners is pursued.

4 Random Ensembles of Random Classifiers

To limit the computational requirement, we consider an error-correcting-output-encoded ensemble (ECOC ensemble). We adopt simple perceptrons as base classifiers. However:

- the number of base learners is high (> 1000);
- base learners are random, except for one bit;
- output codes are random.

The rationale for this set of extreme choices is as follows. To feature the error-correcting property, output codes should be well separated in terms of Hamming distance. Due to the well-known concentration effect on distances in high dimensionality, randomly generated codewords feature remarkably uniform mutual distances [2]:

$$\lim_{d \rightarrow \infty} \frac{\text{var}[d]}{E[d]^2} = 0 \Rightarrow P[\max d < (1 + \varepsilon) \min d] = 1, \quad (1)$$

where E and var are the expectation and variance operators, respectively, d is the distance between any two points in the data space, a random variable, and ε is an arbitrarily small quantity.

Output codes are strings of bits as long as the number of base learners, which, in our case, is larger than 1000. In the experimental results presented here, we used 5000 base classifiers. As a result, although the obtained codes are not optimal, we have a statistical guarantee that they will be sufficiently differentiated as to permit an efficient error-correcting behavior.

To experimentally check this point, we generated sets of 6 random codes of growing size N (100, 1000, 10000) and recorded the maximum and minimum inter-code distances. The result of one experiment, fully representative of the general behavior, is as follows: for $N = 100$, distance was in the range [38, 52]; for $N = 1000$, distance \in [408, 459]; and for $N = 10000$, distance \in [4109, 4324].

The distances are contained within a range which depends on the dimension, but is equal or less than 5% of the whole range of distance values.

Regarding the base learners (simple perceptrons), we adopted the following strategy. The data are normalized in the unit hypercube. The perceptron weights are generated at random, with only a statistical bias to slightly increase the probability that the separating hyperplanes cross the center of the data area.

The base learners are then evaluated on the training set, and for those whose percentage of correct classification does not exceed 50%, *the sign of the weights is reversed*.

This simple-minded procedure ensures that all base classifiers are trained, although they are definitely weak. A similar procedure was used in [18] for the synthesis of output codes. However, in that case, this resulted in changing the encoding of the problem, requiring an iterative procedure.

The novelty of our approach lies in the fact that, rather than trying to overcome the design problem, we simply exploit the collective properties of a large number of components. We feel that this places our work in line with the original ideas underlying Rosenblatt's work [11]. The experimental results will show that, although the performance is not optimal and is subject to a statistical oscillation, it is remarkably close to that of a state-of-the-art classifier, a multi-class support vector machine.

A refinement to the procedure is introduced by *filtering out* the base classifiers depending on their performance level. If, after the sign-flipping, a base classifier still does not exceed a given performance threshold, it is deactivated and not used in the ensemble. This reduces the size of the ensemble itself. The results show that the two properties, ensemble size and individual learner's accuracy, balance up to a certain quality threshold, after which the overall performance decreases.




Fig. 1. Examples of falls. Left: on knees. A is the first part of the fall; B is the first impact with the ground. Right: abrupt sitting on hard chair. Note that temporal profiles are similar.

5 Experiments

5.1 Data Acquisition

Thanks to the enormous progress made in the field of electronic miniaturization and production, nowadays many smartphones include MEMS inertial sensors. Following several other researchers [13,16], we have employed a smartphone as a fall detecting device, specifically an HTC Desire Z model A7272 smartphone with Qualcomm MSM 7230 CPU with 800 MHz clock, and Android operating system. This phone includes a three-axis BMA150 accelerometer with G-sensor capable of measuring the static gravitational acceleration. Suitable software was obtained from the “Android Market” repository in the form of “apps” (Android application programs).

The accelerometer was tested using the *Accelerometer Checker* app by Paul Thomas. The electronic noise is on average much less than 1% of the measured values of accelerations. The systematic error is below 5%. For data collection we used the *Accelerometer* app by Chris Pearson, working with a sampling rate of 65 Hz. The collected data were then accessed through USB.

We obtained our data set from 6 healthy volunteers, aged between 25 and 27, practitioners of the athletic discipline Parkour and therefore trained to deal with several types of falls. They were requested to fall in a safe way, wearing the accelerometer placed on the chest.

The main fall categories considered in the literature as the most frequent/dangerous are as follows:

1. Forward falls, e.g., when one stumbles and falls.
2. Backward falls, e.g., when one slides and falls.
3. Backward falls on the place, e.g., in the event of temporary loss of balance.
4. Lateral falls (left or right), e.g., in case of ankle’s sprain or weakness of a foot.
5. Falls on knees, e.g., in case of illness or failure of the ankles.

In addition, we have collected some additional cases:

6. Forward fall with unsuccessful attempt to hold to a side support with both hands.
7. Abrupt sitting on hard chair (not a true fall).




Fig. 2. Average number of classifiers as a function of the filtering threshold

For each type of fall we collected the data from at least three volunteers. Each simulated fall was repeated five times per subject. Each fall is represented by 111 samples spaced by 15 ms, for a total of about 1.7 s. Start of recording was triggered by hand; a threshold criterion can also be used. Our data set contains about a dozen acquisitions for each type of simulated fall. Following the recent literature, we only consider acceleration magnitude rather than the full acceleration vector.

For each type of fall we noticed a very small inter-individual and intra-individual variability. In Fig. 1 on the left is the temporal profile of the acceleration during a fall on knees. The abrupt sitting on hard chair (right) profile looks quite similar.

After cleaning, the training set is composed of 9 cases for each of the 7 classes (six different fall types, one non-fall), for a total of 54 patterns of dimension 111.

5.2 The Classification Tasks

All experiments are performed for a range of values of the filtering threshold. For each value the statistics on 10 runs are recorded. There are two different tasks:

1. **Fall detection:** detect a fall.

This task implies discriminating one type of fall (here class 1) from the non-fall class (class 7). We only employed one type of fall for this experiment to keep classes more balanced and on the basis of the experimentally verified similarity between all types of fall.

This is an easy, two-class problem.

2. **Fall type discrimination:** detect the type of a fall.

This task is very hard, given the simple encoding of the data: as explained above we used the sampled values directly. It involves the discrimination of each type of fall from the others (classes 1–6). The patterns are very similar and performance of a state-of-the-art method is not very good, as will be shown in the results.

This is a complex, multi-class problem.

In both cases, the number of base classifiers was reduced as the quality threshold grew (Fig. 2). This accounts for the decrease in quality visible in the graphs.




Fig. 3. Result on the fall detection task. Minimum, first quartile, median, third quartile, maximum of the percentage of correct classifications, as a function of the filtering threshold

Fall detection: the results shown in the first graph (Fig. 3) confirm that two-class fall detection is easy. Note that the results are statistically variable, as expected, due to the random nature of the method. A support vector machine is able to separate the two classes, although this requires an optimization step and parameter tuning.

Fall type discrimination: the experiments show that the problem is much more difficult. Note, however, that in this phase of the research only the acceleration magnitude was used, and raw samples were feed into the classifier without feature extraction or other refinements. The results are again variable, but, with respect to the previous task, they are both less stable and worse.

However, the best leave-one-out results obtained with a support vector machine are not dramatically better. Using the multi-class method implemented in MSVMpack [9] the experimental performance is as reported in Table 1.

Table 1. Performance of SVM

Model	Correctly classified	Nr. of support vectors
SVM - linear kernel	26 (48.1%)	53
SVM - Gaussian kernel	25 (46.3%)	53

Note that these result required several iterations to select the best parameters and also a difficult optimization, leading to the need to reduce the optimization accuracy parameter to obtain convergence in reasonable time. Note also that the number of support vectors is always equal to the number of data points (53 for leave-one-out training with a training set size of 54), a clear indication of the problem complexity.

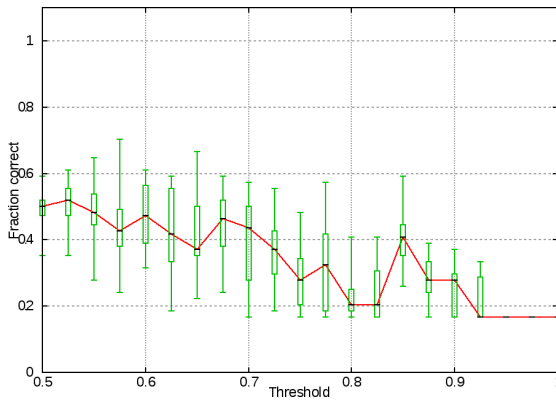


Fig. 4. Result on the fall type discrimination task. Minimum, first quartile, median, third quartile, maximum of the percentage of correct classifications, as a function of the filtering threshold

6 Discussion and Conclusion

The fall detection application has served as a testbed to explore the proposed technique. The two tasks provide different scenarios. The first one is clearly easy, with fairly good performance achieved by standard methods. The other one is much harder, not easy to tackle even with state-of-the-art techniques. Accordingly, our random technique features much worse results. Still, given the extremely simple-minded training criterion implemented, the obtained results can be considered good.

The methods indicates that the sheer number of computing elements can compensate for their (lack of) accuracy. This opens the way for interesting theoretical work and for exploring the connection with other approaches.

References

1. Bagnasco, A., Scapolla, A., Spasova, V.: Design, implementation and experimental evaluation of a wireless fall detector. In: Proceedings of the 4th International Symposium on Applied Sciences in Biomedical and Communication Technologies, p. 65. ACM (2011)
2. Beyer, K., Goldstein, J., Ramakrishnan, R., Shaft, U.: When Is Nearest Neighbor Meaningful? In: Beeri, C., Bruneman, P. (eds.) ICDT 1999. LNCS, vol. 1540, pp. 217–235. Springer, Heidelberg (1998)
3. Brown, G., Wyatt, J., Harris, R., Yao, X.: Diversity creation methods: a survey and categorisation. *Information Fusion* 6(1), 5–20 (2005)
4. Dietterich, T., Bakiri, G.: Solving multiclass learning problems via error-correcting output codes. *Journal of Artificial Intelligence Research* 2(263), 286 (1995)
5. Dietterich, T.G.: Ensemble Methods in Machine Learning. In: Kittler, J., Roli, F. (eds.) MCS 2000. LNCS, vol. 1857, pp. 1–15. Springer, Heidelberg (2000)
6. Freund, Y., Schapire, R.: A decision-theoretic generalization of on-line learning and an application to boosting. In: Computational Learning Theory, pp. 23–37. Springer (1995)

7. Hamming, R.W.: Error detecting and error correcting codes. *Bell System Technical Journal* 29(2), 147–160 (1950)
8. Kuncheva, L.: Combining pattern classifiers. In: *Methods and Algorithms*. Wiley, Chichester (2004)
9. Lauer, F., Guermeur, Y.: MSVMpack: a multi-class support vector machine package. *Journal of Machine Learning Research* 12, 2269–2272 (2011)
10. Masulli, F., Valentini, G.: Effectiveness of error correcting output coding methods in ensemble and monolithic learning machines. *Pattern Analysis & Applications* 6(4), 285–300 (2004)
11. Rosenblatt, F.: *Principles of Neurodynamics*. Spartan, New York (1962)
12. Rumelhart, D.E., McClelland, J.L.: *Parallel Distributed Processing*. MIT Press, Cambridge (1986)
13. Sposaro, F., Tyson, G.: iFall: An Android application for fall monitoring and response. In: *Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC 2009*, pp. 6119–6122. IEEE (2009)
14. Stevens, J., Corso, P., Finkelstein, E., Miller, T.: The costs of fatal and nonfatal falls among older adults. *Injury Prevention* 12, 290–295 (2006)
15. Struck, M., Dinh, C.: A new real-time fall detection approach using fuzzy logic and a neural network. In: *Proceedings of the 6th International Workshop on Wearable, Micro and Nano Technologies for the Personalised Health, pHealth* (2009)
16. Tacconi, C., Mellone, S., Chiari, L.: Smartphone-based applications for investigating falls and mobility. In: *Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, pp. 258–261. IEEE (2011)
17. Yuwono, M., Moulton, B., Su, S., Celler, B., Nguyen, H.: Unsupervised machine-learning method for improving the performance of ambulatory fall-detection systems. *Bio Medical Engineering OnLine* 11(1), 9 (2012)
18. Zor, C., Yanıkoğlu, B., Windeatt, T., Alpaydın, E.: FLIP-ECOC: a greedy optimization of the ECOC matrix. In: Gelenbe, E. (ed.) *Computer and Information Science: Proceedings of the 25th International Symposium on Computer and Information Sciences*, vol. 62, Springer-Verlag New York Inc. (2010)

Part II

Signal Processing

PM₁₀ Forecasting Using Kernel Adaptive Filtering: An Italian Case Study

Simone Scardapane, Danilo Comminiello, Michele Scarpiniti,
Raffaele Parisi, and Aurelio Uncini

Department of Information Engineering, Electronics and Telecommunications (DIET),
“Sapienza” University of Rome, Via Eudossiana 18, 00184, Rome
simonescardapane@gmail.com,
{danilo.comminiello,michele.scarpiniti,
raffaele.parisii}@uniroma1.it, aurel@ieee.org

Abstract. Short term prediction of air pollution is gaining increasing attention in the research community, due to its social and economical impact. In this paper we study the application of a Kernel Adaptive Filtering (KAF) algorithm to the problem of predicting PM₁₀ data in the Italian province of Ancona, and we show how this predictor is able to achieve a significant low error with the inclusion of chemical data correlated with the PM₁₀ such as NO₂.

Keywords: Air pollution, Nonlinear adaptive filtering, Kernel Adaptive Filters, Least Mean Square.

Introduction

Due to its adverse effects on human health, airborne particulate matter has gained over the last years increasing attention, both inside the research community and in the parliaments worldwide. Numerous studies have found a direct link between inhalation, long term exposure to particulate matter, and increase in mortality rates, particularly for lung cancers [1]. Moreover, a continuous presence of particulate matter lead to a constant decrease in visibility inside the cities and to the deposition of trace elements [2].

To deal with these aspects, the European Commission issued on 22 April 1999 the Council Directive 1999/30/EC, that set a roof for daily concentration of particulate matter with an aerodynamic diameter of up to 10 μ m (PM₁₀), and obliged member states to issue a warning every time this roof is reached, entering an “*attention state*”. Nevertheless, last data coming from the European Environmental Agency [3] shows that, although global pollution has decreased, some pollutants such as PM₁₀ have stayed more or less stable and constantly exceeds the required threshold. For this reason, the development of a solid system that is able to accurately predict the daily levels of PM₁₀ has been, and still is, a constant priority for local administrations.

¹ <http://www.eea.europa.eu/publications/air-quality-in-europe-2011>

Since environmental data presents a great complexity and an ample deal of hidden factors, numerous researchers have investigated the possibility of automatically forecasting the daily concentration of PM_{10} using learning systems. In particular, neural networks [3] have been found to achieve a moderate error with small amount of data, and thus represent a good solution for implementing a robust prediction system [4]. Recently, a moderate deal of effort has been put into applying other methodologies such as Support Vector Machines [5] to the problem.

In this paper, we study the application of Kernel Adaptive Filtering (KAF) techniques [6] to the task of predicting PM_{10} in the Italian area of Ancona. KAF algorithms are a recent development in the adaptive filtering field, and have been rarely put to use for real world problems. For this reason, it is interesting to see how they provide a fast and efficient solution to our problem. In addition, comparisons with other standard techniques provide similar results, but the proposed approach is characterized by a low computational cost.

We also investigate the inclusion in the learning process of chemical data such as NO_2 , showing how this lead to a strong decrease in the prediction error, probably compensating for hidden factors that occasionally falsify the instruments readings. Several studies, for example [7,8,9,10], have addressed the problem of PM_{10} forecasting using cross-prediction with different chemical agents. Cross-prediction, in fact, can provide a robust estimate in all those cases, in which the nature of the time-series to be predicted can have several contributes. For example, because the area of Ancona is on the sea-side, many external factors, such as the sea salt, can give a misunderstood contribution on PM_{10} values, falsifying the results and alarming people even if the concentration is less than that measured. In order to obtain a more robust estimation of the PM_{10} values we choose to use cross-prediction in this work.

The paper is organized as follows: Section 1 introduces the Kernel Adaptive Filtering approach. Section 2 describes the experimental setup, while Section 3 shows the experimental results. Finally Section 4 concludes the work.

1 Kernel Adaptive Filtering

Kernel Adaptive Filtering (KAF) [6] is a recent family of learning algorithms that combines the simplicity of Linear Adaptive Filtering with the nonlinear modeling capabilities of other learning techniques such as Neural Networks and Support Vector Machines. KAF are kernel methods, meaning that the original input $\mathbf{x} \in \mathbb{R}^s$ to the problem is transformed into a highly dimensional feature vector $\psi(\mathbf{x})$ through the use of a positive definite kernel function:

$$\kappa(\mathbf{x}, \mathbf{x}') : \mathbb{R}^s \times \mathbb{R}^s \rightarrow \mathbb{R}.$$

Choosing an appropriate kernel function, the newly created input vector is linearly separable, and is used to train a linear adaptive filter. Practically, as long as the original linear algorithms can be formulated in terms of inner products, the use of the kernel trick [11] allows to rewrite them in terms only of kernel evaluations, thus avoiding the explicit computation of the transformed vector.

Algorithm 1. Summary of the KLMS algorithm

Initialize: Training Set T , η , κ

- 1: $a[1] = \eta d[1]$
 - 2: **while** $(\mathbf{x}_n, d[n]) \in T$ available **do**
 - 3: $y[n] = \sum_{j=1}^{n-1} a[n-j] \kappa(\mathbf{x}_n, \mathbf{x}_{n-j})$
 - 4: $e[n] = d[n] - y[n-1]$
 - 5: $a[n] = \eta e[n]$
 - 6: **end while**
-

Starting from linear filtering theory, a number of kernel filters have been devised, such as the Kernel Least-Mean Square [12] that is used in this work. They all provide nonlinear approximation to any function and moderate complexity, with respect to other methods such as recurrent neural networks. The main problems arising in the use of a KAF are the choice of a proper kernel, the need for regularization and the fact that the network grows linearly with the number of processed inputs. This last problem is not addressed in the current work. For a review of methods to efficiently curtail the growth of the network, we recommend to the interested reader [12].

1.1 Kernel Least Mean Square

Kernel Least Mean Square (KLMS) [12] is the simplest training algorithm in the KAF family, and is derived by Least Mean Square (LMS) [3] by the use of the kernel trick. Despite its simplicity, it has been shown to own a self regularizing property.

The pseudo-code of KLMS is briefly summarized in Algorithm 1. It takes as input a training set T composed of pairs $\{(\mathbf{x}_n, d[n])\}$, where $\mathbf{x}_n \in \mathbb{R}^s$ is the input vector to the system, and $d[n] \in \mathbb{R}$ the desired scalar output. There is a need to select a proper *step size* η and the kernel function $\kappa: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$, then, for every processed input, the algorithm stores the current input vector \mathbf{u}_n and the corresponding weight $a[n] = \eta e[n]$. The step size η balances between speed and convergence, and it can be chosen using a reasoning like the classical LMS algorithm. The output of the filter at time n is given by

$$y[n] = \sum_{j=1}^{n-1} a[n-j] \kappa(\mathbf{x}_n, \mathbf{x}_{n-j}). \quad (1)$$

In equation (1) $\kappa(\mathbf{u}_n, \mathbf{u}_{n-j})$ refers to the value of the kernel. In this work we used the Gaussian kernel, that, for two generic input vectors \mathbf{u} and \mathbf{u}' becomes:

$$\kappa(\mathbf{u}, \mathbf{u}') = \exp(-\gamma \|\mathbf{u} - \mathbf{u}'\|^2), \quad (2)$$

where γ is the *kernel bandwidth* or *kernel parameter*. In addition $\kappa(\mathbf{u}, \mathbf{u}')$ has the interesting property of being shift-invariant.

As a measure of the error we used the mean-square error, which is defined over the training set T as:

$$\text{MSE}_T(\mathbf{a}_n) = \frac{1}{N} \sum_{(\mathbf{x}_n, d[n]) \in T} (d[n] - y[n])^2. \quad (3)$$




Fig. 1. Hourly observations of PM_{10} for the year 2010 (in total 8764 observations) in Marina di Ancona

2 Experimental Setup

In order to test the proposed predictor, we used hourly observations of PM_{10} and NO_2 for the year 2010 in Marina di Ancona (Italy)². These data are interesting since, as for every maritime region, accuracy of the sensors can be falsified by the presence of external factors such as sea salt. We compensate this aspect with the inclusion of chemical data given by NO_2 .

Data is normalized between 0 and 1, and the resulting PM_{10} series can be seen in Figure 1. We observe that the concentration of particulate matter has stayed more or less stable around the year, while some observations are clearly outliers.

In the first experiment, the input vector \mathbf{x}_n is an embedding of the last 168 elements of the PM_{10} $\{x[n-1], \dots, x[n-\tau]\}$, corresponding to the previous 7 days of observations, while the desired output is the concentration of PM_{10} for twelve p.m. of next day. In the second experiment, the corresponding observations of the NO_2 are added to each input vector, leading to a 336 elements.

From the original series we extracted randomly 700 input vectors and corresponding outputs for training, and 150 elements for testing, thus generating a statistically inde-

² Data can be freely downloaded from

http://ww3.provincia.ancona.it/dati_ambiente/riepiloghi/

pendent testing set. Experiments are averaged on 50 different runs, and are conducted using an Intel i3 3.07 GHz processor at 64 bit, with 4 GB of RAM available.

2.1 Correlation between PM₁₀ and NO₂

NO₂ was chosen as a supporting time series since it has been found to be highly correlated with the particulate matter evolution. In Figure 2 we can visually see the correlation between the two time series for a brief period. To confirm this dependence, we computed the correlation coefficient ρ of the covariance matrix $\mathbf{C} = \text{cov}(\mathbf{X})$, where \mathbf{X} is a 2×8760 matrix containing all the hourly observations. The correlation coefficient ρ was found to be greater than 0.4, thus confirming our initial hypothesis.




Fig. 2. Correlation between PM₁₀ and NO₂ for a brief period of the year 2010 in Marina di Ancona

Figure 2 clearly shows that, around sample 210 (emphasized with a black cross), the values of measured PM₁₀ are greater than the values of NO₂, differently from other samples. This fact could be due to external factors, such as the sea salt. This latter assumption is argued by means of the hourly weather condition in that day. In fact there was a strong wind from the sea towards the city.

3 Results

Figure 3 shows the evolution of the MSE for the two cases, with and without the inclusion of NO₂ chemical correlated data. The learning rate η was set to 0.2 for the first case and to 0.05 for the second, while the kernel parameter has remained constant at $\gamma = 1$. The error is also compared with a standard LMS filter with $\eta = 0.001$.




Fig. 3. Evolution of the mean square error (MSE) for KLMS in PM₁₀ prediction with and without the inclusion of NO₂ correlated data

We see that KLMS achieves a low error over the testing set in less than 700 iterations, a result that further improve in the second case, with the inclusion of the NO₂ observations. The predictions of the trained KLMS network are shown in Figure 4.

In addition, Figure 4, which has a slightly different scale on the y-axis due to the normalization of the data itself, confirms the hypothesized outliers in data around sample 210. In fact, now the predicted value for PM₁₀ is less than the measured one, according to values of NO₂ and justifying the external nature of that sample. We can conclude that this and similar outliers are due to the sea salt.

Note that the resulting network is constituted of 700 nodes, corresponding to the 700 training examples, each of which is composed of 360 elements. In a real world application, for computational requirements, there would be a need of a method for actively curtailing the growth of the resulting system, such as a pruning algorithm.




Fig. 4. Predicted and measured data of trained KLMS network

We have also compared the proposed approach with a Multilayer Perceptron (MLP) [3], a Support Vector Machine [3,13] and a Nonlinear Autoregressive (NAR) model [14], with and without the inclusion of NO₂ correlated data. The MLP has one hidden layer with 50 neurons, while the learning rate is set to 0.2 for each neuron. Results obtained with these standard architectures are numerical equivalent to those of Figures 3 and 4 in terms of MSE and predicted data (for this reason are not reported in the paper). This fact suggests that the proposed approach is able to well perform the prediction of correlated environmental data and obtains similar results found in literature, but it presents the advantage of a lower computational cost, with respect MLP and SVM approaches.

4 Conclusions

We demonstrated how KAF algorithms provides a fast and efficient way of training a simple network to predict environmental data such as PM₁₀. Our system achieves a significant low error over a testing set with the inclusion of correlated data and shows a robust behavior with respect to outliers due to external factors. By iterated application, this system can become an handy tool for local legislators to accurately predict the evolution of particulate matter in their region, in the effort to strongly reduce the global concentration of PM₁₀ over the year. Although comparisons with other standard

techniques provide similar results, the proposed approach is characterized by a low computational cost.

References

1. Pope, C.A., Burnett, R., Thun, M.J., Calle, E.E., Krewskik, D., Ito, K., Thurston, G.: Lung cancer, cardiopulmonary mortality, and long term exposure to fine particulate air pollution. *Journal of the American Medical Association* 287, 1132–1141 (2002)
2. Hooyberghs, F., Mensink, C., Dumont, G., Fierens, F., Brasseur, O.: A neural network forecast for daily average PM10 concentrations in Belgium. *Atmospheric Environment* 39(8), 3279–3289 (2005)
3. Haykin, S.: *Neural Networks and Learning Machines*, 2nd edn. Pearson Publishing (2009)
4. Gardner, M.W., Dorling, S.R.: Artificial neural networks (the multilayer perceptron): a review of applications in the atmospheric sciences. *Atmospheric Environment* 32, 2627–2636 (1998)
5. Aramongsanuwat, S., Meesad, P.: Prediction of PM10 using Support Vector Regression. In: 2011 International Conference on Information and Electronics Engineering IPCSIT, vol. 6, pp. 120–124 (2011)
6. Liu, W., Principe, J.C., Haykin, S.: *Kernel Adaptive Filtering: A Comprehensive Introduction*. Wiley (2010)
7. Aurangojeb, M.: Relationship between PM₁₀, NO₂ and particle number concentration: validity of air quality controls. *Procedia Environmental Sciences* 6, 60–69 (2011)
8. Ito, K., Thurston, J.D., Nadas, A., Lippmann, M.: Monitor-to-monitor temporal correlation of air pollution and weather variables in the North-Central U.S. *Journal of Exposure Analysis and Environmental Epidemiology* 11(1), 21–32 (2001)
9. Lippmann, M., Ito, K., Nadas, A., Burnett, R.T.: Association of particulate matter components with daily mortality and morbidity in urban populations. *Journal of Exposure Analysis and Environmental Epidemiology* 9(5), 5–72 (2000)
10. Rakesh, K., Abba, E.J.: Air pollution concentrations of PM_{2.5}, PM₁₀ and NO₂ at ambient and kerbsite and their correlation in metro city - Mumbai. *Environmental Monitoring and Assessment* 119(1-3), 191–199 (2006)
11. Scholkopf, B.: The kernel trick for distances. In: XIII Proceedings of the 2000 Conference on Advances in Neural Information Processing Systems, vol. 13, pp. 301–307 (2000)
12. Liu, W., Pokharel, P.P., Principe, J.C.: The Kernel Least-Mean-Square algorithm. *IEEE Transactions on Signal Processing* 56(2), 543–554 (2008)
13. Vapnik, V.N.: *The Nature of Statistical Learning Theory*. Springer (1995)
14. Conover, W.J.: *Practical Nonparametric Statistics*. Wiley (1971)

A Collaborative Filter Approach to Adaptive Noise Cancellation

Michele Scarpiniti, Danilo Comminiello, Raffaele Parisi, and Aurelio Uncini

Department of Information Engineering, Electronics and Telecommunications (DIET),
“Sapienza” University of Rome, Via Eudossiana 18, 00184, Rome
{michele.scarpiniti,danilo.comminiello,
raffaele.paris}@uniroma1.it, aurel@ieee.org

Abstract. In this paper we propose a filter combination for the adaptive noise cancellation (ANC) problem in nonlinear environment. The architecture consists in a convex combination of two adaptive filters: a classical filter and a nonlinear filter based on Functional Links. While the convergence of the linear filter is very fast, the convergence of the nonlinear one might be slower, even if it provides a more accurate solution. The convex combination of both filters allows to reach good performances in terms of convergence and speed. In addition a variable step size is used in order to obtain better performance. Several experimental results, in different reverberant conditions, demonstrate the effectiveness of the proposed approach.

Keywords: Noise Cancellation, Adaptive Filters, Functional Link Network, Convex Combination, Variable Stepsize.

Introduction

Adaptive Noise Cancellation (ANC) is an evergreen application useful in a wide range of scenarios [1], such as conference rooms. Although ANC architectures are well known in literature and applications, some issues are still open. In fact, in real-time applications two main contrasting characteristics are very important: high convergence speed and good convergence performance. These characteristics can be controlled by the step-size used in adaptive algorithms. Unfortunately a small value of this parameter provides good convergence performance, i.e. small error for the obtained clean signal, but results in a slow convergence speed. In contrast a large value for the stepsize provides a fast convergence but poor quality. In order to meet these conflicting requirements, the stepsize needs to be controlled, for example, by sequentially scaling its value [2,3]. In addition, signals are often deteriorated by nonlinear distortions, due to low-cost audio equipment, such as amplifiers and loudspeakers, or to vibrating structures and chests. These distortions usually neutralize the work of linear ANC systems, making them useless in most applications.

Recently a convex combination filter approach was proposed [4,5,6]. This approach can simply solve the ambiguity on the choice of the stepsize. In fact, the adaptive filter results in a combination of two filters, the first with a large stepsize and the latter with

a small one. This solution overcomes the dichotomy between speed and convergence: while the second filter provides good convergence performances, the first one speeds up the convergence of the overall filter [5]. More recently such a combination was successfully applied to nonlinear acoustic echo cancellation problems [7,8].

In this paper we propose the use of this convex filter combination to implement an ANC architecture that can exploit both positive improvement of such an architecture: achieving good performance in a small amount of time and recovering the nonlinear distortion. In addition a variable stepsize is used in order to enhance the performance obtained by the algorithm [3]. These two adaptive filters collaborate toward the convergence: in this way we name this architecture as *collaborative* ANC (CANC).

Several experimental results conducted in an office environment varying the reverberation time T_{60} in a wide range, demonstrate the effectiveness of the proposed approach.

This paper is organized as follows: section 1 describes the proposed architecture, while section 2 shows some experimental results. Finally section 3 draws our conclusions.

1 The Proposed Architecture

The standard ANC architecture consists in an adaptive filter trained by a desired signal $d[n] = s_f[n] + r[n]$, where $s_f[n] = s[n] * h_s[n]$, $s[n]$ is the desired clean speech signal, $h_s[n]$ is the room impulse response (RIR) between the speech source and the microphone, and $r[n]$ is the background noise captured by the microphone. The filter input signal $x[n]$ is a version of the background noise not altered by the environmental effects: the microphone recording $x[n]$ must be placed close to the noise source $b[n]$, $x[n] = b[n]$. The contribution to the desired signal is $r[n] = b[n] * h_n[n]$, where $h_n[n]$ is the RIR between the noise source and the microphone. At convergence the adaptive filter coefficients must be an estimation of this impulse response. In addition we suppose that the signal received by the microphones is distorted, due to a non ideal behavior of the transducer and its electronic equipment.

The collaborative ANC, or CANC, consists in a convex combination of a linear adaptive filter (LAF) and a nonlinear filter, implemented by a Functional Link Network (FLN). In this way the proposed system is constituted by the architecture shown in Figure 1. A single FLN adaptive filter was already used for ANC in [9], obtaining promising results. In order to consider always present the path due to the linear filter, the combination is performed by summing the output of the linear filter with a scaled version of the output of the nonlinear filter:

$$y[n] = y_L[n] + y_{NL}[n] = y_L[n] + \lambda[n]y_{FL}[n], \quad (1)$$

where $0 \leq \lambda \leq 1$, $y_L[n]$ is the output of the linear filter, evaluated simply as

$$y_L[n] = \mathbf{w}_n^T \cdot \mathbf{x}_n, \quad (2)$$

where $\mathbf{w}_n = [w_1, \dots, w_L]^T$ is the vector of the L coefficients for the linear filter at instant n and $\mathbf{x}_n = [x[n], x[n-1], \dots, x[n-L+1]]^T$.

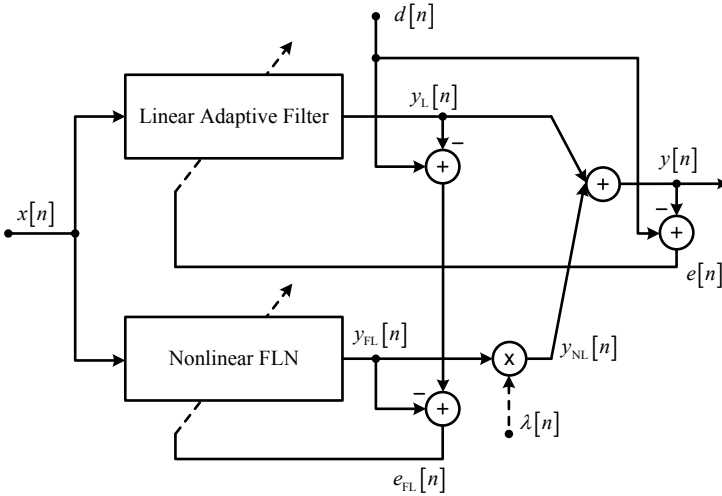


Fig. 1. The proposed collaborative combination of adaptive filters

In [10] the Functional Link Network (FLN) is described as an artificial neural network (ANN) with a single layer, that is able to process an input signal using a functional expansion block. This block consists of a series of linearly independent functions that accept the input signal as argument. The functional expansion projects the input in a higher dimension space.

The functional expansion used in [7] is a trigonometric series expansion composed of both linear and nonlinear elements. However, in order to implement a purely nonlinear kernel, we will consider that the enhanced pattern $\mathbf{z}[n]$ is composed just by some nonlinear transformations of the original variables:

$$\mathbf{z}[n] = [1, \mathbf{z}_0^T[n], \dots, \mathbf{z}_{L_{in}-1}^T[n]]^T, \quad (3)$$

where $\mathbf{z}_i[n] = [z_{i,0}[n], \dots, z_{i,Q-1}[n]]$, whose elements $z_{i,j}[n]$ assume the following values:

$$z_{i,j}[n] = \begin{cases} \sin(p\pi x[n-i]) & \text{for } j = 2p \\ \cos(p\pi x[n-i]) & \text{for } j = 2p + 1 \end{cases} \quad (4)$$

in which $p = 0, 1, \dots, P-1$, being P the expansion order and $x[n]$ represents the input sample at n -th instant, while $i = 0, 1, \dots, L_{in}-1$ is the input sample index and L_{in} is the length of the input signal block; finally $j = 0, 1, \dots, Q-1$ is the enhanced pattern index and $Q = 2P$ is the number of functional links for the i -th input. Therefore, the enhanced pattern is composed of $L_{en} = 2QL_{in} + 1$ elements.

It is now possible to write the output of the nonlinear FLAF simply at the instant n as the output of a FIR filter:

$$y_{NL}[n] = \mathbf{h}_n^T \cdot \mathbf{z}_n, \quad (5)$$

where $\mathbf{h}_n = [h_0, h_1, \dots, h_{L_{en}-1}]^T$.

1.1 Adaptation of the Proposed Architecture

Let us pose $e[n] = d[n] - y[n]$ the overall error of the proposed architecture. Then the adaptation of each filter is performed minimizing the overall square error $e^2[n]$ by using the classical *Normalized Least Mean Square* (NLMS) algorithm [11]:

$$\mathbf{w}_{n+1} = \mathbf{w}_n + \mu_L[n] \frac{e[n] \mathbf{x}_n}{\delta + \mathbf{x}_n^T \mathbf{x}_n}, \quad (6)$$

$$\mathbf{h}_{n+1} = \mathbf{h}_n + \mu_{NL}[n] \frac{e[n] \mathbf{x}_n}{\delta + \mathbf{z}_n^T \mathbf{z}_n}, \quad (7)$$

where δ is a regularization parameter.

In order to take into account the under-modeling scenario we adopt a variable step-size (VSS) [3]. The variable stepsize parameter $\mu_j[n]$ in (6) and (7), after M iterations, is chosen according to [12], :

$$\mu[n] = \begin{cases} \mu_f, & n \leq M \\ \left| 1 - \frac{\sqrt{|\hat{\sigma}_d^2[n] - \hat{\sigma}_y^2[n]|}}{\hat{\sigma}_e^2[n] + \xi} \right|, & n > M \end{cases}, \quad (8)$$

where ξ is a regularization parameter, and the general parameter $\hat{\sigma}_\theta^2[n]$ represents the power estimate of the sequence $\theta[n]$:

$$\hat{\sigma}_\theta^2[n] = \gamma \hat{\sigma}_\theta^2[n-1] + (1 - \gamma) \theta^2[n], \quad (9)$$

where $\gamma = 1 - 1/(kM)$ is a weighted factor with $k > 1$. The initial value is $\hat{\sigma}_\theta^2[0]$ and is set to 1.

In order to reduce the gradient noise and to keep the mixing parameter in the range $(0, 1)$, the adaptation of $\lambda[n]$ can be carried out through the adaptation of another parameter, $a[n]$, related to $\lambda[n]$ by the following equation:

$$\lambda[n] = \left(1 + e^{-a[n]} \right)^{-1}. \quad (10)$$

The update of $a[n]$ is then given by [8]:

$$a[n] = a[n-1] + \frac{\mu_a}{p[n]} \lambda[n] (1 - \lambda[n]) e[n] \varepsilon[n], \quad (11)$$

where $\varepsilon[n] = y_1[n] - y_2[n]$; the term $p[n] = \beta p[n-1] + (1 - \beta) \varepsilon^2[n]$ is the estimated power of $\varepsilon[n]$, and β is a threshold close to one [13].

The proposed architecture is robust to any nonlinearity level, since when the convolutive path is merely linear $\lambda[n]$ will converge towards 0 and the whole scheme will behave like a purely linear filter, thus avoiding any gradient noise from the nonlinear filter. On the other hand, when the convolutive path presents nonlinearities the mixing parameter will approach 1 according to the nonlinearity level in the path. Moreover, after convergence, the value of $\lambda[n]$ represents a measure of the degree of nonlinearity present in the system.

2 Experimental Results

The experimental tests are conducted in an office room of standard dimension of $6 \times 4 \times 3$ m. The location of microphones and sources is depicted in Figure 2. Different reverberation times are simulated in the range $0 \div 350$ ms, in order to test the performance of the proposed architecture versus the reverberation time. The room impulse responses are simulated using the Matlab tool RoomSim¹ [14]. The performance of the proposed algorithm is measured in terms of output SNR, defined by the following equation:

$$\text{SNR}_{out} = 10 \log_{10} \frac{E \{s_f^2[n]\}}{E \{(r[n] - y[n])^2\}}. \quad (12)$$

Eq. (12) is justified by the fact that

$$e[n] = d[n] - y[n] = (s[n] * h_s[n] + x[n] * h_n[n]) - y[n] = s_f[n] + (r[n] - y[n]),$$

where the first term in the right side is the desired clean speech signal, while the terms in round brackets is the residual noise, that converges to zero.




Fig. 2. The proposed experimental set-up

The filter lengths L depend on the reverberation time used, and it is listed in Table 1, while $L_{en} = 256$. The expansion order is set to $P = 2$. The signal test is a female speech, sampled at 8 kHz.

¹ Roomsim is a MATLAB simulation of shoe-box room acoustics for use in teaching and research. Roomsim is available from

<http://media.paisley.ac.uk/~campbell/Roomsim/>

Table 1. Reverberation time T_{60} and filter length L used in simulations

T_{60} [ms]	L
Anechoic	128
50	512
100	720
200	1024
350	2048

Table 2. Summary of output SNR [dB] for the proposed experimental tests with a female speech signal at different T_{60} and input SNR, using the NLMS algorithm and a mild distortion

SNR_{IN} [dB] \ T_{60} [ms]	Anechoic	50	100	200	350
10	25.33	20.13	18.93	16.49	14.16
5	23.78	20.27	17.42	15.45	12.33
0	20.88	17.72	16.36	13.99	10.91
-5	19.15	17.01	14.49	11.28	8.76
-10	18.13	15.32	11.41	7.81	5.89




Fig. 3. Comparison of output SNR versus T_{60} and input SNR for the NLMS algorithm in the case of a mild distortion

The first experimental test is conducted using the NLMS algorithm (6). The other parameters used are $\mu_L[0] = \mu_N[0] = 0.2$, $\mu_a = 5 * 10^{-2}$, $\delta = 10^{-6}$, $\beta = 0.9$, $\gamma = 0.9$, while the parameter $a[n]$ is initialized as $a[0] = 1$ and the vectors of filter parameters \mathbf{w}

Table 3. Summary of output SNR [dB] for the proposed experimental tests with a female speech signal at different T_{60} and input SNR, using the NLMS algorithm and a strong distortion

SNR_{IN} [dB] \ T_{60} [ms]	Anechoic	50	100	200	350
10	24.30	20.09	18.63	15.51	12.88
5	22.73	19.88	17.22	14.82	11.43
0	19.79	17.27	15.35	13.19	10.87
-5	17.58	15.95	13.45	10.16	8.24
-10	16.49	13.69	10.15	7.66	4.18




Fig. 4. Comparison of output SNR versus T_{60} and input SNR for the NLMS algorithm in the case of a strong distortion

and \mathbf{h} are initialized to $\delta[n]$. A mild loudspeaker distortion is also considered, simulated as $d_{NL}[n] = \tanh(\beta d[n])$ and using $\beta = 0.5$. As background noise we use a white Gaussian noise (WGN). The power of the WGN is adjusted in order to have five different input SNR, precisely 10, 5, 0, -5 and -10 dB. Results for the considered test in all reverberant environments are summarized in Table 2. These results are averaged over 100 trials. Results presented in Table 2 are also depicted in Figure 3 for a more evident interpretation. As we can see from Figure 3 the performance in terms of output SNR is decreasing by increasing the reverberation time, due to the longer length of adaptive filters. This effect is more evident for high input SNR.

A second experimental test is conducted in the same simulated environment as the previous test, but using a stronger distortion, setting $\beta = 4$. The architecture parameters are set as in the previous case. Results for this simulations, averaged over 100 trials, are summarized in Table 3. These results are also depicted in the Figure 4. This figure




Fig. 5. Comparison of output SNR for two different input SNR at $T_{60} = 100$ ms in a changing environment

confirms that performances are decreasing by increasing the reverberation time and this trend is constant by decreasing the input SNR.

In a third experimental test, we use a changing environment. In the first half of the experiment the scenario is purely linear, then a strong distortion (with $\beta = 4$) is introduced. The reverberation time of this simulation is set to $T_{60} = 100$ ms and the architecture is run for an input SNR of 10 dB and -10 dB, respectively.

The output SNRs are evaluated averaging results over 250 samples. These results are depicted in the Figure 5. The profiles of graphics in Figure 5 show the same trend as the previous simulations, but evidence that the performance falls down when the scenario is abruptly changed, while it quickly returns to steady-state level. This fact confirms that we can take a great advantage from using a combination of a linear and nonlinear algorithms in the proposed collaborative architecture.

These simulations show that the proposed architecture is able to reach a good noise suppression in terms of SNR, even in the case of a strong nonlinear distortion and a changing environment.

3 Conclusions

In this paper an adaptive filter combination is proposed for the solution of Adaptive Noise Cancellation in nonlinear environment. The architecture described in this work is able to solve this hard problem in reverberant environment, as demonstrated by several experimental results. In particular the filter adaptation is conducted by the NLMS algorithm. It is shown that we can take advantage from the convex combination in the

proposed collaborative architecture, reaching good performance in terms of both convergence and speed, and nonlinear distortion compensation.

References

1. Widrow, B., Glover, J.R., McCool, J.M., Kaunitz, J., Williams, C.S., Hearn, R.H., Zeidler, J.R., Dong, E.J., Goodlin, R.C.: Adaptive noise cancellation: Principles and applications. *Proceedings of the IEEE* 63(12), 1692–1716 (1975)
2. Ang, W.P., Farhang-Boroujeny, B.: A new class of gradient adaptive step-size LMS algorithms. *IEEE Transactions on Signal Processing* 49(4), 805–810 (2001)
3. Hwang, J.K., Li, Y.P.: Variable step-size LMS algorithm with a gradient-based weighted average. *IEEE Transactions on Signal Processing* 16(12), 1043–1046 (2009)
4. Arenas-García, J., Figueiras-Vidal, A.R.: Adaptive combination of normalised filters for robust system identification. *Electronics Letters* 41(15), 874–875 (2005)
5. Arenas-García, J., Figueiras-Vidal, A.R., Sayed, A.H.: Mean-square performance of a convex combination of two adaptive filters. *IEEE Transactions on Signal Processing* 54(3), 1078–1090 (2006)
6. Arenas-García, J., Gómez-Verdejo, V., Figueiras-Vidal, A.R.: New algorithms for improved adaptive convex combination of LMS transversal filters. *IEEE Transactions on Instrumentation and Measurement* 54(6), 2239–2249 (2005)
7. Communiello, D., Scarpiniti, M., Parisi, R., Uncini, A.: A functional link based nonlinear echo canceller exploiting sparsity. In: *Proc. of International Workshop on Acoustic Echo and Noise Control (IWAENC 2010)*, Tel Aviv, Israel, August 30–September 2 (2010)
8. Communiello, D., Azpicueta-Ruiz, L.A., Scarpiniti, M., Uncini, A., Arenas-García, J.: Functional link based architectures for nonlinear acoustic echo cancellation. In: *Proc. of Hands-free Speech Communication and Microphone Arrays (HSCMA 2011)*, Edinburgh, UK, May 30–June 1, pp. 180–184 (2011)
9. Sicuranza, G.L., Carini, A.: A generalized FLANN filter for nonlinear active noise control. *IEEE Transactions on Audio, Speech, and Language Processing* 19(8), 2412–2417 (2011)
10. Pao, Y.H.: *Adaptive Pattern Recognition and Neural Networks*. Addison-Wesley (1989)
11. Haykin, S.: *Adaptive Filter Theory*, 4th edn. Prentice-Hall (2001)
12. Paleologu, C., Ciochină, S., Benesty, J.: Variable step-size NMLS algorithm for under-modeling acoustic echo cancellation. *IEEE Signal Processing Letters* 15, 5–8 (2008)
13. Azpicueta-Ruiz, L.A., Zeller, M., Figueiras-Vidal, A.R., Arenas-García, J., Kellermann, W.: Adaptive combination of volterra kernels and its application to nonlinear acoustic echo cancellation. *IEEE Transactions on Audio, Speech, and Language Processing* 19(1), 97–110 (2011)
14. Campbell, D., Palomaki, K., Brown, G.: A Matlab simulation of shoebox room acoustics for use in research and teaching. *Computing and Information Systems* 9(3), 48 (2005)

Waveform Variation of the Explosion-Quakes as a Function of the Eruptive Activity at Stromboli Volcano

Antonietta M. Esposito, Luca D'Auria, Flora Giudicepietro, and Marcello Martini

Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Napoli Osservatorio Vesuviano,
Napoli, Italy

{antonietta.esposito,luca.dauria,flora.giudicepietro,
marcello.martini}@ov.ingv.it

Abstract. In the period from June to September 2011, the Stromboli volcano was affected by an activity characterized by an increase of the volcanic tremor amplitude, in the magnitude of explosions and with some lava overflows. In order to examine and understand in more detail this particular phase of the volcano, we present here an unsupervised investigation of the waveform variation of the explosion-quakes recorded during this period. The aim is to identify a possible relationship between the temporal changes of these events and the volcano seismic activity. The analysis is performed on a dataset of about 8400 explosion-quakes by using a SOM neural network. This technique works well with large datasets allowing to find out unpredicted characteristics among them. The SOM clustering highlights sudden changes occurring at the end of July and of August and a permanent variation between June and September reflecting a modification in the volcano activity. These results could be interesting for focusing the analysis of the seismological dataset in these intervals in order to evidence minor, but important variations, which were previously undetected and to improve the knowledge on the explosive dynamics of the volcano.

Keywords: Explosion-quakes, SOM neural network, unsupervised clustering, volcano dynamics.

1 Introduction

The Stromboli volcano, located in the Tyrrhenian Sea, has a persistent but usually moderate volcanic activity, called Strombolian, occurring at the vents on the top of the volcanic edifice. Its monitoring is continuously carried out by a dense network of 13 broadband stations [2], installed on the island by the Istituto Nazionale di Geofisica e Vulcanologia (INGV) after the eruptive crisis in December 2002 [1].

Its seismicity is characterized by volcanic tremor and explosion-quakes. Volcanic tremor is a continuous seismic signal, recorded on active volcanoes, usually related to moderate persistent degassing of the volcano. Explosion-quakes, on the other hand, are transient signals related to discrete explosions.

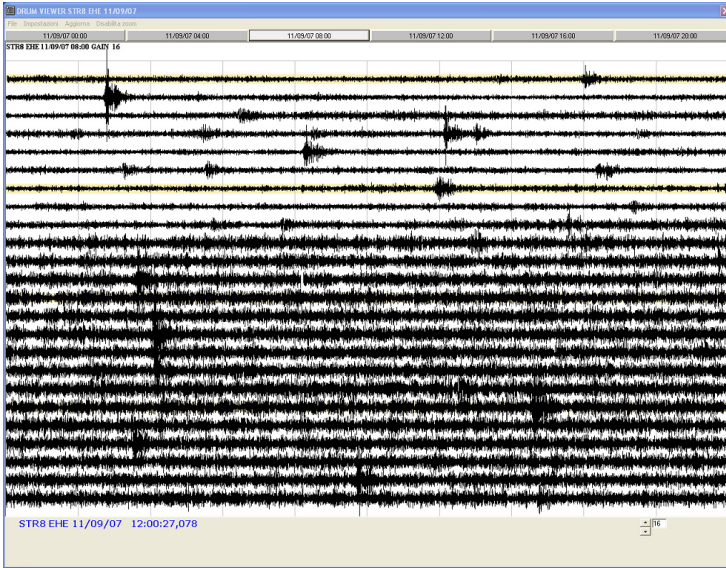


Fig. 1. An example of increase, from low to high values, of the volcanic tremor amplitude recorded on September 7, 2011 by the STR8 station. The image shows a window of 4 hours of recording where each line is a 10-minute signal.




Fig. 2. An example of powerful explosion-quakes recorded by the STR8 station on the August 25, 2011. The picture illustrates a window of 4 hours of recording where each line is a 10-minute signal.

During June-September, 2011 Stromboli has shown an unusual behavior characterized by an increase in the volcanic activity with an increase in the volcanic tremor amplitude (Figure 1) and in the magnitude of the explosion-quakes (Figure 2). Moreover, two high frequency signals, related to the formation of small lava flows from the NE crater of the volcano [6], were recorded: the first one at 20:52 UTC on August 1, with duration of about 30 minutes, and the second one at 04:14 UTC on August 2, which lasted about two hours.

Figure 3 shows images of a relatively strong explosion and of the lava flow along the “Sciara del Fuoco” flank of the volcano occurred on August 2, 2011. They are taken by the thermal and visible cameras used for the volcano monitoring and located at 400m above the sea level, north of the craters area.




Fig. 3. Frames from the thermal and visible cameras showing a strong explosion and the lava flow along the “Sciara del Fuoco” occurred on August 2, 2011 (courtesy by INGV Catania). The color bar indicates the temperature.

In order to better understand this particular phase of the volcano, we present here an unsupervised investigation of a dataset of explosion-quakes recorded during this period. The analysis was performed using the SOM neural network to highlight a possible relationship between the variation in time of the waveforms of these events and the explosive dynamics of the volcano.

The results show that the waveform changes are clearly visible and interpretable on the SOM map and that there exists a good association between the clustering in time of the events and the positions on the map.

In the following, the dataset of the explosion-quakes is described, the applied clustering technique is illustrated, and finally the SOM results are discussed.

2 Dataset Overview

Our dataset consists of 8409 explosion-quakes recorded by the vertical component of a seismic station (STR6, located 500m south of the craters) and sampled at 50 Hz. This station is located not far from the craters and operated almost continuously during the analyzed period with a good quality of the recorded signals.

An explosion-quake is a seismic signal generated by an explosion, with frequency content mostly in the range 1-6 Hz and a complex and highly variable waveform. These signals usually exhibit no distinct seismic phases (Figure 4). The typical Strombolian activity consists in 10-13 individual explosions per hour emitting gases and pyroclastic fragments. Instead, in occasional episodes of increased activity the occurrence rate of these events can reach a value of 20-30 events per hour.

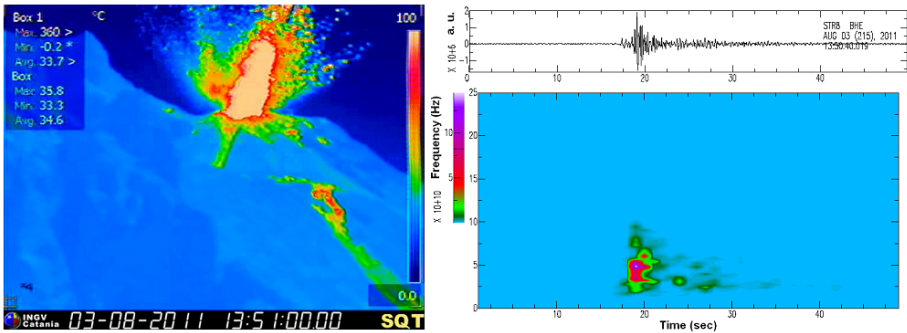


Fig. 4. A frame of a strong explosion (on the left) recorded by the thermal camera (Courtesy by INGV Catania) on August 3, 2011 and the associated waveform (on the top) and spectrogram (on the bottom). The color bar indicates the temperature.

In our analysis the signals have been selected on the daily seismograms using a manual picking procedure [9]. A signal window of 4 sec (200 samples) (Figure 5) has been used to represent each event since, this interval, includes the onset of the signal and allows its discrimination. Finally, the amplitude has been normalized respect to the Root Mean Square (RMS).




Fig. 5. An explosion-quake signal of 1200 samples (20s) (blue) and its corresponding vector of 200 samples (4s) (red) considered for the analysis

3 The SOM Technique

When a large data set is available and no information about which similarity measure is best suited to group them, then it may be useful to apply an unsupervised technique [8]. Classic methods of cluster analysis often make assumptions regarding linearity, normal distribution or inherent clustering tendencies in the data. For these reasons they can fail for data with complex structure.

In this work we adopted an unsupervised SOM (Self-Organizing Map) network [10, 11] to perform the waveforms clustering of the explosion-quakes. This method has already been used for the analysis of seismic data recorded at Stromboli [3, 4, 5, 7].

Contrary to the classic methods, the SOM technique provides easy visualization, it is not critically dependent on its parameters and is able to handle large datasets to detect isolate patterns and structures in the data. Through a process called self-organization, the SOM algorithm maps the high-dimensional input space into a bi-dimensional grid of processing units to each of which has an associated prototype vector. The learning is competitive: this means that the algorithm identifies at each iteration and for each input the winning node, i.e. the unit whose connection weights are the closest to the input vector in terms of Euclidean distance. Then it updates its weights and those of the adjacent nodes (mathematical details in [12]). In this way, not only the winner but also its whole neighborhood is moved closer to the input pattern. The resulting map preserves the topological properties of the original data: points that are near each other in the input space are mapped to nearby map units in the SOM. Thus, the SOM can be used as a clustering tool as well as a visualization tool for high-dimensional data.

4 Results

In this work we used a SOM map with 84 (12x7) nodes, a local hexagonal structure and a global toroid shape. In Figure 6 (on the left) we show the SOM map obtained for the examined dataset. Individual nodes are visualized as yellow hexagons. The toroidal topology is represented as a sheet in order to have a more direct representation of the clusters organization. The size of the map nodes specifies the number of signals which fall in that node (i.e. the data density). The gray hexagons describe the Euclidean distances between the nodes using a gray level scale [10, 11]. On the right of Figure 6 are illustrated the prototype vectors associated with each node. Since the analysis is realized by using information on the waveform of the explosion-quakes, then the SOM prototypes preserve this information providing a simple visualization of the clustering and consequently of the results.




Fig. 6. The toroid SOM map with $12 \times 7 = 84$ nodes (on the left) and the associated prototype vectors (on the right)

In Figure 7 we show the temporal distribution of the events on the SOM map (top panels) and the occurrence time of the events in the 84 SOM clusters (graph on the bottom). In the first graph, the SOM maps of different colors indicate the signals belonging to different months of the analyzed period: red for June, green for July, blue for August and cyan for September. In the second graph, the red lines indicate the beginning and the end of the examined time interval; the orange line identifies the lava flows; while the blue markers indicate changes in the clusters distribution which could correspond to variations in the volcano activity.

From this last graph, it is possible to make three important observations: first, the beginning and the end of the unusual activity are marked by a change in the clusters distribution; second, also the main variations occurred during the considered period are clearly identifiable as variations in the clusters; finally, between June and September, i.e. the pre and post anomalous activity, a persistent modification happens that corresponds to a persistent change in the dynamics of the volcano. However, both the graphs show that there is a good correspondence between the positions on the map and the clustering in time of the events, i.e. the clustering of the map matches appropriately a clustered temporal distribution of the events.

To reinforce this concept, Figure 8 demonstrates how different areas of the map correspond to different clusters in time and so different phases of the unusual activity. The colored circles on the SOM map (on the left) highlight different families of signals which are distributed differently over time (on the right) meaning that different physical processes of the volcano are involved.

We conclude that, the waveform variations of the explosion-quakes are clearly visible and interpretable on the SOM map.

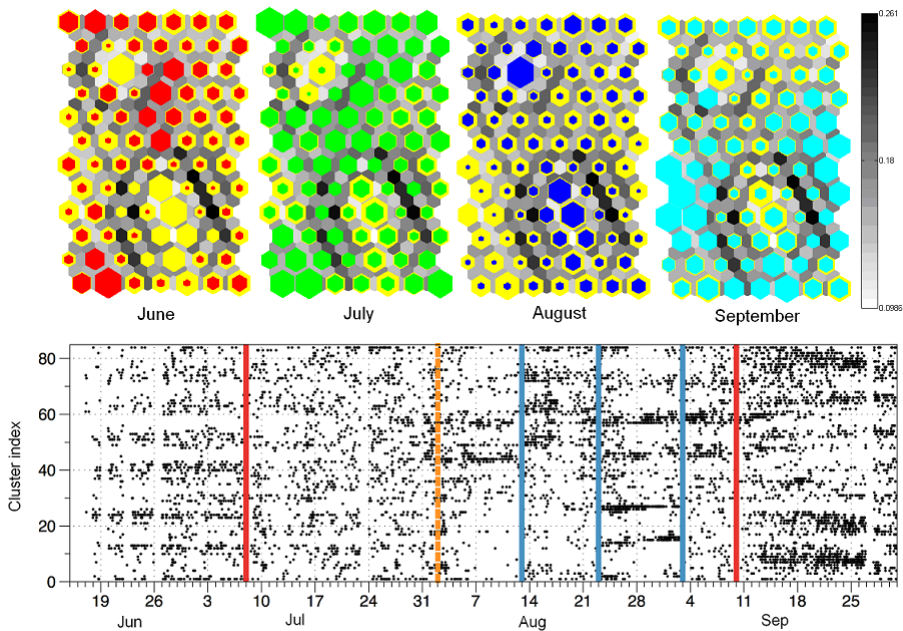


Fig. 7. The temporal distribution of the events on the SOM map (upper graph) and the occurrence time of the events in the 84 SOM clusters (lower graph) are visualized. In the first graph, on the map in yellow, related to the entire selected period, are plotted the different analyzed months by using different colors. In the second graph, the red lines specify the beginning and the end of the unusual period, the orange line indicates the lava flows, while the blue markers point out changes in the cluster distribution which could correspond to variations in the volcano activity.




Fig. 8. The SOM map (on the left) and the occurrence time of the events in the 84 SOM clusters (on the right) are displayed. The colored circles on the map identify regions to which correspond clusters in time.

5 Conclusions

This work presents an unsupervised investigation of the waveforms of a large dataset (8409 recordings) of explosion-quakes recorded at Stromboli during June-September 2011, when a phase of anomalous activity was ongoing.

The analysis was realized using a SOM neural network. This method works well with large dataset, does not depend on its parameters and is particularly useful when it is desirable to organize and present complex data in a form which is easier to understand. Moreover, being the analysis based on the waveforms of the signals, this allows to better visualize where the separation between the clusters is and which its nature is.

In our case, this technique has proved to be able to identify and clearly displaying the waveforms variations of the analyzed signals. Thus, from the graphs of the Figures 7 and 8, it is possible to observe that, tracking the change in the distribution of events on the map, is indicative of variations in the explosive dynamics. Changes in the shape of the waveforms of the explosion-quakes correspond to changes in the positions on the map.

A significant result is that a persistent change is visible between June and September, i.e. before and after the analyzed period. This means that the explosion-quakes have no longer the same waveforms as before. Probably, during this interval, there has been a persistent variation in the volcanic system that appears clear with the SOM map representation.

However, further studies are needed on the events belonging to the main areas of the SOM map and the corresponding active vents to correlate them with volcanological phenomena in order to improve the knowledge on the explosive dynamics of the volcano.

References

1. Calvari, S., Spampinato, L., Lodato, L., Harris, A.J.L., Patrick, M.R., Dehn, J., Burton, M.R., Andronico, D.: Chronology and complex volcanic processes during the 2002-2003 flank eruption at Stromboli volcano (Italy) reconstructed from direct observations and surveys with a handheld thermal camera. *J. Geoph. Res.* 110, B02201 (2005), doi:10.1029/2004JB003129
2. De Cesare, W., Orazi, M., Peluso, R., Scarpato, G., Caputo, A., D'Auria, L., Giudicepietro, F., Martini, M., Buonocunto, C., Capello, M., Esposito, A.M.: The broadband seismic network of Stromboli volcano, Italy, *Seismol. Res. Lett.* 80, 435–439 (2009), doi:10.1785/gssrl.80.3.435
3. Esposito, A.M., Scarpetta, S., Giudicepietro, F., Masiello, S., Pugliese, L., Esposito, A.: Nonlinear Exploratory Data Analysis Applied to Seismic Signals. In: Apolloni, B., Marinaro, M., Nicosia, G., Tagliaferri, R. (eds.) WIRN 2005 and NAIS 2005. LNCS, vol. 3931, pp. 70–77. Springer, Heidelberg (2006)
4. Esposito, A.M., Giudicepietro, F., Scarpetta, S., D'Auria, L., Marinaro, M., Martini, M.: Automatic discrimination among landslide, explosion-quake and microtremor seismic signals at Stromboli volcano using Neural Networks. *Bull Seismol. Soc. Am (BSSA)* 96(4), 1230–1240 (2006), doi:10.1785/0120050097
5. Esposito, A.M., Giudicepietro, F., D'Auria, L., Scarpetta, S., Martini, M., Coltelli, M., Marinaro, M.: Unsupervised neural analysis of very-long-period events at Stromboli volcano using the self-organizing maps. *Bull Seismol. Soc. Am (BSSA)* 98(4) 2449–2459 (2008), doi:10.1785/0120070110, ISSN: 0037-1106
6. Esposito, A.M., D'Auria, L., Giudicepietro, F., Peluso, R., Martini, M.: Landslides as precursors of the 2007 effusive eruption of Stromboli volcano: a neural network based detector. Submitted to the *Journal of Volcanology and Geothermal Research* (2011)
7. Esposito, A.M., D'Auria, L., Giudicepietro, F., Longobardi, M., Martini, M.: Clustering of Hybrid Events at Stromboli Volcano (Italy). In: Apolloni, B., et al. (eds.) *Frontiers in Artificial Intelligence and Applications, Neural Nets WIRN 2011 - Proceedings of the 21th Italian Workshop on Neural Nets*, vol. 234. IOS Press (2011) ISBN: 978-1-60750-971-4
8. Everitt, B., Landau, S., Leese, M.: *Cluster Analysis*. Oxford University Press, New York (2001)
9. Giudicepietro, F., Lo Bascio, D., Caputo, T., D'Auria, L.: Drum_Pkev: un programma per la costituzione di cataloghi supervisionati di eventi sismici ad elevata frequenza di accadimento, *Rapporti Tecnici NGV*, Numero 124 (2010)
10. Kohonen, T.: *Self Organizing Maps*. Springer (1995)
11. Kohonen, T., Hynninen, J., Kangas, J., Laaksonen, J.: SOM_PAK: The self-organizing map program package, Report A31, Helsinki University of Technology, Laboratory of Computer and Information Science, Espoo, Finland (1996), http://www.cis.hut.fi/research/som_lvq_pak.shtml
12. Kohonen, T.: *Self-Organizing Maps*, Series in Information Sciences 30, 2nd edn. Springer, Heidelberg (1997)

Artificial Neural Network (ANN) Morphological Classification of Magnetic Resonance Imaging in Multiple Sclerosis

Alessia Bramanti^{1,2}, Lilla Bonanno², Placido Bramanti², and Pietro Lanzafame²

¹ Faculty of Engineering, University of Messina

² IRCCS Centro Neurolesi "Bonino-Pulejo"

alessia.bramanti@gmail.com

Abstract. Multiple Sclerosis (MS) is an autoimmune condition in which the immune system attacks the Central Nervous System. Magnetic Resonance Imaging (MRI) is today a crucial tool for diagnosis of MS by allowing in-vivo detection of lesions. New lesions may represent new inflammation; they may increase in size during acute phase to contract later while the disease severity is reduced. To monitor evolution in time of lesions and to correlate this to MS phases, we focused on the application of Artificial Neural Network (ANN) based classification of MS lesions. An euclidean distance histogram, representing the distribution of edge inter-pixel distances, is used as input. In this work, we have extended the study already published, increasing to 21 the number of images. We can observe that the percentage of correct results on 21 images (93.81%) increased if compared to the study performed on 13 images (92.31%). This methodology could be used to monitor evolution in time of lesions of each patient and to correlate this to MS phases (i.e. to know if the lesions change their form).

Keywords: Multiple Sclerosis, Magnetic Resonance Imaging, Artificial Neural Network based classification, Euclidean Distance Histogram.

1 Introduction

Multiple Sclerosis (MS) is an autoimmune condition in which the immune system attacks the Central Nervous System (CNS) [1].

Magnetic Resonance Imaging (MRI) has become the most sensitive paraclinical test in diagnosis, assessment of disease evolution and treatment of the effects in MS. MRI is used as a prognostic tool at the first presentation of symptoms, suspicious of brain demyelination. Multiple hyperintense lesions on T2-weighted sequences are the characteristic MR appearance of MS. The majority of lesions are small, although, can occasionally measure several centimeters in diameter. MS lesions are usually small with intermediate high signal intensity with less severe degree of inflammation [2]. MS lesions tend to have an ovoid configuration with the major axis perpendicular to the ventricular borders (Dawson's fingers) (fig. 1).



Fig. 1. Axial proton density (PD), T2-weighted and Fluid Attenuated Inversion Recovery (FLAIR) images of a patient with MS demonstrate multiple hyperintense lesions with periventricular predominance

Most lesions, especially in the early stages of the disease, are evident on conventional MRI but diffuse irregular hyperintensities have also been demonstrated in the later stages of the disease. These areas with poorly defined borders, are usually seen around the ventricles and called dirty appearing white matter (DAWM).

2 Materials and Methods

The images have been acquired at IRCCS Centro Neurolesi "Bonino Pulejo" of Messina in DICOM format and transformed in png format. The cropped images, containing only a lesion, were used. These smaller images are the input of the algorithm (example in fig. 2). We used MRI images from the MS patients which were stored in a database to be read by the algorithm automatically and sequentially [3]. A preliminary study was published in Communication to SIMAI Congress, "Artificial Neural Network (ANN) Morphological Classification by Euclidean Distance Histograms for Prognostic Evaluation of Magnetic Resonance Imaging in Multiple Sclerosis", when we considered 13 image [4].

In this work, we have extended the work already published, increasing to 21 the number of images.




Fig. 2. Original image

We used a Multilayer Perceptron network (fig. 3). This simple type of network is interesting because the hidden units are free to construct their own representations of the input. The weights between the input and hidden units determine when each hidden unit is active, and so by modifying these weights, a hidden unit can choose what it represents [5].

We considered multi-layer architectures because units are often numbered by layer, instead of following a global numbering.




Fig. 3. Multilayer Perceptron

For the construction of training set, it must be obtained the contour of the lesion from the corresponding cropped image, through the following [6]:

- applying a thresholding global operation obtaining a segmented binary image;
- realizing uniform the internal of the lesion;
- applying to the obtained image the Laplacian operator, which can be implemented as the filter

$$H = \begin{vmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{vmatrix}$$

This formulation of Laplacian operator is called the 4-neighbours and allows us to obtain as a result the contour of the lesion [7] (fig. 4).

Once obtained the contour, we can calculate found edge inter-pixel distances and construct a histogram have been normalized to ten values which represents the number of occurrences of distances (y axis in the histogram) that have a certain value (axis x in the histogram) (fig. 5).




Fig. 4. (a) Segmented binary image; (b) Binary image after morphological operation; (c) Contour of lesion

The histogram have been normalized to ten values by the following formula (1):

$$dI = \frac{D_e \cdot d_{i_{max}}}{d_{max}} \quad (1)$$

where $d_{i_{max}} = 10$ is the values number of the normalized histogram, $d_{max} = 35$ is the maximum distance of the relative contour, D_e is the euclidean distance and dI is the correspondent index to the calculate distance.




Fig. 5. Histogram normalize to ten values

The examples that the administrator must provide are the values, normalized between 0 and 1, of the built histogram. All examples chosen for the training phase of the network are included in the training set, which is a file of "training". For each lesion the euclidean distance histogram of its contour is calculated, by the following formula (2).

$$D_e(p, q) = \sqrt{[(x - s)^2 + (y - t)^2]} \quad (2)$$

where D_e is the euclidean distance, $p(x, y)$ and $q(s, t)$ are two pixels of the contour. These values are stored on a file. Each row represents data of a contour and it is composed by thirteen value: in the first ten columns there are values of the obtained histogram and in the last three columns values that the network should produce. Since we consider three classes of contours, that of type *irregular*, *rounded* and *elongated*, these values are respectively 1 0 0 for values belonging to the first class (*irregular contours*), 0 1 0 for second class (*rounded contours*) and 0 0 1 for the third class (*elongated contours*).

2.1 Contour Recognition

Giving the neuronal network, an image containing a lesion of MS as input, produces output results. Through the study of output values, depending on how the neural network were trained, it can be established if the examined lesion belongs to a particular class. For example, providing to the network the image of a lesion, it can be got as output values [0.00, 0.26, 8.91].

By establishing a threshold this vector can be rounded to $[0, 0, 1]$. This means that the network has recognized the lesion, classifying it as lengthened contour.

3 Results

The neural network was trained using the data of seven pictures for each class of contour, while for verification have used data of seven images for the classes *elongated*, *rounded* and *irregular*. The figure (fig. 6- 8) shows some examples of training data.




Fig. 6. (a) Irregular contour; (b) Normalised euclidean distances histogram




Fig. 7. (a) Rounded contour; (b) Normalised euclidean distances histogram




Fig. 8. (a) Elongated contour; (b) Normalised euclidean distances histogram

From the table 1, we see that the results are correct at 93.81% and positive falses are 6.19%.

Table 1. The results obtained

Image	Membership class	Output	Membership class output
n.1	Elongated	0.00 – 0.26 – 0.91	Elongated
n.2	Elongated	0.00 – 0.05 – 1.00	Elongated
n.3	Elongated	0.00 – 0.05 – 1.00	Elongated
n.4	Elongated	0.00 – 0.33 – 0.90	Elongated
n.5	Elongated	0.00 – 0.28– 0.96	Elongated
n.6	Elongated	0.00 – 0.30 – 1.00	Elongated
n.7	Elongated	0.00 – 0.33 – 0.90	Elongated
n.8	Rounded	0.20 – 0.12 – 0.92	Uncertain
n.9	Rounded	0.17 – 0.71 – 0.05	Rounded
n.10	Rounded	0.21 – 0.74 – 0.03	Rounded
n.11	Rounded	0.22 – 0.74 – 0.03	Rounded
n.12	Rounded	0.18 – 0.70 – 0.04	Rounded
n.13	Rounded	0.20 – 0.74 – 0.03	Rounded
n.14	Rounded	0.19 – 0.75 – 0.05	Rounded

Table 2. (continued)

n.15	Irregular	0.75 – 0.38 – 0.00	Irregular
n.16	Irregular	0.96 – 0.00 – 0.25	Irregular
n.17	Irregular	0.94 – 0.00 – 0.35	Irregular
n.18	Irregular	0.73 – 0.40 – 0.01	Irregular
n.19	Irregular	0.85 – 0.00 – 0.19	Irregular
n.20	Irregular	0.77 – 0.00 – 0.25	Irregular
n.21	Irregular	0.90 – 0.30 – 0.21	Irregular

4 Conclusions

The tests conducted during the verification of the classification algorithm by neural network have produced positive results in all cases: the neural network is capable to distinguish the different contours analysed, highlighting their class membership. MS lesions obtained by MRI images are simply an example of a possible application. In this work, we have extended the study already published, increasing to 21 the number of images. We can observe that the percentage of correct results on 21 images (93.81%) increased if compared to the study performed on 13 images (92.31%).

The set of made checks showed that the application meets the targets proposed and how it could be a useful support for the neurologist. This methodology could be used to monitor evolution in time of lesions of each patient and to correlate this to MS phases (i.e. to know if the lesions change their form).

References

1. Sahraian, M.A., Shakouri Rad, A., Motamedi, M., Pakdaman, H., Radue, E.H.: Magnetic Resonance Imaging abnormalities in multiple sclerosis: a review Iran. J. Radiol (Summer 2007)
2. Anastasi, G., Bramanti, P., Di Bella, P., et al.: Volume rendering based on magnetic resonance imaging: advances in understanding the three-dimensional anatomy of the human knee. J. Anat. 211(3), 399–406 (2007)
3. Kanellopoulos, I., Wilkinson, G.G.: Strategies and best practice for neural network image classification. International Journal of Remote Sensing 18(4), 711–725 (1997)
4. Celona, A., Lanzafame, P., Bonanno, L., Marino, S., Spanò, B., Grasso, G., Puccio, L., Bramanti, P.: Artificial Neural Network (ANN) morphological classification by euclidean distance histograms for prognostic evaluation of magnetic resonance imaging in multiple sclerosis. In: Communications To Simai Congress, vol. 3, pp. 283.1–283.9 (2009)
5. Smeraldi, G.: Introduzione alle reti neurali. Technical report (1996)
6. Pratt, W.K.: Digital image processing. PIKS Inside (2001)
7. Daugman, J.G.: Complete discrete 2-D Gabor transforms by neural networks for image analysis and compression, Acoustics, Speech and Signal Processing. IEEE Transactions on Acoustics, Speech and Signal Processing 36(7), 1169–1179 (1988)

Neural Moving Object Detection by Pan-Tilt-Zoom Cameras

Alessio Ferone^{1,*}, Lucia Maddalena², and Alfredo Petrosino¹

¹ Department of Applied Science
University of Naples Parthenope, Italy

{alessio.ferone,alfredo.petrosino}@uniparthenope.it

² Institute for High-Performance Computing and Networking
National Research Council, Italy
lucia.maddalena@cnr.it

Abstract. Automated video surveillance using video analysis and understanding technology has become an important research topic in the area of computer vision. Most cameras used in surveillance are fixed, allowing to only look at one specific view of the surveilled area. Recently, the progress in sensor technologies is leading to a growing dissemination of Pan-Tilt-Zoom (PTZ) cameras, that can dynamically modify their field of view. Since PTZ cameras are mainly used for object detection and tracking, it is important to extract moving object regions from images taken with this type of camera. However, this is a challenging task because of the dynamic background caused by camera motion.

After reviewing background subtraction-based approaches to moving object detection in image sequences taken from PTZ cameras, we present a neural-based background subtraction approach where the background model automatically adapts in a self-organizing way to changes in the scene background. Experiments conducted on real image sequences demonstrate the effectiveness of the presented approach.

Keywords: Visual Surveillance, Motion Detection, Background Subtraction, Self Organization, Artificial Neural Network, PTZ Camera.

1 Introduction

Moving object detection is an important research problem in the field of video surveillance because it provides a focus of attention for recognition, classification, and activity analysis, allowing the analysis only of moving pixels [1].

Background subtraction is one of the most common approaches to moving object detection by static cameras (see surveys in [2,3,4,5]). It consists in constructing and updating a model of the fixed background, and detecting moving objects as those that do not belong to the model. Compared to other approaches, such as optical flow, no assumptions about the velocity of the object are made

* Corresponding author.

and it does not suffer of the foreground aperture problem. Moreover, background subtraction approaches are usually computationally affordable for real world applications. Nevertheless, the background subtraction approach is highly sensitive to dynamic scene changes due to lighting and extraneous events, with the consequent need for a suitable adaptation of the background model [6,7].

In order to overcome the limitations of static cameras, in the recent years the progress in sensor technologies has lead to a growing adoption of PTZ cameras, that can change their field of view through the use of panning, tilting, and zooming (i.e., moving left and right, up and down, closer and farther away), thus enabling to focus the attention on automatically selected areas of interest [8,9,10,11]. However this type of camera has introduced new challenges, because even pixels belonging to static objects appear to move in the camera frame. This effect is known as *ego-motion* and its estimation and compensation represents one of the main objectives of the research in the active video area [10].

Extensive research on moving object detection by PTZ cameras relies on background subtraction, where the scene background is obtained either by a mosaic image of the background or by ego-motion compensation of the background (see section 2). In this paper we embed an ego-motion compensation technique into a neural-based background subtraction approach to moving object detection. The background model automatically adapts in a self-organizing way to changes in the scene background. Background variations arising in a usual stationary camera setting are accurately handled by the neural background model originally proposed for this type of setting [12], while handling of variations due to the PTZ camera movement is ensured by having the neural background model to automatically compensate the eventual ego-motion, estimated at each time instant.

The paper is organized as follows. In section 2 we review the literature concerning background subtraction techniques applied to PTZ cameras. In section 3 we sketch the neural self-organizing model for image sequences and describe how to embed the background compensation in order to handle image sequences taken from PTZ cameras. In section 4 we present some preliminary results achieved with the implementation of the presented approach. Section 5 includes concluding remarks.

2 Background Subtraction Approaches for PTZ Cameras

Approaches to moving object detection for PTZ cameras based on background subtraction can be classified into two categories: mosaiced background-based and background compensation-based approaches.

The *mosaiced background*-based approaches [13], also referred to as *frame to background* methods [14], create and maintain a mosaic image of the whole scene background, on which background subtraction techniques are applied to extract moving object regions [15,16,17,18,19,20,21,22,23,24,25,26,27]. A mosaic image, also referred to as panorama, is a compound image built through properly composing multiple images of the same scene taken from different viewpoints and

warping them into a common reference coordinate system. The result consists in a single image of the scene, with greater resolution or spatial extent [28]. The basic idea behind mosaiced background-based approaches to moving object detection consists in the use of traditional background subtraction techniques on a mosaic image of the observed scene built by the PTZ camera images. Subsequent video frames are registered into the mosaic coordinate system first, in order to locate the subset of the mosaic background model corresponding to the current frame. This subset is then adopted to perform background subtraction and update.

The *background compensation*-based approaches [13], also named *frame-by-frame* methods [14], estimate the transformation parameters between time related images by using corresponding features extracted from these images, and create a difference image and/or a motion-compensated background in order to detect moving object regions [29,30,31,32,14,33,34,35,36,13,37]. These approaches first estimate the apparent motion of the static background due to the camera movement, and then compensate the previous image for the estimated motion. Moving objects are usually detected by frame differencing, i.e., by subtracting the compensated previous image and the current image.

Both mosaiced background-based and background compensation-based approaches need to register (or align) various pairs (or collections) of images. Therefore, assuming a motion model (e.g., affine or projective) relating pixel coordinates in one image to pixel coordinates in another, image matching techniques should be envisaged to detect and match salient features among images, with the final aim of estimating motion parameters [38].

Most methods deal with PT cameras, and some of them are applicable to a PTZ camera. Indeed, in the case of zooming, the motion parallax problem arises, where the apparent motion of objects closer to the image plane is higher than that of objects that are further away. Instead, when there is no motion parallax, the apparent motion of all objects in the scene does not depend on their distance from the camera. This can be guaranteed by rotating the camera around its optical center (approach taken by many commercial systems) and holds also for most cameras when objects are far from it [20].

According to the adopted image matching techniques, the algorithms can be classified into two main families: intensity-based (or direct) methods and feature-based methods. *Intensity-based methods* usually attempt at iteratively estimating the transformation parameters by minimizing an error function based on pixel-by-pixel brightness differences in overlapping areas. Exploiting the information associated with every single pixel, these methods achieve highly accurate registration; however, they are computationally demanding. *Feature-based methods* first extract distinctive features (i.e., regions, lines, histograms, intensity projection profiles, and keypoints) from each image, then match these features to establish a global correspondence, and finally estimate the geometric transformation between the images. Outlier detection helps to filter out bad tracked features or mismatched points.

Moreover, in the case of mosaiced background-based approaches, techniques must be developed to compute a globally consistent set of alignments among multiple images existing in the panorama, and to efficiently discover which images overlap one another [38]. Two classes of mosaicing algorithms may be distinguished, with regard to the number of frames that are simultaneously combined: global and sequential registration methods. *Global registration methods* compute the best alignment among several images by simultaneously minimizing the misregistration between all the overlapping pairs of images. They achieve accurate geometric reconstruction, but are computationally intensive and require all the images to be known in advance. *Sequential algorithms* allow the construction of a mosaic by continuously combining new images as soon as they become available. Every new image is aligned with the previous one (frame-to-frame registration) or with the mosaic built thus far (frame-to-mosaic registration). These methods allow faster computation and do not need all the images in advance; however, the achieved registration is only locally optimal, and may lead to error accumulation.

From the above analysis we can conclude that both the approaches to moving object detection by PTZ cameras have their pros and cons. Background compensation-based approaches usually require less computational cost and memory storage, while the main advantage of creating a mosaic model is to save the information of the scene background that is readily available whenever the camera moves to a new position or/and returns to a previous captured location.

It should be mentioned that other approaches to moving object detection by PTZ cameras exist that are not based on background subtraction. This is the case of the so-called *optical flow clustering*-based approaches, which compute clusters of dense or sparse optical flows in order to identify regions of movement [39,40,41]. Other methods exploit further information of the scene settings. This is the case of the works in [42,43], where the authors consider PTZ cameras performing a guard tour, following a predefined set of positions covering the area under surveillance; and the works in [44,45], where motion parameters are measured by using specialized hardware. Finally, a comprehensive introduction of PTZ camera networks, highlighting how cameras cooperation and reconfiguration can be exploited for active surveillance, has been recently provided [10].

3 Self-Organizing Background Subtraction for PTZ Cameras

Relying on the recent SOBS algorithm [12], we build the sequence background model by learning in a self-organizing manner image sequence variations, seen as trajectories of pixels in time. A neural network mapping method is proposed to use a whole trajectory incrementally in time fed as an input to the network. Each neuron computes a function of the weighted linear combination of incoming inputs, and therefore can be represented by a weight vector, obtained collecting the weights related to incoming links. An incoming pattern is mapped to the neuron whose set of weight vectors is most similar to the pattern, and weight vectors in a neighborhood of such node are updated. Differently from [12], at

each time instant the neural background model automatically compensates the eventual ego-motion due to the PTZ camera, leading to what will be called the PTZ-SOBS algorithm.

3.1 Neural Model Representation

Given an image sequence $\{I_t\}$, for each pixel \mathbf{p} in the image domain D , we build a neuronal map consisting of $n \times n$ weight vectors $m_t^{i,j}(\mathbf{p})$, $i, j = 0, \dots, n - 1$, which will be called a *model* for pixel \mathbf{p} and will be indicated as $M_t(\mathbf{p})$:

$$M_t(\mathbf{p}) = \left\{ m_t^{i,j}(\mathbf{p}), i, j = 0, \dots, n - 1 \right\}. \quad (1)$$

If every sequence frame has N rows and P columns, the complete set of models $M_t(\mathbf{p})$ for all pixels \mathbf{p} of the t -th sequence frame I_t is organized as a 2D neuronal map B_t with $n \times N$ rows and $n \times P$ columns, where the weight vectors $m_t^{i,j}(\mathbf{p})$ for the generic pixel $\mathbf{p} = (x, y)$ are at neuronal map position $(n \times x + i, n \times y + j)$, $i, j = 0, \dots, n - 1$:

$$B_t(n \times x + i, n \times y + j) = m_t^{i,j}(\mathbf{p}), i, j = 0, \dots, n - 1. \quad (2)$$

This configuration of the whole neuronal map B_t allows to easily take into account the spatial relationship among pixels and corresponding weight vectors.

3.2 Background Subtraction and Neural Model Update

At each subsequent time step t , background subtraction is achieved by comparing each pixel of the t -th sequence frame I_t with the model for that pixel. In the general case of image sequences taken from PTZ cameras, the incoming pixel \mathbf{p} of the t -th sequence frame I_t could have moved as compared to the previous time $t - 1$. Therefore, the current model $M_{t-1}(\mathbf{p})$, whose weight vectors are stored in B_{t-1} as described in Eq. (2), could be an improper model for actual pixel \mathbf{p} . In order to keep track of such spatial movements, we compute the homography H between sequence frames I_{t-1} and I_t , that allows to obtain, for each pixel \mathbf{p}' of I_{t-1} , the corresponding pixel $\mathbf{p} = H\mathbf{p}'$ of I_t . This information is exploited in order to address the proper model for current pixel \mathbf{p} , to perform background subtraction, and to update the proper model through a selective weighted running average analogous to [12].

4 Experimental Results

Several experimental tests have been conducted to validate our approach to moving object detection in image sequences taken from PTZ cameras. In the following, qualitative and quantitative results will be described for the *Lab1* sequence, that represents a typical indoor situation critical for video surveillance systems. The *Lab1* sequence, publicly available in the download section

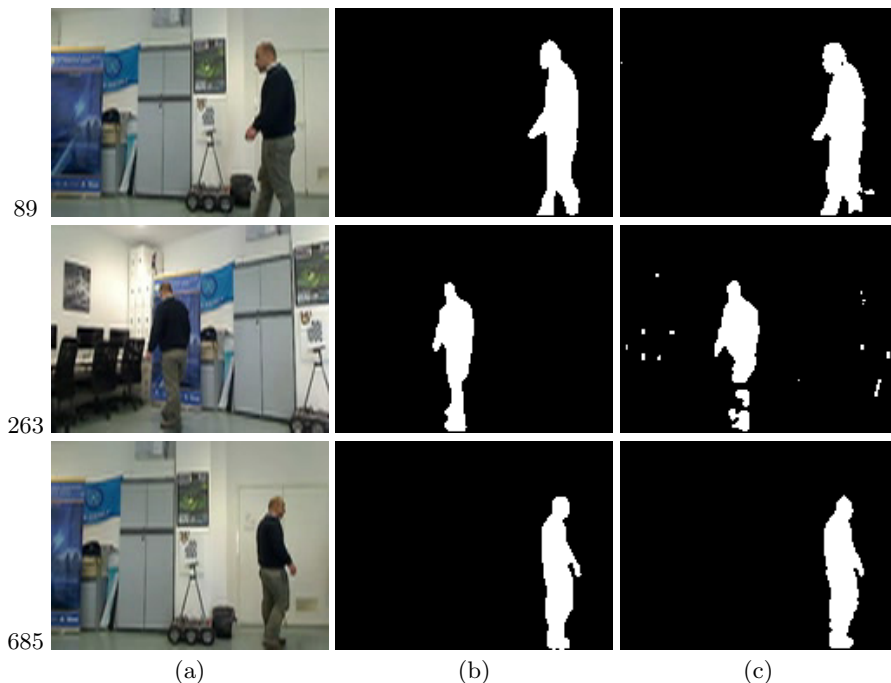


Fig. 1. Results of moving object detection for the frames 89, 263, and 685 of the *Lab1* sequence: (a) original frames; (b) ground truth masks; (c) moving object detection masks computed by the PTZ-SOBS algorithm

of <http://cvprlab.uniparthenope.it>, is an indoor sequence consisting of 886 frames of 160×120 spatial resolution. The scene consists of an office, where a person comes in and walks, and the PTZ camera follows the movement of the person. Representative frames together with obtained results are reported in Fig. 1, where we report the original sequence frames no. 89, 263, and 685 (first column), the corresponding ground truth masks (second column), and the moving object detection mask computed by the PTZ-SOBS algorithm (third column). In frame 89 a person has entered the office, and the camera is panning from right to left in order to follow him; the compensated background provides a good representation of the real background, and the corresponding detection mask is quite accurate. In frame 263 a person has moved further left, while the camera is still following him. The compensated background still provides a good representation of the real background. The corresponding detection mask is quite accurate, although few white pixels of the moving person and of the background have been misdetections. In frame 685 the person is going back from left to right, still followed by analogous movement of the camera. The compensated background still provides a good representation of the real background, as shown by the quite accurate corresponding detection mask.

The average segmentation accuracy results achieved by the proposed approach on the considered sequence, in terms of Precision, Recall and F_1 measure, are quite encouraging. Indeed, the achieved $Recall=0.94$ value ensures that most of the moving pixels are indeed detected as moving; at the same time, the high $Precision=0.88$ value indicates that only few of the pixels detected as moving are instead background pixels. The consequent high value of $F_1=0.91$, that is the weighted harmonic mean of $Precision$ and $Recall$, allows us to deduce the overall high segmentation accuracy of the proposed approach.

5 Conclusions

We present an approach to the problem of moving object detection in image sequences taken from PTZ cameras based on the idea of exploiting the available knowledge concerning the self-organized learning behavior of the brain, which is the foundation of human visual perception, and traducing it into models and algorithms that can accurately solve the problem. A neural self-organizing background model is presented, that automatically adapts to variations in the scene background, both arising in a usual stationary camera setting and due to the PTZ camera movement. The compensated background model is adopted for accurate moving object detection, as demonstrated by experimental results on real image sequences.

References

1. Collins, R.T., Lipton, A.J., Kanade, T., Fujiyoshi, H., Duggins, D., Tsin, Y., Tolliver, D., Enomoto, N., Hasegawa, O., Burt, P., Wixson, L.: A system for video surveillance and monitoring. Technical Report CMU-RI-TR-00-12, Carnegie Mellon University, Pittsburgh, PA (2000)
2. Cheung, S.C.S., Kamath, C.: Robust techniques for background subtraction in urban traffic video. In: Panchanathan, S., Vasudev, B. (eds.) Proc. Visual Communications and Image Processing, SPIE, vol. 5308, pp. 881–892 (2004)
3. Elhabian, S., El Sayed, K., Ahmed, S.: Moving object detection in spatial domain using background removal techniques: State-of-art. Recent Patents on Computer Science 1(1), 32–54 (2008)
4. Piccardi, M.: Background subtraction techniques: a review. In: Proc. IEEE SMC, vol. 4, pp. 3099–3104 (October 2004)
5. Radke, R.J., Andra, S., Al-Kofahi, O., Roysam, B.: Image change detection algorithms: A systematic survey. IEEE Trans. Image Process. 14, 294–307 (2005)
6. Toyama, K., Krumm, J., Brumitt, B., Meyers, B.: Wallflower: principles and practice of background maintenance. In: Proc. ICCV, vol. 1, pp. 255–261 (1999)
7. Huang, K., Wang, L., Tan, T., Maybank, S.: A real-time object detecting and tracking system for outdoor night surveillance. Pattern Recognition 41(1), 432–444 (2008)
8. Cristani, M., Farenzena, M., Bloisi, D., Murino, V.: Background subtraction for automated multisensor surveillance: a comprehensive review. EURASIP J. Adv. Signal Process, 43:1–43:24 (February 2010)

9. Elgammal, A.: Figure-ground segmentation-pixel-based. In: Moeslund, T.B., Hilton, A., Krüger, V., Sigal, L. (eds.) *Visual Analysis of Humans*, pp. 31–51. Springer, London (2011)
10. Micheloni, C., Rinner, B., Foresti, G.: Video analysis in pan-tilt-zoom camera networks. *EEE Signal Processing Magazine* 27(5), 78–90 (2010)
11. Sheikh, Y., Javed, O., Kanade, T.: Background subtraction for freely moving cameras. In: 12th IEEE International Conference on Computer Vision, pp. 1219–1225 (October 2009)
12. Maddalena, L., Petrosino, A.: A self-organizing approach to background subtraction for visual surveillance applications. *IEEE Trans. Image Process* 17(7), 1168–1177 (2008)
13. Suhr, J.K., Jung, H.G., Li, G., Noh, S.I., Kim, J.: Background compensation for pan-tilt-zoom cameras using 1-d feature matching and outlier rejection, pp. 371–377. IEEE Computer Society Press, Los Alamitos (March 2011)
14. Micheloni, C., Foresti, G.L.: Real-time image processing for active monitoring of wide areas. *Journal of Visual Communication and Image Representation* 17(3), 589–604 (2006)
15. Bartoli, A., Dalal, N., Horaud, R.: Motion panoramas. *Computer Animation and Virtual Worlds* 15(5), 501–517 (2004)
16. Bevilacqua, A., Azzari, P.: A fast and reliable image mosaicing technique with application to wide area motion detection. In: Kamel, M.S., Campilho, A. (eds.) *ICIAR 2007*. LNCS, vol. 4633, pp. 501–512. Springer, Heidelberg (2007)
17. Del Bimbo, A., Lisanti, G., Masi, I., Pernici, F.: Continuous recovery for real time pan tilt zoom localization and mapping. In: 8th IEEE International Conference on Advanced Video and Signal-Based Surveillance (AVSS), pp. 160–165 (September 2011)
18. Hayman, E., Eklundh, J.O.: Statistical background subtraction for a mobile observer. In: *Proceedings Ninth IEEE International Conference on Computer Vision 2003*, vol. 1, pp. 67–74 (October 2003)
19. Jin, Y., Tao, L., Di, H., Rao, N., Xu, G.: Background modeling from a free-moving camera by multi-layer homography algorithm. In: 15th IEEE International Conference on Image Processing (ICIP), pp. 1572–1575 (October 2008)
20. Mittal, A., Huttenlocher, D.: Scene modeling for wide area surveillance and image synthesis. In: *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, vol. 2, pp. 160–167 (2000)
21. Monari, E., Pollok, T.: A real-time image-to-panorama registration approach for background subtraction using pan-tilt-cameras. In: 8th IEEE International Conference on Advanced Video and Signal-Based Surveillance (AVSS), pp. 237–242 (September 2011)
22. Prati, A., Seghedoni, F., Cucchiara, R.: Fast dynamic mosaicing and person following. In: *Proceedings of the 18th International Conference on Pattern Recognition, ICPR 2006*, pp. 920–923. IEEE Computer Society, Washington, DC (2006)
23. Ren, Y., Chua, C.S., Ho, Y.K.: Statistical background modeling for non-stationary camera. *Pattern Recognition Letters* 24(1-3), 183–196 (2003)
24. Sankaranarayanan, K., Davis, J.W.: PTZ Camera Modeling and Panoramic View Generation via Focal Plane Mapping. In: Kimmel, R., Klette, R., Sugimoto, A. (eds.) *ACCV 2010, Part II*. LNCS, vol. 6493, pp. 580–593. Springer, Heidelberg (2011)
25. Sugaya, Y., Kanatani, K.: Extracting moving objects from a moving camera video sequence. *Mem. Fac. Eng. Okayama Univ.* 39, 56–62 (2005)

26. Xue, K., Liu, Y., Chen, J., Li, Q.: Panoramic background model for PTZ camera. In: 3rd International Congress on Image and Signal Processing (CISP), vol. 1, pp. 409–413 (October 2010)
27. Zhang, J., Wang, Y., Chen, J., Xue, K.: A framework of surveillance system using a ptz camera. In: 3rd IEEE International Conference on Computer Science and Information Technology (ICCSIT), vol. 1, pp. 658–662 (July 2010)
28. Bevilacqua, A., Azzari, P.: High-quality real time motion detection using PTZ cameras. In: IEEE International Conference on Video and Signal Based Surveillance, vol. 23 (November 2006)
29. Araki, S., Matsuoka, T., Takemura, H., Yokoya, N.: Real-time tracking of multiple moving objects in moving camera image sequences using robust statistics. In: Proceedings Fourteenth International Conference on Pattern Recognition, vol. 2, pp. 1433–1435 (August 1998)
30. Bin, L., Qiang, Z., Huanxia, L.: Research on background motion estimation and compensation in image sequences. In: 2011 International Conference on Mechatronic Science, Electric Engineering and Computer (MEC), pp. 1370–1373 (August 2011)
31. Cai, Q., Mitiche, A., Aggarwal, J.: Tracking human motion in an indoor environment. In: International Conference on Image Processing, vol. 1, pp. 215–218 (October 1995)
32. Lee, K.W., Ryu, S.W., Lee, S.J., Park, K.T.: Motion based object tracking with mobile camera. *Electronics Letters* 34(3), 256–258 (1998)
33. Nguyen, T.T., Jeon, J.W.: Real-Time Background Compensation for PTZ Cameras Using GPU Accelerated and Range-Limited Genetic Algorithm Search. In: Ho, Y.-S. (ed.) PSIVT 2011, Part I. LNCS, vol. 7087, pp. 85–96. Springer, Heidelberg (2011)
34. Pham, X.D., Cho, J.U., Jeon, J.W.: Background compensation using hough transformation. In: IEEE International Conference on Robotics and Automation (ICRA), pp. 2392–2397 (May 2008)
35. Rao, N., Di, H., Xu, G.: Joint correspondence and background modeling based on tree dynamic programming. In: 18th International Conference on Pattern Recognition (ICPR), vol. 2, pp. 425–428 (2006)
36. Shun, Z., Xiuqin, S., Liyin, X.: Global motion compensation for image sequences and motion object detection. In: 2010 International Conference on Computer Application and System Modeling (ICCASM), vol. 1, pp. 406–409 (October 2010)
37. Tordoff, B., Murray, D.: Reactive control of zoom while fixating using perspective and affine cameras. *IEEE Trans. Pattern Anal. Mach. Intell.* 26(1), 98–112 (2004)
38. Szeliski, R.: Image alignment and stitching: a tutorial. *Found. Trends. Comput. Graph. Vis.* 2(1), 1–104 (2006)
39. Jung, Y.K., Lee, K.W., Ho, Y.S.: Feature-Based Object Tracking with an Active Camera. In: Chen, Y.-C., Chang, L.-W., Hsu, C.-T. (eds.) PCM 2002. LNCS, vol. 2532, pp. 1137–1144. Springer, Heidelberg (2002)
40. Kim, J., Ye, G., Kim, D.: Moving object detection under free-moving camera. In: 17th IEEE International Conference on Image Processing (ICIP), pp. 4669–4672 (September 2010)
41. Varcheie, P., Bilodeau, G.A.: Adaptive fuzzy particle filter tracker for a PTZ camera in an IP surveillance system. *IEEE Transactions on Instrumentation and Measurement* 60(2), 354–371 (2011)

42. Guillot, C., Taron, M., Sayd, P., Pham, Q.C., Tilmant, C., Lavest, J.M.: Background Subtraction for PTZ Cameras Performing a Guard Tour and Application to Cameras with Very Low Frame Rate. In: Koch, R., Huang, F. (eds.) ACCV Workshops 2010, Part I. LNCS, vol. 6468, pp. 33–42. Springer, Heidelberg (2011)
43. Guillot, C., Taron, M., Sayd, P., Pham, Q.C., Tilmant, C., Lavest, J.M.: Background subtraction adapted to PTZ cameras by keypoint density estimation. In: Proc. BMVC., vol. 34, pp. 1–10 (2010)
44. Murray, D., Basu, A.: Motion tracking with an active camera. *IEEE Trans. Pattern Anal. Mach. Intell.* 16(5), 449–459 (1994)
45. Xu, C., Liu, J., Kuipers, B.: Motion segmentation by learning homography matrices from motor signals. In: Proceedings of the 2011 Canadian Conference on Computer and Robot Vision, CRV 2011, pp. 316–323. IEEE Computer Society, Washington, DC (2011)

Control of Coffee Grinding with General Regression Neural Networks

Luca Mesin, Diego Alberto, and Eros Pasero

Dipartimento di Elettronica e Telecomunicazioni
Politecnico di Torino

C.so Duca degli Abruzzi 24, 10129 Torino, Italy
{luca.mesin, diego.alberto, eros.pasero}@polito.it
<http://www.polito.it>

Abstract. Standardization and the assessment of the quality of the final product is fundamental in food industry. Coffee particle properties are monitored continuously during coffee beans grinding. Operators control the grinders in order to keep coffee particle granulometry within specific thresholds. In this work, a general regression neural network approach is used to learn to control two grinders used for coffee production at LAVAZZA factory, obtaining average control error of the order of a few μm . The results appear promising for the future development of an automatic decision support system.

Keywords: General Regression Neural Network, Partial mutual information, coffee, grinding.

1 Introduction

European regulations requires continuous controls in food industry to assess the quality of food [1], [2]. The quality of the products must be maintained at a high level during the complete production chain, even if many external factors can influence the quality of the final product: material, mashing, fermentation, maturation and blending conditions. Thus, automated controls are needed to guarantee high and stable production quality [3], [4]. The control should be adaptive, in order to fit the specific property of the product [5], [6]. For example, Artificial Neural Networks (ANN) were applied to food industry resulting as an effective and promising adaptive processing tool [7], [8]. As an alternative, General Regression Neural Networks (GRNN) have also been proposed to estimate a dependent variable on the basis of noisy measurements [9] and were proposed for prediction, modelling, mapping, interpolation or control applications.

In this paper, we present an application of GRNNs to the control of two grinders used for coffee production in LAVAZZA factory [10]. A few neural methods have been proposed in the literature for similar applications [11]-[14]. Information on granulometry and density of the coffee beans are recorded along

the production line. On the basis of such data, human operators regulate the grinders in order to maintain the characteristics of the product within thresholds which are specific for the considered blend. Our aim is to provide to the operators indications on the most secure regulation, in order to avoid unwanted breakdown of the complete production chain that may occur if some of the most important parameters go out of range.

2 Methods

2.1 Data Acquisition

Granulometry and densitometry are recorded along the production line, examining coffee particles obtained after grinding. Time sampling is not fixed: a sample is usually taken when the operator suspects that some parameter is going out of range. Coffee granulometry is obtained using a laser diffraction analyser, providing a histogram of size with 17 bins from which 4 indexes are computed, as detailed below. Moreover, coffee bulk density is measured using a densitometer. Finally, the temperature was measured close to the grinders. This parameter is important as a high temperature could indicate that there is a mistake between the regulation set up and the actual distance between the wheels of the grinders. Then, data are loaded by the operators and managed by the LAVAZZA Quality Control System which, by means of a Statistical Process Control software, allows data storing, searching and simple analysis, like identification of linear trends. The size distribution is about bimodal. Fig. 1 shows the four parameters that are extracted to describe this probability density function (PDF): x_{10} , x_{50} , x_{90} , and coffee dust (or simply dust). The first three values represent the upper limits of the interval in the PDF that contain the 10%, 50%, and 90%, respectively, of all coffee particles dimension, they are expressed in μm . The fourth parameter (dust) estimates the particles percentage with a size lower than $100 \mu\text{m}$. The production line recorded only four variables extracted from the granulometry PDF and the density of the product. Such variables were estimated as the most significant in characterizing the final product. Operators were instructed to manually regulate the grinders when these variables went out of specific ranges. Our full dataset consisted of 2570 patterns. On the basis of these measures, the operator decides if a change in the regulation of burr grinders is required. Sometimes, the granulometry variables are close to threshold values which requires the product to be monitored with a high frequency. If such thresholds are undergone, production is stopped.

2.2 Data Pre-processing

Data Interpolation. As mentioned above, data were not regularly sampled, but the time sample interval varied. Data were resampled at a constant sampling time of 10 minutes (using cubic spline interpolation) within time windows in which the intervals between consecutive samples were lower than 30 minutes.

Actual data on density and granulometry were used together with the last regulation of the two grinders were used as measures on the basis of which the control of the grinders was to be estimated.

Selection of Input Features. The selection of optimal input features for the regression algorithm is of great importance, in order to improve performance, to reduce the measurement noise and to counteract the difficulties of facing a problem with large dimension. We used a method based on the selection of the input data providing maximal information on the output. In order to avoid redundancy among input features, the algorithm proposed in [15]-[17] was applied. It determines the interdependencies between candidate variables computing the Partial Mutual Information (PMI).

The mutual information $MI(X; Y)$ of two random variables is defined as the reduction in uncertainty with respect to Y due to observation of X :

$$MI(X; Y) = H(X) + H(Y) - H(X, Y) . \quad (1)$$

Where H is the information entropy:

$$H(Z) = - \int f_Z(z) \ln f_Z(z) dz . \quad (2)$$

Mutual Information is nonnegative. It is zero only when X and Y are independent variables. Thus, it is used as a sort of measure of distance from independence. An important issue in estimating the mutual information is the estimation of joint and marginal probability density functions. A robust estimation was performed using Parzen method [18], [19]:

$$\hat{f}(x) = \frac{1}{n} \sum_{i=1}^n K_h(x - x_i) . \quad (3)$$

Where n is the number of dimension of the considered random variable and the kernel K was assumed to be Gaussian with standard deviation h (kernel bandwidth):

$$K_h(x - x_i) = \frac{1}{(\sqrt{2\pi}h)^d \sqrt{|S|}} \exp\left(-\frac{\|x - x_i\|^2}{2h^2}\right) . \quad (4)$$

Where S is the sample covariance matrix. Thus, the following expression for the joint probability density function is obtained:

$$\hat{f}_{XY}(x, y) = \frac{1}{(\sqrt{2\pi}h)^d \sqrt{|S|}} \frac{1}{n} \sum_{i=1}^n \exp\left(-\frac{\|x - x_i\|^2}{2h^2}\right) \exp\left(-\frac{\|y - y_i\|^2}{2h^2}\right) . \quad (5)$$

Where, in the following, x and y represent the feature vector and the output variable (the control of the burr grinders), respectively. The value of the bandwidth greatly affects the kernel density estimation. Too small values lead to data

under-smoothing, whereas too large bandwidth determines an over-smooth of the sample probability density-function. Different algorithms have been developed to search for an optimal value of the bandwidth h , but it is usually specific for the application. Given the expression (5), the discrete estimation of (1) is the following:

$$MI(X; Y) = \frac{1}{n} \sum_{i=1}^n \ln \frac{\widehat{f}_{XY}(x_i, y_i)}{\widehat{f}_X(x_i) \widehat{f}_Y(y_i)}. \quad (6)$$

Where x_i and y_i are the i^{th} bivariate sample data pair in a sample of size n , and $\widehat{f}_X(x_i)$, $\widehat{f}_Y(y_i)$ and $\widehat{f}_{XY}(x_i, y_i)$ are respective marginal and joint probability-densities estimated at the data points. Given the mutual information between two variables, the PMI can be defined. Assume that Y is a scalar random variable with observations y , X is a candidate input with observations x , and Z is the set of already selected input features. The PMI between X and Y is computed using the residuals of both variables once the effects of the selected inputs Z have been considered. The dependence between variables is removed by computing for each x and y the residuals:

$$x' = x - \mathbf{E}[x|z] \quad y' = y - \mathbf{E}[y|z]. \quad (7)$$

Where $\mathbf{E}[\cdot|z]$ is the regression of the chosen variable based on the features already selected. The regression of the input X and the output Y are computed using the kernel density estimation approach:

$$\begin{aligned} \mathbf{E}[x|z] &= \frac{\int_{-\infty}^{+\infty} x \cdot f_{XZ}(x, z) dx}{\int_{-\infty}^{+\infty} f_{XZ}(x, z) dx} = \\ &= \frac{1}{n} \frac{\sum_{i=1}^n \int_{-\infty}^{+\infty} x \cdot \exp\left(-\frac{\|x-x_i\|^2}{2h^2}\right) \exp\left(-\frac{\|z-z_i\|^2}{2h^2}\right) dx}{\sum_{i=1}^n \int_{-\infty}^{+\infty} \exp\left(-\frac{\|x-x_i\|^2}{2h^2}\right) \exp\left(-\frac{\|z-z_i\|^2}{2h^2}\right) dx} \cong \frac{1}{n} \frac{\sum_{i=1}^n x_i K_h(z - z_i)}{\sum_{i=1}^n K_h(z - z_i)}. \end{aligned} \quad (8)$$

$$\begin{aligned} \mathbf{E}[y|z] &= \frac{\int_{-\infty}^{+\infty} y \cdot f_{YZ}(y, z) dy}{\int_{-\infty}^{+\infty} f_{YZ}(y, z) dy} = \\ &= \frac{1}{n} \frac{\sum_{i=1}^n \int_{-\infty}^{+\infty} y \cdot \exp\left(-\frac{\|y-y_i\|^2}{2h^2}\right) \exp\left(-\frac{\|z-z_i\|^2}{2h^2}\right) dy}{\sum_{i=1}^n \int_{-\infty}^{+\infty} \exp\left(-\frac{\|y-y_i\|^2}{2h^2}\right) \exp\left(-\frac{\|z-z_i\|^2}{2h^2}\right) dy} \cong \frac{1}{n} \frac{\sum_{i=1}^n y_i K_h(z - z_i)}{\sum_{i=1}^n K_h(z - z_i)}. \end{aligned} \quad (9)$$

For the problem at hand, the output Y was a vector of two variables: the control of the first and second grinder. The possible candidates were the previous regulation of the grinders, the density and granulometry of coffee beans and the temperature of the grinders. On the basis of the PMI criterion, the previous regulation of the grinders were the variable most related to the output. But such

an information was not useful to make a change of regulation, needed to maintain good properties of the final product in a situation in which density and granulometry data are drifting out of range. Thus, density and granulometry data were also included. The temperature of the grinders had a low mutual information with the output and was excluded. A portion of data, indicated with X_1 and Y_1 , was used to train the method, which was then applied on another portion of data, indicated with X_2 , for validation or test purposes. Specifically, the estimated control Y_2 corresponding to X_2 was estimated using the following expression:

$$Y_2(X_2) = \frac{1}{n} \frac{\sum_{i=1}^n Y_1^i \exp\left(-\frac{\|X_2 - X_1^i\|^2}{2h^2}\right)}{\sum_{i=1}^n \exp\left(-\frac{\|X_2 - X_1^i\|^2}{2h^2}\right)}. \quad (10)$$

To check the performances of the method, experimental data were divided into training (70% of data), validation (15% of data) and test sets (15% of data), choosing randomly the samples associated to each set, for 10 times. The network with highest average generalization capability (i.e., the minimum error) on the validation set was selected, and the error on the test set was further considered. All the analyses have been carried out using a MATLAB platform [20].

2.3 Development of an Automatic Control of the Burr Grinders

The regression of a dependent variable, Y , on an independent variable, X , is the computation of the most probable value of Y for each value of X based on a finite number of possibly noisy measurements of X and the associated values of Y . Once selected the input features, the output was estimated using equation 5. The estimation of the probability density function relating the input features to the output control variables allowed computing also the standard deviation of the estimate, which can be considered as an indication of the reliability of the estimation provided. Input and output data were normalized with respect to their standard deviation after subtracting the mean. After processing, data were de-normalized with the inverse procedure. Different bandwidths of the kernel were investigated and the one providing optimal generalization to the validation set was selected (equal to 0.65).

3 Results

Fig. 1 shows the raw data considered. The granulometry of a sample of the product (Fig. 1A) is used to extract the variables (x_{10} , x_{50} , x_{90} , and dust, Fig. 1B), which, together with the density of the product, constitute the input features for the neural control systems. The control of the grinders (which is the output to be learned by the control system) is the distance between the wheels constituting the two burr grinders (Fig. 1C).

An example of processing is shown in Fig. 2, where a possible choice of training, validation and test sets are considered. The estimate of the control is superimposed to raw data. Indication of mean standard deviation of the estimate is

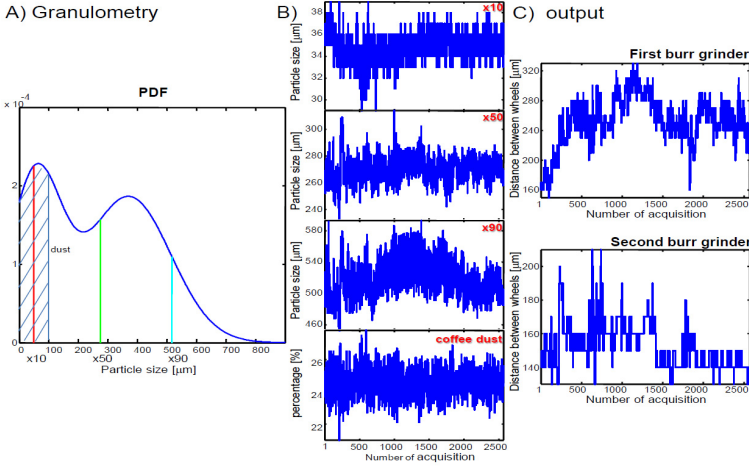


Fig. 1. A) Granulometry (PDF) obtained fitting real data with the sum of two Gaussians. Data extracted from the PDF are indicated: x_{10} , x_{50} and x_{90} are the 10th, the 50th and the 90th percentile; *dust* is the integral of the PDF from 0 to 100 μm . B) Data extracted from granulometry (x_{10} , x_{50} , x_{90} , and *dust*), which, together with the density of the product, constitute the input time series for the neural control systems. C) First and second burr grinding output values (output to be learned by the neural control systems).

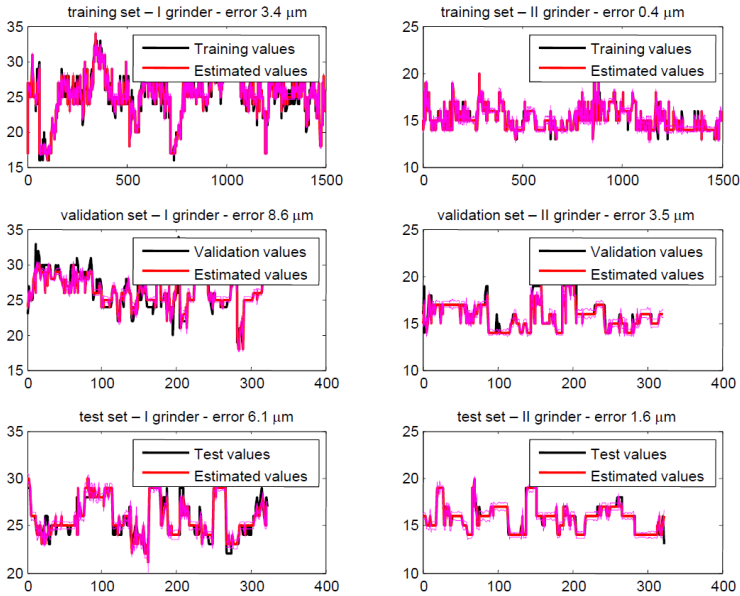


Fig. 2. Example of training, validation and test sets, with estimate superimposed to raw data. Indication of standard deviation of the estimate is provided by magenta lines, corresponding to mean \pm standard deviation. Mean error is also provided.




Fig. 3. Test Set error histograms for predicted I burr grinding (left) and II grinding values (right)

provided. Moreover, the mean estimation error is indicated. The distribution of the mean control error on the test set is shown in Fig. 3.

4 Discussion

A neural system was developed to learn how to control the two grinders used for coffee bean grinding in LAVAZZA. The input variables are obtained by measuring density and granulometry of samples. Control errors of a few μm were obtained (with a control step of $10 \mu\text{m}$). The results of our automatic control were also checked by an experienced operator, indicating that they agree with his choice. Moreover, they are in line with a recent work performed on the same dataset using multi-layer perceptron [21]. Thus, our results are promising for a future integration of the control system in the production line, to support operators in controlling burr grinders.

Acknowledgments. The authors are deeply indebted to Alberto Cabilli (R&D, LAVAZZA, Turin, Italy) for the selection of the experimental data and interpretation of results.

References

1. <http://www.food.gov.uk/foodindustry/regulation/>
2. Leslie, R.B., Oliveira, J.C., Medin, A.G.: Food Forum: a Research Forum for an Innovative and Globally Competitive European Food Industry. *Food Research International* 33(3-4), 295–297 (2000)

3. Linko, S.: Expert Systems - what can they do for the food industry? *Trends in Food Science & Technology* 9(1), 3–12 (1998)
4. Ilyukhin, S.V., Haley, T.A., Singh, R.K.: A survey of automation practices in the food industry. *Food Control* 12(5), 285–296 (2001)
5. Banga, J.R., Bsa-Canto, E., Moles, C.J., Alonso, A.A.: Improving food processing using modern optimization methods. *Trends in Food Science & Technology* 14(4), 131–144 (2003)
6. Bhuvaneswari, N.S., Uma, G., Rangaswamy, T.R.: Adaptive and optimal control of a non-linear process using intelligent controllers. *Applied Soft Computing* 9(1), 182–190 (2009)
7. Torrecilla, J.S., Otero, L., Sanz, P.D.: Artificial neural networks: a promising tool to design and optimize high-pressure food processes. *Journal of Food Engineering* 69(3), 299–306 (2005)
8. Guyer, D., Yang, X.: Use of genetic artificial neural networks and spectral imaging for defect detection on cherries. *Computers and Electronics in Agriculture* 29(3), 179–194 (2000)
9. Specht, D.F.: A General Regression Neural Network. *IEEE Transactions on Neural Networks* 2(6), 568–576 (1991)
10. LAVAZZA S.p.A., Strada Settimo 410, Torino
11. Tominaga, O., Ito, F., Hanai, T., Honda, H., Kobayashi, T.: Sensory Modeling of Coffee with a Fuzzy Neural Network. *Journal of Food Science* 67(1), 363–368 (2001)
12. Hernandez, J.A., Heyd, B., Trystram, B.: Prediction of brightness and surface area kinetics during coffee roasting. *Journal of Food Engineering* 89(2), 156–163 (2008)
13. Yip, D.H.F., Yu, W.W.H.: Classification of Coffee using Artificial Neural Network. In: *IEEE International Conference on Evolutionary Computation*, pp. 655–658. IEEE Press, Nagoya (1996)
14. Singh, S., Hines, E.L., Gardner, J.W.: Fuzzy neural computing of coffee and tainted water data from an electronic nose. *Sens. Actuators B* 30(3), 190–195 (1996)
15. Sharma, A.: Seasonal to interannual rainfall probabilistic forecasts for improved water supply management: Part 1 - A strategy for system predictor identification. *Journal of Hydrology* 239(1-4), 232–239 (2000)
16. Mesin, L., Orione, F., Taormina, R., Pasero, E.: A feature selection method for air quality forecasting. In: *20th IEEE International Conference on Artificial Neural Networks*, pp. 489–494. IEEE Press, Thessaloniki (2010)
17. May, R.J., Maier, H.R., Dandy, G.C., Gayani Fernando, T.M.K.: Non-linear variable selection for artificial neural networks using partial mutual information. *Envir. Mod. and Soft.* 23, 1312–1326 (2008)
18. Parzen, E.: On Estimation of a Probability Density Function and Mode. *Annals of Math. Statistics* 33, 1065–1076 (1962)
19. Costa, M., Moniaci, W., Pasero, E.: INFO: an artificial neural system to forecast ice formation on the road. In: *IEEE International Symposium on Computational Intelligence for Measurement Systems and Applications*, pp. 216–221 (2003)
20. MATLAB, The language of technical computing, <http://www.mathworks.com>
21. Mesin, L., Alberto, D., Pasero, E., Cabilli, A.: Control of coffee grinding with Artificial Neural Networks. In: *22nd IEEE International Conference on Artificial Neural Networks*. IEEE Press, Brisbane (2012)

Defects Detection in Pistachio Nuts Using Artificial Neural Networks

Paolo Motto Ros and Eros Pasero

Laboratorio di Neuronica, Dipartimento di Elettronica,
Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Torino, Italy
{paolo.mottoros,eros.pasero}@polito.it

Abstract. In-line automated inspection of raw materials is one of the major concerns in food industry. The aim of this paper is to devise a method to sort pistachio nuts, in order to reject the bad ones. X-ray images are used to compute a set of fuzzy features, with membership functions automatically inferred from the positive samples. A functional-link neural network is then used for the proper classification task. By means of a repeated cross-validation, the proposed solution showed a correct recognition rate of 99.6 %, with a false positive rate of 0.3 % with a single classifier and 0.1 % with a combined one.

Keywords: Image Processing, Pattern Recognition, Fuzzy Systems, Functional-Link Neural Networks.

1 Introduction

Food industry is focusing even more attention on making the overall production process more efficient; often the most cost-effective solution is the automated inspection of all the parts entering in the process. The inspection process is performed in-line, which means that every item (being fruit, nut, or vegetable) has to be verified in a real-time way with a non-destructive test. The first point leads to evaluate those methods which are efficient from the resource requirement point of view (mainly computation time in case of software algorithms) and the second point automatically discards those approaches that make the item useless for the production after the tests.

The above considerations make machine vision applications the perfect match for the issues in this field [1], since in almost all the systems the collected data can be considered images. Optical analysis is only one possible approach; color changes can be good indicators of defects inside a nut, for example. Image analysis is often used to have a first information about external defects. Main limitation of this approach is that internal defects can not be visible and changes in color are not always due to defects in nuts; a high number of good products are therefore rejected. Other approaches are required to inspect internal defects, like visible/near-infrared, fluorescence, magnetic resonance spectroscopy

or X-rays [2]. Among these, X-ray histogram features and their spatial derivatives have been used for detection of insect-infested fruits, showing initially not exceptional performances. But recent technological improvements (in terms of cost-effectiveness) [3] have proven to be a reliable solution for insect-infested pistachio [4,5,6,7,8], even if it is not always the selected approach [9,10].

These machines are composed by a conveyor belt moving the raw materials inside a main unit (closed in a lead shielding), where there is a X-ray source tube on one side and a X-ray linear sensor on the opposite, with the belt, carrying inspected products, in the between [11,1]. Almost all this kind of inspection machines share this structure, but what makes the difference are the involved image processing algorithms. In this respect, traditional techniques are still useful and effective [9,11,7,8], but a growing interest has been focused on Artificial Neural Networks (ANN) [12,5,4,6], or, generally speaking, data-driven approaches (e.g., a k -nearest neighbor classifier, as in [13]). A data-driven approach can boost the performances and the ease of setup/tune the inspection machine due to the ability of automatically learn from real sample images. Instead of focusing the efforts in defining or trimming several complex and inter-related low-level parameters (thresholds, filter coefficients and so on), the emphasis is on defining a set of meaningful features for the classification, the quality and size of the dataset (used to train/test the system), and, in the end, its ability to generalize the expected behavior with previously unknown data. From a practical standpoint, it enables the inspection process not to be tight to a particular product variety and to be easily adapted (automatic learning) to broader set of inspection needs/requirements.

The goal of this study is to investigate the flexibility of X-ray techniques and ANN to separate good pistachio nuts (i.e., to be used as raw material for other food products) from the bad ones; rejected pistachios can be divided into damaged and infected. The first ones are those which have been somehow either damaged in the previous production stages or not good at all as nuts; the second ones are those infected by insects or similar. All the nuts have been manually verified to assess their quality; afterwards they have been scanned with a X-ray scanner to create the database for this study. The current dataset is composed by 100 good, 75 damaged and 100 infected pistachio nuts.

The proposed solution is structured as follows:

1. Feature extraction
 - (a) Segmentation
 - (b) Histogram computation
 - (c) Fuzzy membership inference
 - (d) Fuzzy histogram comparison
 - (e) Score sorting
2. Classification
 - (a) Chebyshev functional link
 - (b) Normalization
 - (c) Artificial neural network
 - (d) Descaling and thresholding

Next section will present the feature extraction algorithm we used to improve the discriminant functions to distinguish good from bad pistachio nuts. Section 3 will describe ANN we used to classify the nuts. Section 4 will comment the obtained results and present a comparison with other methods. Section 5 will resume our approach and show the performances.

2 Feature Extraction

The choice of the set of features is not a straightforward task, since we should ensure that the representation of the problem domain in the feature space is complete and sound. Considering the pistachio nuts data set, features can be roughly divided into two classes: geometric ones and statistics about the image content. The first ones can be the shape, the computed volume, relations between segmented sub-parts (“inner” or “outer” nutmeat) or any measurements based on them, like the fitness of the shape respect to a given geometric shape, the regularity of the border and so on. However, this type of information is only useful if the need is to grade the nuts, e.g., to develop a sorter devoted to distinguish among different varieties. Furthermore, they strongly depend on how each nut is laid on the conveyor belt, making almost impossible to formulate any reasonable assumption. On the opposite, using statistics about the image content allows to have features invariant respect to rotation, scale, shape, position, but still able to capture the “fitness” of the inspected product. The most used statistics relies on the concept of histogram, and it can be either a set of different histograms, each one related to lower level filter results [5,6], a set of statistics computed on the classical intensity histogram, or a mix of both [4,13].




Fig. 1. Histogram comparison membership function

In our case, we studied a different approach, relying on the assumption that the intensity histogram can be regarded as a fingerprint of the pistachio, so that we could compute a prototype histogram for each nut class and then measure, for each sample to be tested, the similarity respect to the reference data. It is clear that it will be almost impossible to have a perfect match, simply because such reference is created artificially, but nevertheless we have a score upon which we can train a proper classifier. This idea directly leads to the fuzzy logic, where

boolean information are substituted by real numbers (usually ranging between 0 and 1) expressing the degree of membership $\mu_p(x)$ of a variable x to a given property p . In the proposed method we have one membership function (with a trapezoidal shape, showed in Fig. [1](#)) μ_i for each histogram bin h_i :

$$\mu_i(x) = \begin{cases} 0 & \text{if } x \leq (1 - \gamma)P_i^\beta, \\ \frac{x - (1 - \gamma)P_i^\beta}{(1 - \alpha)P_i^{50} - (1 - \gamma)P_i^\beta} & \text{if } (1 - \gamma)P_i^\beta < x \leq (1 - \alpha)P_i^{50}, \\ 1 & \text{if } (1 - \alpha)P_i^{50} < x \leq (1 + \alpha)P_i^{50}, \\ \frac{x - (1 + \alpha)P_i^{50}}{(1 + \gamma)P_i^{(100-\beta)} - (1 + \alpha)P_i^{50}} & \text{if } (1 + \alpha)P_i^{50} < x \leq (1 + \gamma)P_i^{(100-\beta)}, \\ 0 & \text{if } x > (1 + \gamma)P_i^{(1-\beta)}. \end{cases} \quad (1)$$

where α and γ are coefficients ranging between 0 and 1, P_i^n is the n th percentile computed for the i th histogram bin, among the values obtained from the reference images (i.e., the positive samples in the dataset), and β determines which percentiles will be computed. These inferred membership functions are then used to extract the features from all the images to be analyzed. This feature extraction uses a data-driven approach, effectively learning and embedding partial knowledge of the problem domain in the first stage of the overall classification process.

As output we have a set of fuzzy numbers (one for each histogram bin) which can be used in a fuzzy rule system for the decision process. One idea could be to use a logical conjunction, assuming that if *all* the bins of a test image are similar to the reference ones, then such image should belong to the positive set. In practice, this is a too restrictive way to solve the problem, leading to too many misclassifications. One way to loosen such rule is to replace the logical operator with the disjunction, meaning that if *any* of the test bins is similar to the reference ones, then the outcome will be positive. Again, we would have too many misclassifications, mainly because we have loosen too much the rule.

From the practical standpoint, in fuzzy logic, usually, the conjunction corresponds to take the minimum value, while disjunction to the maximum one. We can see that they are at the opposite ends of the ordered set of such values, so if we would like to “control” the degree of conjunction/disjunction represented by the result, then we can resort to the order statistics computed on such set. Ideally, choosing the median value could be a good trade-off between being too restrictive and too loose. However, since the next stage is based on a data-driven approach (ANN), then we can simply use the ordered set of the fuzzy results as the features to be used for the classification; it will be the next stage to automatically learn (and implicitly embed) the fuzzy rules and the defuzzification process.

The only issue left to be determined is how to segment the nut images. For this task we can safely assume that the X-ray device has been previously calibrated, so that the background is pretty uniform and, from the numerical standpoint, near to 0 (given that values are proportional to the X-ray absorption). This means that a plain global threshold is the simplest and more effective way to accomplish the goal. In synthesis, the feature extraction stage is as follows:

1) segmentation; 2) histogram computation; 3) membership inference on positive samples; 4) histogram fuzzification; 5) results sorting.

3 Classification

It is nowadays well-known that Artificial Neural Networks (ANN) are an effective way to address complex classification (and regression) problems. However, even if they are inherently non-linear, it is not always easy to design a network which is able to properly cover the problem domain, mainly due to its order or structure and the first-order structure of common multi-layer perceptrons. The usual solution is to increase the size of the network, in terms of number of connections, neurons or hidden layers, but as long as we increase the size, we increase also the risk of overfitting. Instead, designing High-Order Neural Network (HONN) has the main advantage to directly embed a much more complex behavior in the network (e.g., invariance to some differences in the data), without requiring a complex or large topology [14]. Despite the advantages, the main drawback of a pure HONN approach is that as long as we increase the order, then the number of inputs for the network increases exponentially, thus making useful to look for smarter solutions [15].

One alternative is to use functional-link networks [16], where each input is replaced by an expanded input through a set of predefined functions. This functional link is used as a pre-processing stage for the neural network, and due to its nature, it is not required anymore to have a hidden layer in the network. The great advantage is that we have a simpler and smaller network, with less weights to be trained (so that the training phase is much faster), and a smaller risk of overfitting, retaining the same abilities of generalization, given that the functional link is an orthogonal set of high-order non-linear functions [17]. The main difficulty is to identify the proper set of functions (both the family of functions and the size of the set) maximizing the trade-off between complexity and performances. As an example, in [18] an evolutionary approach has been used to find the polynomial terms making the functional link. In that case, the problem was not only to find out the right degree of the polynomials, but, more important, to select and limit the set of polynomials to be used, since the possibilities arise exponentially as long as the degree increases.

Another opportunity is to rely on a well known family of functions, and then just choose the right degree to limit the size of the functional link. Among the others, Chebyshev polynomials have received great attention, thanks to their capability of universal approximator while still retaining the advantages discussed before [19]. From the application standpoint they have proven to be a reliable solution both in regression and classification tasks [20,21,22]; in synthesis, this type of network seems to be the right choice for the problem we are addressing.

Considering at the classification dataflow, after the feature extraction stage, we have the functional link: each x_i feature of the input space is replaced by its Chebyshev expansion:

$$\{x_0, x_1, \dots, x_n\} \leftarrow \{ C_1(x_0), C_2(x_0), \dots, C_k(x_0), \\ C_1(x_1), C_2(x_1), \dots, C_k(x_1), \\ \dots, \\ C_1(x_n), C_2(x_n), \dots, C_k(x_n) \}, \quad (2)$$

where $C_j(y)$ is the Chebyshev polynomial defined by the recursive formula given by

$$\begin{aligned} C_{t+1}(y) &= 2yC_t(y) - C_{t-1}(y) \\ C_1(y) &= x \\ C_0(y) &= 1. \end{aligned} \quad (3)$$

After the functional link there is the normalization stage; it is required because even if the extracted features have values between 0 and 1, it is not guaranteed to have a similar range after the expansion. Respect to common normalization approaches, based on mean and variance, here we preferred to apply order statistics, with the following equation:

$$y_i = \left(\frac{x_i - P_i^{50}}{P_i^{(100-\delta)} - P_i^\delta} * \frac{Y_i^{max} - Y_i^{min}}{2} \right) + \frac{Y_i^{max} + Y_i^{min}}{2}, \quad (4)$$

where x_i is the i th output of the functional link, P_i^n is its n th percentile, δ determines which percentiles will be computed, and Y_i^{min} , Y_i^{max} are the desired minimum/maximum values for such variable, in order to better control the values fed into the following stage. The reasons for such kind of normalization are twofold: first, it is well known that order statistics is more robust than “classical” statistics; second, in this way we do not make any assumption regarding the distribution of the data.

The expanded normalized values are then used as input to the proper network, consisting in just one input layer and the output layer with one neuron (no hidden layer). The output is then de-normalized (with a formula similar to the previous one) to match the desired range of values and then thresholded to decide whenever the input sample belongs to the class for which the system has been trained.

The overall classification approach has been tested in three situations, to distinguish between 1) good nuts against damaged/infected ones; 2) damaged against good; 3) infected against good. The real purpose of this system is to reduce the chances to select as good a bad (damaged or infected) nut. Ideally, the first classifier would be everything we need, however, even if very low, there is a small risk of having False Positives (FP), see Table [II](#). Using the other two different classifiers and then combining the results together (with a logical disjunction), we can try to further improve the results, as the corresponding False Negative (FN) column (second and third row) in Table [II](#) would suggest.

Table 1. Performance results for different classification tests, measured with a 10×10 repeated cross-validation: recognition rate (RR), False Positives (FP), False Negatives (FN)

Test	RR (%)	FP (%)	FN (%)
Good vs damaged/infected	99.6	0.3	0.1
Damaged vs good	99.5	0.4	0.1
Infected vs good	99.8	0.2	0.0

4 Discussion

Performances are shown in Table 1 and Table 2. For the second set of results, the reported statistics are defined as follows:

$$SN = \frac{f^{TP}}{f^{TP} + f^{FN}} \times 100 \quad (5)$$

$$SP = \frac{f^{TN}}{f^{TN} + f^{FP}} \times 100 \quad (6)$$

$$ACC = \frac{f^{TP} + f^{TN}}{f^{TP} + f^{TN} + f^{FP} + f^{FN}} \times 100 \quad (7)$$


$$PR = \frac{f^{TP}}{f^{TP} + f^{FP}} \times 100, \quad (8)$$

where SN is the sensitivity, SP the specificity, ACC the accuracy, PR the precision; f^{TP} is the true positive frequency, f^{TN} the true negative frequency, f^{FP} the false positive frequency and f^{FN} the false negative frequency. Further, in Fig. 2 the Receiver Operating Characteristic (ROC) curve is reported, showing how the classifier behavior changes respect to the threshold used to discriminate between positive and negative outcomes. All the tests have been done with a 10×10 repeated cross-validation, which ensures a good reliability of those results, being able to overcome the common problem of assessing the generalization ability.


Looking at the ROC curves in Fig. 2 we can see that they are above those presented in [5,13], meaning that the proposed system is better able to discriminate between the positive samples and the negative ones. Regarding the other performances, despite the accuracy is usually the most representative information, in this type of problems the ability not to classify as good a bad nut (corresponding to the false positive rate) is more important. With the proposed solution we have an average accuracy of 99.6%, with a FP of 0.3%, if we use just one classifier, or 0.1%, with a double classifier (as discussed in the previous section). If we consider a similar solution, but with images acquired in the visible range (not X-ray) [9], then we have a FP of 1.4% (FN 2.3%). Restricting the field to the classification of X-ray images, in [13] similar features (histogram-based statistics) with a k -NN classifier have been proposed resulting in an accuracy of about

Table 2. Performance results for different classification tests, measured with a 10×10 repeated cross-validation: sensitivity (SN), specificity (SP), accuracy (ACC), precision (PR)


Test	SN (%)	SP (%)	ACC (%)	PR (%)
Good vs damaged/infected	99.8	99.5	99.6	99.1
Damaged vs good	99.8	99.3	99.5	99.1
Infected vs good	100	99.6	99.8	99.6
Average	99.9	99.4	99.6	99.3



(a) Good vs damaged/infected



(b) Damaged vs good



(c) Infected vs good

Fig. 2. Response operating characteristic curve for different classification cases

90%. In [7] the purpose was to identify the defect location, and the resulting FP is 1% using a scanned film, or 12% with a linear detector (as in our case). Further restricting the comparison to those solutions employing ANN, in [6] the sensitivity is about 90% (in our solution 99.9% on average), while in [4], which is the most similar approach (using histogram-based statistics and ANN), has both FP and FN around 1%.

5 Conclusion

A neuro-fuzzy classifier for pistachio nuts by means of their X-ray images has been proposed. The aim is to identify the good nuts and to reject the bad (damaged or infected) ones. Images are first segmented, then a set of sorted fuzzy features are computed from the histograms. During the learning stage, fuzzy membership functions are inferred from the positive samples. A Chebyshev functional-link neural network is used as the proper classifier; its result is thresholded to have the final result.

Performances have been measured with a repeated cross-validation method, to emphasize their reliability, showing a correct classification rate (good nuts against bad ones) of 99.6%, and a false positive rate of 0.3% which can be lowered to 0.1% with two combined classifiers.

References

1. Cubero, S., Aleixos, N., Moltó, E., Gómez-Sanchis, J., Blasco, J.: Advances in machine vision applications for automatic inspection and quality evaluation of fruits and vegetables. *Food and Bioprocess Technology* 4, 487–504 (2011)
2. Ruiz-Altisent, M., Ruiz-García, L., Moreda, G., Lu, R., Hernandez-Sanchez, N., Correa, E., Diezma, B., Nicolai, B., García-Ramos, J.: Sensors for product characterization and quality of specialty crops — a review. *Computers and Electronics in Agriculture* 74(2), 176–194 (2010)
3. Haff, R., Toyofuku, N.: X-ray detection of defects and contaminants in the food industry. *Sensing and Instrumentation for Food Quality and Safety* 2, 262–273 (2008)
4. Casasent, D.A., Sipe, M.A., Schatzki, T.F., Keagy, P.M., Lee, L.C.: Neural net classification of X-ray pistachio nut data. *LWT - Food Science and Technology* 31(2), 122–128 (1998)
5. Keagy, P.M., Parvin, B., Schatzki, T.F.: Machine recognition of navel orange worm damage in X-ray images of pistachio nuts. *LWT - Food Science and Technology* 29(1-2), 140–145 (1996)
6. Sim, A., Information, B.P., Division, C.S.: Invariant representation and hierarchical network for inspection of nuts from X-ray images. *International Journal of Imaging Systems and Technology* 7, 231–237 (1996)
7. Kim, S., Schatzki, T.: Detection of pinholes in almonds through X-ray imaging. *Transactions of the ASAE* 44(4), 997–1003 (2001)
8. Lu, Q., Cai, J., Li, Y., Wang, F.: Real-time nondestructive inspection of chestnuts using X-ray imaging and dynamic threshold. In: *World Automation Congress, WAC 2010*, pp. 365–368 (September 2010)

9. Pearson, T.C., Doster, M.A., Michailides, T.J.: Automated detection of pistachio defects by machine vision. *Applied Engineering in Agriculture* 17(5), 729–732 (2001)
10. Omid, M., Mahmoudi, A., Omid, M.H.: An intelligent system for sorting pistachio nut varieties. *Expert Systems with Applications* 36(9), 11528–11535 (2009)
11. Chuang, C.L., Ouyang, C.S., Lin, T.T., Yang, M.M., Yang, E.C., Huang, T.W., Kuei, C.F., Luke, A., Jiang, J.A.: Automatic X-ray quarantine scanner and pest infestation detector for agricultural products. *Computers and Electronics in Agriculture* 77(1), 41–59 (2011)
12. Marini, F.: Artificial neural networks in foodstuff analyses: Trends and perspectives— a review. *Analytica Chimica Acta* 635(2), 121–131 (2009)
13. Talukder, A., Casasent, D., Woon Lee, H., Keagy, P.M., Schatzki, T.F.: A new feature extraction method for classification of agricultural products from X-ray images. In: *Proc. SPIE* (1998)
14. Giles, C.L., Maxwell, T.: Learning, invariance, and generalization in high-order neural networks. *Appl. Opt.* 26(23), 4972–4978 (1987)
15. Artyomov, E., Yadid-Pecht, O.: Modified high-order neural network for invariant pattern recognition. *Pattern Recognition Letters* 26(6), 843–851 (2005)
16. Pao, Y.H., Takefuji, Y.: Functional-link net computing: theory, system architecture, and functionalities. *Computer* 25(5) () 25(5), 76–79 (1992)
17. Klassen, M., Pao, Y., Chen, V.: Characteristics of the functional link net: a higher order delta rule net. In: *IEEE International Conference on Neural Networks 1988*, vol. 1, pp. 507–513 (July 1988)
18. Sierra, A., Macias, J., Corbacho, F.: Evolution of functional link networks. *IEEE Transactions on Evolutionary Computation* 5(1), 54–65 (2001)
19. Lee, T.T., Jeng, J.T.: The Chebyshev-polynomials-based unified model neural networks for function approximation. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* 28(6), 925–935 (1998)
20. Patra, J., Kot, A.: Nonlinear dynamic system identification using Chebyshev functional link artificial neural networks. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* 32(4), 505–511 (2002)
21. Purwar, S., Kar, I., Jha, A.: On-line system identification of complex systems using Chebyshev neural networks. *Applied Soft Computing* 7(1), 364–372 (2007)
22. Misra, B.B., Dehuri, S.: Functional link artificial neural network for classification task in data mining. *Journal of Computer Science* 3, 948–955 (2007)

Part III

Applications

LVQ-Based Hand Gesture Recognition Using a Data Glove

Francesco Camastra* and Domenico De Felice

Department of Applied Science, University of Naples Parthenope,
Centro Direzionale Isola C4, 80143 Naples, Italy
francesco.camastra@uniparthenope.it, domenico.defelice@alice.it

Abstract. This paper presents a real-time hand gesture recognizer based on a *Learning Vector Quantization (LVQ)* classifier. The recognizer is formed by two modules. The first module, mainly composed of a data glove, performs the feature extraction. The second module, the classifier, is performed by means of LVQ. The recognizer, tested on a dataset of 3900 hand gestures, performed by people of different gender and physique, has shown very high recognition rate.

1 Introduction

Gesture is one of the means that humans use to send informations. In general, the information amount conveyed by gesture increases when the information quantity sent by the human voice decreases [1]. Besides, the hand gesture for some people, e.g., the disabled people, is one among main means for sending information.

The aim of this work is the development of a real-time hand gesture recognizer. The real-time requirement is motivated by associating to the gesture recognition the performing of an action, for instance the opening of a multimedia presentation, the starting of an internet browser and other similar actions.

The paper presents a real-time hand gesture recognition system based on *Learning Vector Quantization (LVQ)* classifier that can work on a netbook. The system is formed by two modules. The first module, a feature extractor, is performed by a *data glove* [2] and represents the hand gesture by means of a five-dimensional feature vector. The second module, the classifier, is performed by LVQ.

The paper is organized as follows: the feature extraction module is discussed in Section 2; a review of Learning Vector Quantization is provided in Section 3; Section 4 reports some experimental results; in Section 5 some conclusions are drawn.

* Corresponding author.

2 Feature Extraction

Several approaches were proposed for gesture recognition [3,4]. In gesture recognition, it is common to use a camera in combination with an image recognition system. These systems have the disadvantages that the image/gesture recognition is very sensitive to illumination, hand position and hand orientation [5]. In order to avoid these problem we decided to use an approach borrowed by Virtual Reality [2] applications where the movements of the hands of people are tracked asking them to wear data gloves [6]. A data glove is a particular glove that has sensors, typically magnetic or in optic fiber, inside that allow the track the movement of the fingers of hand.

Having said that, the feature extraction module of our system is essentially composed of the DG5 Vhand 2.0¹ data glove, shown in figure 1, that has to be worn by the person, whose gesture has to be recognized. The data glove, whose updating time of about 20 ms, measures the flexure of each finger, providing a real value between 0 and 100, that represents the finger complete closure and opening, respectively. To get more robustness to the system, the flexure for each finger is obtained averaging five consecutive measures. Therefore a new feature is generated every 100 ms. Moreover, in our experiments, considering thousands of gestures performed by people of different gender and physique, the flexure measure of finger takes value in a narrower range, typically between 10 and 70. Therefore, the values provided by the data gloves are subsequently normalized between 0 and 1, so that the values of 10 and 70 corresponds to 0 and 1, respectively. Summarizing the hand gesture is represented by means of a vector of five normalized real values generated every 100 ms.



Fig. 1. DG5 VHand 2.0 Data Glove

¹ DG5 Vhand 2.0 is a registered trademark of DGTech Engineering Solutions.

3 The Classifier

In recent years *Support Vector Machine (SVM)* [7,8,9] has been one of the most effective classification algorithm. Since SVM is a binary classifier, if we want to use SVM when the number of classes K is larger than two it is necessary to use specific strategies. The simplest strategy, *one-versus-all method* [10], requires an ensemble of K SVM classifiers, i.e., one classifier for each class against all the other classes. If we use other strategies the number of classifiers increases to $K(K-1)/2$. Therefore, SVMs require computational resources that are not compatible with a real-time recognizer.

3.1 Learning Vector Quantization

Having said that, we have chosen *Learning Vector Quantization (LVQ)* [11] as classifier in the gesture recognizer since it requires moderate computational resources. In last years LVQ was successfully applied in many different fields such as handwritten digit [12] and cursive character recognition [13], the identification of the Head-and-Shoulders patterns in the financial time series [14], the analysis of sensor array data [15] and automatic recognition of whale calls for real-time monitoring [16].

We pass to describe LVQ and first fix the notation. Let $\mathcal{D} = \{\mathbf{x}^1, \dots, \mathbf{x}^\ell\} \subseteq \mathbb{R}^N$ be a data set. We call *codebook* the set $M = \{\mathbf{m}^1, \dots, \mathbf{m}^P\} \subset \mathbb{R}^N$ and $P \ll \ell$. Vector quantization aims to yield codebooks that represent as much as possible the input data \mathcal{D} .

LVQ is a supervised version of vector quantization and generates codebook vectors (*codevectors*) to produce *near-optimal decision boundaries* [17]. LVQ consists of the application of a few different learning techniques, namely LVQ1, LVQ2 and LVQ3.

LVQ1 uses for classification the nearest-neighbour decision rule; it chooses the class of the nearest codebook vector. LVQ1 learning is performed in the following way. We denote with \mathbf{m}_t^c the value of \mathbf{m}^c at time t and with $\mathcal{C}(\mathbf{x})$ and $\mathcal{C}(\mathbf{m}_t^c)$ the class of \mathbf{x} and \mathbf{m}_t^c , respectively. If \mathbf{m}_t^c is the nearest codevector to the input vector \mathbf{x} , then

$$\left\{ \begin{array}{l} \mathbf{m}_{t+1}^c \leftarrow \mathbf{m}_t^c + \alpha_t [\mathbf{x} - \mathbf{m}_t^c] \text{ if } \mathcal{C}(\mathbf{x}) = \mathcal{C}(\mathbf{m}_t^c) \\ \mathbf{m}_{t+1}^c \leftarrow \mathbf{m}_t^c - \alpha_t [\mathbf{x} - \mathbf{m}_t^c] \text{ if } \mathcal{C}(\mathbf{x}) \neq \mathcal{C}(\mathbf{m}_t^c) \end{array} \right\} \quad (1)$$

where α_t is the learning rate at time t .

The other codevectors are not changed, i.e.

$$\mathbf{m}_{t+1}^i \leftarrow \mathbf{m}_t^i \text{ if } i \neq c \quad (2)$$

In our experiments, we used a particular version of LVQ1, that is the *Optimized-learning-rate (OLVQ1)* [17], a version of the model that provides a different learning rate for each codebook vector. We pass to describe OVLQ1. If \mathbf{m}_t^c is the closest codevector to $\mathbf{x}(t)$, i.e., the input vector at time t , then the equation (1) holds. The equation (1) can be rewritten in the following way:

$$\mathbf{m}^c(t+1) \leftarrow [1 - s(t)\alpha_c(t)]\mathbf{m}^c(t) + s(t)\alpha_c(t)\mathbf{x}(t) \quad (3)$$

where $s(t) = 1$ if the classification is correct and $s(t) = -1$ otherwise.

We can observe that the codevector \mathbf{m}_t^c is statistically independent of $\mathbf{x}(t)$. Therefore the statistical accuracy of the learned codebook is optimal if the effects of the corrections made at different times are of equal weight. We note that $\mathbf{m}^c(t+1)$ keeps a trace from $\mathbf{x}(t)$ through the last term in the equation (3) and traces from the earlier $\mathbf{x}(t')$, with $t' = 1, \dots, t-1$, through \mathbf{m}_t^c . The absolute magnitude of the last trace from $\mathbf{x}(t)$ is scaled down by the factor $\alpha_c(t)$ whereas the trace from $\mathbf{x}(t-1)$ is scaled down by $[1 - s(t)\alpha_c(t)]\alpha_c(t-1)$. If we assume that the two scaling are equal, we have:

$$\alpha_c(t) = [1 - s(t)\alpha_c(t)]\alpha_c(t-1). \quad (4)$$

If this condition holds for all t , by induction it can be shown that the traces stored up to time t from all the previous t' will be scaled down by an equal amount at the end. Therefore the optimal value of the learning rate $\alpha(t)$ is given by the following update rule:

$$\left\{ \begin{array}{l} \alpha_c(t) \leftarrow \frac{\alpha_c(t-1)}{1+\alpha_c(t-1)} \text{ if } \mathcal{C}(\mathbf{x}) = \mathcal{C}(\mathbf{m}_t^c) \\ \alpha_c(t) \leftarrow \frac{\alpha_c(t-1)}{1-\alpha_c(t-1)} \text{ if } \mathcal{C}(\mathbf{x}) \neq \mathcal{C}(\mathbf{m}_t^c) \end{array} \right\} \quad (5)$$

Since both LVQ1 and OLVQ1 tend to push codevectors away from the decision surfaces of the *Bayes rule* [18], it is necessary to apply to the codebook generated a successive learning technique called LVQ2.

LVQ2 tries harder to approximate the Bayes rule by pairwise adjustments of codevectors belonging to adjacent classes. If \mathbf{m}^s and \mathbf{m}^p are nearest neighbours of different classes and the input vector \mathbf{x} , belonging to the \mathbf{m}^s class, is closer to \mathbf{m}^p and falls into a zone of values called *window* (that is defined around the midplane of \mathbf{m}^s and \mathbf{m}^p), the following rule is applied:

$$\left\{ \begin{array}{l} \mathbf{m}_{t+1}^s \leftarrow \mathbf{m}_t^s + \alpha_t[\mathbf{x} - \mathbf{m}_t^s] \\ \mathbf{m}_{t+1}^p \leftarrow \mathbf{m}_t^p - \alpha_t[\mathbf{x} - \mathbf{m}_t^p] \end{array} \right\} \quad (6)$$

It can be shown [11] that the LVQ2 rule produces an instable dynamics. To prevent this behavior as far as possible, the window w within the adaptation rule takes place must be chosen carefully. Moreover, the hypothesis margin μ of the classifier [19] is given by:

$$\mu = \frac{\|\mathbf{x} - \mathbf{m}^s\| - \|\mathbf{x} - \mathbf{m}^p\|}{2}. \quad (7)$$

Hence LVQ2 can be seen as a classifier which aims at structural risk minimization during training, comparable to *Support Vector Machines* [20]. Therefore a very good generalization ability of LVQ2 can be expected also for high dimensional data.

In order to overcome the LVQ2 stability problems, Kohonen proposed a further algorithm (LVQ3). If \mathbf{m}^i and \mathbf{m}^j are the two closest codevectors to input \mathbf{x} and \mathbf{x} falls in the window, the following rule is applied:

$$\left\{ \begin{array}{l} \mathbf{m}_{t+1}^i \leftarrow \mathbf{m}_t^i \quad \text{if } \mathcal{C}(\mathbf{m}^i) \neq \mathcal{C}(\mathbf{x}) \wedge \mathcal{C}(\mathbf{m}^j) \neq \mathcal{C}(\mathbf{x}) \\ \mathbf{m}_{t+1}^j \leftarrow \mathbf{m}_t^j \quad \text{if } \mathcal{C}(\mathbf{m}^i) \neq \mathcal{C}(\mathbf{x}) \wedge \mathcal{C}(\mathbf{m}^j) \neq \mathcal{C}(\mathbf{x}) \\ \mathbf{m}_{t+1}^i \leftarrow \mathbf{m}_t^i - \alpha_t[\mathbf{x}_t - \mathbf{m}_t^i] \quad \text{if } \mathcal{C}(\mathbf{m}^i) \neq \mathcal{C}(\mathbf{x}) \wedge \mathcal{C}(\mathbf{m}^j) = \mathcal{C}(\mathbf{x}) \\ \mathbf{m}_{t+1}^j \leftarrow \mathbf{m}_t^j + \alpha_t[\mathbf{x}_t - \mathbf{m}_t^j] \quad \text{if } \mathcal{C}(\mathbf{m}^i) \neq \mathcal{C}(\mathbf{x}) \wedge \mathcal{C}(\mathbf{m}^j) = \mathcal{C}(\mathbf{x}) \\ \mathbf{m}_{t+1}^i \leftarrow \mathbf{m}_t^i + \alpha_t[\mathbf{x}_t - \mathbf{m}_t^i] \quad \text{if } \mathcal{C}(\mathbf{m}^i) = \mathcal{C}(\mathbf{x}) \wedge \mathcal{C}(\mathbf{m}^j) \neq \mathcal{C}(\mathbf{x}) \\ \mathbf{m}_{t+1}^j \leftarrow \mathbf{m}_t^j - \alpha_t[\mathbf{x}_t - \mathbf{m}_t^j] \quad \text{if } \mathcal{C}(\mathbf{m}^i) = \mathcal{C}(\mathbf{x}) \wedge \mathcal{C}(\mathbf{m}^j) \neq \mathcal{C}(\mathbf{x}) \\ \mathbf{m}_{t+1}^i \leftarrow \mathbf{m}_t^i + \epsilon\alpha_t[\mathbf{x}_t - \mathbf{m}_t^i] \quad \text{if } \mathcal{C}(\mathbf{m}^i) = \mathcal{C}(\mathbf{m}^j) = \mathcal{C}(\mathbf{x}) \\ \mathbf{m}_{t+1}^j \leftarrow \mathbf{m}_t^j + \epsilon\alpha_t[\mathbf{x}_t - \mathbf{m}_t^j] \quad \text{if } \mathcal{C}(\mathbf{m}^i) = \mathcal{C}(\mathbf{m}^j) = \mathcal{C}(\mathbf{x}) \end{array} \right\}, \quad (8)$$

where $\epsilon \in [0, 1]$ is a fixed parameter.

3.2 Classification with Rejection

In practical applications, it is usually associated to the recognition of a gesture the performing of a given action, e.g., the starting of a multimedia performance. In this applicative scenario, it is desiderable that the classifier recognizes a gesture, i.e., classifies, only when the probability of making a mistake is negligible. When the probability of making is not negligible, the classifier has to *reject* the gesture, i.e., it does not classify.

We implemented a rejection scheme in the classifier in the following way. Let d be the Euclidean distance between the input \mathbf{x} and the closest codevector \mathbf{m}^c , the following rule is applied:

$$\text{If } d \leq \rho \text{ then classify else reject} \quad (9)$$

where ρ is a parameter that manages the trade-off between error and rejection.

4 Experimental Result

To validate the recognizer we selected 13 gestures, invariant by rotation and translation. We associated to each gesture a symbol, a letter or a digit, as shown in figure 2. We collected a database of 7800 right hand gestures, performed by people of different gender and physique. The database was splitted with a random into two equal parts, training and test set each containing 3900 gestures. The number of classes used in the experiments was 13, namely the number of the different gestures in our database. In our experiments the three learning techniques, i.e. LVQ1, LVQ2 and LVQ3, were applied. We trained several LVQ nets by specifying different combinations of learning parameters, i.e. different learning rates for LVQ1, LVQ2, LVQ3 and various total number of codevectors.

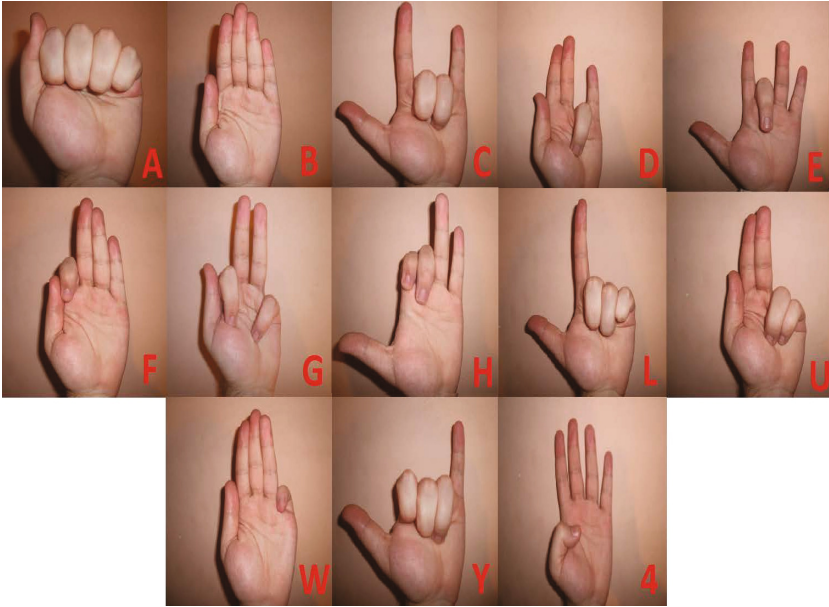


Fig. 2. Gestures represented in the database

Table 1. Recognition rates on the test set, in absence of rejection, for LVQ classifiers

Algorithm	Correct Classification Rate
knn	92.56 %
LVQ1 + LVQ2	99.31 %
LVQ1 + LVQ3	99.31 %

The best LVQ net was selected by means of *10-fold cross-validation*² [21][22]. LVQ trials were performed using *LVQ-pak* [23] software package. We compared LVQ against the *K-nearest-neighbor* (*Knn*) classifier, available in the *LVQ-pak* software package. In Table 1, for different classifiers, the performances on the test set, measured in terms of recognition rate in absence of rejection, are reported.

Our best result in terms of recognition rate is *99.31%*. The confusion matrix for LVQ classifier, on the test set, is shown in table 2.

To our best classifier, i.e., LVQ1+LVQ3, the rejection rule, described in Section 5.2, was applied. The results obtained for different values of the rejection

² In *K-fold cross-validation* the training set is partitioned into *K* subsets. Of the *K* subsets, a single subset is retained as the validation data for testing the model, and the remaining *K - 1* subsets are used as training data. The cross-validation process is then repeated *K* times, i.e. the folds, with each of the *K* subsets used exactly once as the validation data. The *K* results from the folds then can be averaged to produce a single estimation.

Table 2. The confusion matrix for LVQ1+LVQ3 classifier on the test set The values are expressed in terms of percentage rates

	A	B	C	D	E	F	G	H	L	U	W	Y	4
A	100	0	0	0	0	0	0	0	0	0	0	0	0
B	0	100	0	0	0	0	0	0	0	0	0	0	0
C	0	0	97.33	2.67	0	0	0	0	0	0	0	0	0
D	0	0	0	99.33	0	0	0	0	0	0.67	0	0	0
E	0	0.33	0	0	99.67	0	0	0	0	0	0	0	0
F	0	0	0	0	0	100	0	0	0	0	0	0	0
G	0	0	0	0	0	0	97.67	0	0	0	2.33	0	0
H	0	0	0	0	0	0	0	100	0	0	0	0	0
L	0	0	0	0	0	0	0	0	100	0	0	0	0
U	0	0	0	0	0	0	0	0	0	97	3	0	0
W	0	0	0	0	0	0	0	0	0	0	100	0	0
Y	0	0	0	0	0	0	0	0	0	0	0	100	0
4	0	0	0	0	0	0	0	0	0	0	0	0	100

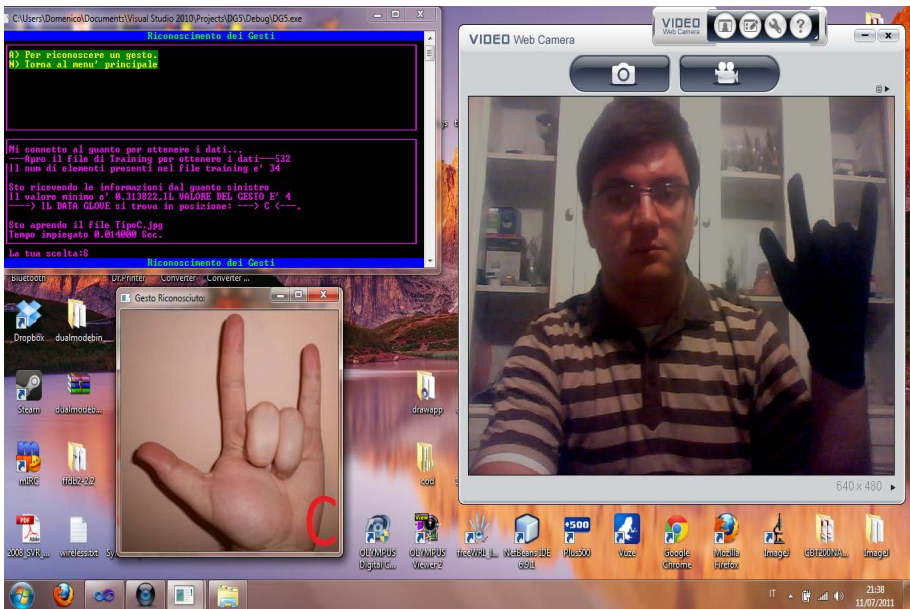


Fig. 3. The Hand Gesture Recognizer

threshold ρ are shown in Table 3. Asking to the recognizer a negligible error, e.g., less than 0.4%, the recognizer can still guarantee a high correct recognition rate, i.e., larger than 97.0%.

Table 3. Correct, Error and Reject rates on the test set of LVQ1+LVQ3 for different rejection threshold ρ

ρ	Correct	Error	Reject
0.33	99.31 %	0.69 %	0.00 %
0.30	97.36 %	0.38 %	2.26 %
0.25	92.38 %	0.23 %	7.39 %
0.20	84.62 %	0.20 %	15.20 %
0.10	41.33 %	0.05 %	58.62 %
0.0992	40.82 %	0.00 %	59.18 %

A demo³ of the recognizer, where the performing of a Powerpoint⁴ command is associated to the recognition of a sequence of gestures, is used by the first author, as data glove demonstrator, in his B. Sc. course of *Virtual Reality*.

The system, shown in figure 3, is implemented in C++ under 64 bit Windows Vista Microsoft and Visual Studio 2010 on a Netbook with Intel Celeron⁵ Processor SU2300 1.20 Ghz, 3 GB RAM and requires 140 CPU⁶ ms to recognize a single gesture. Since the feature extraction requires 100 ms, as underlined in Section 2, the remaining 40 ms are spent for the classification and performing the associated Powerpoint command.

5 Conclusions

In this paper we have described real-time hand gesture recognition system based on a Learning Vector Quantization classifier. The system is formed by two modules. The first module, mainly composed of a data glove, performs the feature extraction and represents the hand gesture by a 5-dimensional feature vector. The second module, the classifier, is performed by means of Learning Vector Quantization. The recognizer, tested on a dataset of 3900 hand gestures, performed by people of different gender and physique, has shown a recognition rate, in absence of rejection, larger than 99%. The system implemented on a netbook requires an average time of 140 CPU ms to recognize a hand gesture. Furthermore, a version of the system where the recognition of a sequence of gestures is associated the carrying out of a multimedia presentation command has been developed. In the next future we plan to investigate its usage in classrooms as tool for helping disabled students.

Acknowledgments. Firstly, the authors wish to warmly acknowledge the *Centro di Calcolo* of University of Naples Parthenope and in particular Prof. Giulio Giunta, their director, for having provided the data gloves, object of the work.

³ The demo is available at <http://dsa.uniparthenope.it/francesco.camastra>

⁴ Powerpoint, Visual Studio 2010 and Windows Vista are registered trademarks by Microsoft Corp.

⁵ Celeron is a registered trademark by Intel Corp.

⁶ CPU time has been measured using the proper C++ function.

The authors wish to thank the anonymous reviewers for their valuable comments.

Part of the work was developed by Domenico De Felice, under the supervision of Francesco Camastra, as final dissertation project for B. Sc. in Computer Science at University of Naples Parthenope.

References

1. Kendon, A.: How gestures can become like words. In: *Crosscultural Perspectives in Nonverbal Communication*, Toronto, Hogrefe, pp. 131–141 (1988)
2. Burdea, G., Coiffet, P.: *Virtual Reality Technology*. John-Wiley & Sons, New York (2003)
3. Mitra, S., Acharya, T.: Gesture recognition: A survey. *IEEE Transactions on Systems, Man and Cybernetics, Part C: Applications and Reviews* 37(3), 311–324 (2007)
4. Chaudhary, A., Raheja, J.L., Das, K., Raheja, S.: A Survey on Hand Gesture Recognition in Context of Soft Computing. In: Meghanathan, N., Kaushik, B.K., Nagamalai, D. (eds.) *CCSIT 2011, Part III. CCIS*, vol. 133, pp. 46–55. Springer, Heidelberg (2011)
5. Weissmann, J., Salomon, R.: Gesture recognition for virtual reality applications using data gloves and neural networks. In: *Proceedings of the IJCNN 1999*, pp. 2043–2046. IEEE Press (1999)
6. Dipietro, L., Sabatini, A., Dario, P.: A survey of glove-based systems and their applications. *IEEE Transactions on Systems, Man and Cybernetics* 38(4), 461–482 (2008)
7. Cortes, C., Vapnik, V.: Support vector networks. *Machine Learning* 20, 1–25 (1995)
8. Schölkopf, B., Smola, A.: *Learning with Kernels*. MIT Press, Cambridge (2002)
9. Shawe-Taylor, J., Cristianini, N.: *Kernels Methods for Pattern Analysis*. Cambridge University Press (2004)
10. Herbrich, R.: *Learning Kernel Classifiers*. MIT Press, Cambridge (2004)
11. Kohonen, T.: Learning vector quantization. In: *The Handbook of Brain Theory and Neural Networks*, pp. 537–540. MIT Press (1995)
12. Ho, T.: Recognition of handwritten digits by combining independent learning vector quantizations. In: *Proceedings of the Second International Conference on Document Analysis and Recognition*, pp. 818–821. IEEE (1993)
13. Camastra, F., Vinciarelli, A.: Cursive character recognition by learning vector quantization. *Pattern Recognition Letters* 22(6-7), 625–629 (2001)
14. Zapranis, A., Tsinaslanidis, P.: Identification of the Head-and-Shoulders Technical Analysis Pattern with Neural Networks. In: Diamantaras, K., Duch, W., Iliadis, L.S. (eds.) *ICANN 2010, Part III. LNCS*, vol. 6354, pp. 130–136. Springer, Heidelberg (2010)
15. Ciosek, P., Wróblewski, W.: The analysis of sensor array data with various pattern recognition techniques. *Sensors and Actuators B: Chemical* 114(1), 85–93 (2006)
16. Mouy, X., Bahoura, M., Simard, Y.: Automatic recognition of fin and blue whale calls for real-time monitoring in the st. lawrence. *Journal of the Acoustical Society of America* 126(6), 2918–2928 (2009)
17. Kohonen, T.: *Self-Organizing Maps*. Springer, Berlin (1997)
18. Duda, R., Hart, P., Stork, D.: *Pattern Classification*. John-Wiley & Sons, New York (2001)

19. Crammer, K., Gilad-Bachrach, R., Navot, A., Tishby, N.: Margin analysis of the lvq algorithm. In: *Advances in Neural Information Processing Systems*, pp. 109–114. MIT Press (2002)
20. Vapnik, V.: *Statistical Learning Theory*. John Wiley and Sons, New York (1998)
21. Stone, M.: Cross-validatory choice and assessment of statistical prediction. *Journal of the Royal Statistical Society* 36(1), 111–147 (1974)
22. Hastie, T., Tibshirani, R., Friedman, R.: *The Elements of Statistical Learning*. Springer (2001)
23. Kohonen, T., Hynninen, J., Kangas, J., Laaksonen, J., Torkkola, K.: Lvq-pak: The learning vector quantization program package. Technical Report A30, Helsinki University of Technology, Laboratory of Computer and Information Science (1996)

Investigation of Single Nucleotide Polymorphisms Associated to Familial Combined Hyperlipidemia with Random Forests

Antonino Staiano¹, Maria Donata Di Taranto², Elena Bloise³,
Maria Nicoletta D'Agostino^{3,4}, Antonietta D'Angelo^{3,4}, Gennaro Marotta⁵,
Marco Gentile⁵, Fabrizio Jossa⁵, Arcangelo Iannuzzi⁶,
Paolo Rubba⁵, and Giuliana Fortunato^{3,4}

¹ Dipartimento di Scienze Applicate, Università di Napoli "Parthenope",
Isola C4, Centro Direzionale, 80143, Napoli, Italy

antonino.staiano@uniparthenope.it

² IRCCS Fondazione SDN, Napoli, Italy

³ CEINGE S.C.a r.l. Biotecnologie Avanzate, Napoli, Italy

⁴ Dipartimento di Biochimica e Biotecnologie Mediche, Università degli Studi di
Napoli Federico II, Napoli, Italy

⁵ Dipartimento di Medicina Clinica e Sperimentale, Università degli Studi di Napoli
Federico II, Napoli, Italy

⁶ Centro per la diagnosi e la terapia delle dislipidemie familiari, Azienda Ospedaliera
A. Cardarelli, Napoli, Italy

Abstract. Single nucleotide polymorphisms (SNPs) are the foremost part of many genome association studies. Selecting a subset of SNPs that is sufficiently informative but still small enough to reduce the genotyping overhead is an important step towards disease-gene association. In this work, a Random Forest (RF) approach to informative SNPs selection in Familial Combined Hyperlipidemia (FCH) is proposed. FCH is the most common form of familial hyperlipidemia. Affected patients have elevated levels of plasma triglycerides and/or total cholesterol and show increased risk of premature coronary heart disease. In order to identify susceptibility markers for FCH we perform the analysis of 21 SNPs in ten genes associated with high cardiovascular risk. RF appears to be a useful technique in identifying gene polymorphisms involved in FCH: the identified SNPs confirmed some variants in a couple of genes as genetic markers of FCH as proved in various studies in scientific literature and lead us to report for the first time a further gene association with FCH. This result could be promising and encourages to further investigate on the role of the identified gene in the development of FCH phenotype.

Keywords: Machine learning, feature selection, random forests, genome association, single nucleotide polymorphisms, familial combined hyperlipidemia.

1 Introduction

Genome and genome-wide association studies (GA) have revolutionized genetics making it possible the search for genetic factors that influence common complex traits and the characterization of the effects of those factors [5]. Single nucleotide polymorphisms (SNPs) are at the forefront of many disease-gene association studies. SNPs are mutations at a single nucleotide position that occurred during evolution and were passed on through heredity, accounting for most of the genetic variation among different individuals. The number of SNPs has been estimated to be around seven million in the human genome [8]. Thus, in GA studies, selecting a subset of SNPs that is sufficiently informative but still small enough to reduce the genotyping overhead is an important step towards disease-gene associations [15]. Machine learning techniques [1] have been applied with ever growing interest and success in GA for feature selection purposes. The success is mostly due to their natural suitability in assessing many loci jointly which may be more informative than testing associations at individual markers and to deal with large amount of data. In fact, most of the studies, in this context, have used a single-locus analysis strategy, in which each variant is tested individually for association with a specific phenotype. However, as often stated in literature, a reason for the lack of success in genetic studies of complex disease is the existence of interactions between loci. If a genetic factor functions primarily through a complex mechanism that involves multiple other genes and, possibly, environmental factors, the effect might be missed if the gene is examined in isolation without allowing for its potential interactions with these other unknown factors [5]. A huge amount of diverse approaches have been developed in the last decade to informative SNPs selection (see [5], [15] and references therein for an overview) for different disease-gene associations, ranging from multiple sclerosis [3] to chick mortality [11]. Nonetheless, the complexity of GA studies and the peculiarity of each disease-gene association make it impossible to figure out the “best” technique for each application. Random Forest [2], an ensemble of a machine learning techniques based on Decision Trees (DTs), is one candidate model to informative SNP selection. DT [6], is a very popular technique in the machine learning community. DT is efficient from a computational point of view, works well with both categorical or continuous data, copes fine with missing values thanks to the notion of surrogate variables, and its resulting structure after training is easily interpretable even for an user with no machine learning background. On the other hand, DT suffers from high variance. Often a small change in the data can result in a very different tree structure, making interpretation somewhat arbitrary. This motivates the use of RF, which grows a set of many DTs and combines their results through bagging [6] resulting in a very effective reduction of variance in each tree in the ensemble. Aim of this work is a RF informative SNPs selection in Familial Combined Hyperlipidemia (FCH) [12]. FCH is the most common form of familial hyperlipidemia [9]. Affected patients have elevated levels of plasma triglycerides and/or total cholesterol and show increased risk of premature coronary heart disease [13]. In order to identify susceptibility markers for FCH we perform the analysis of 21 SNPs in 10 genes associated

with high cardiovascular risk: *Lipoprotein Lipase* (LPL), *Cholesteryl ester transfer protein* (CETP), *HMG-CoA reductase* (HMGCR), PCSK9, ApoA5, ApoC3, *Upstream transcription factor 1* (USF1), *Peroxisome proliferator-activated receptor gamma* (PPARG), *Gap junction alpha-4* (GJA4) and *Kinesin-like protein 6* (KIF6).

The rest of the paper is organized as follows: Section 2 describes the fundamentals of RF, while Section 3 illustrates the FCH sample data used, the experimental settings and the obtained results, respectively. Finally, conclusions in Section 4 close the paper.

2 Random Forests

The RF technique is an ensemble learning technique for conducting classification (or regression) analyses. It constructs a collection of classification trees to aggregate them into one robust classifier. The approach is based on CART (Classification and Regression Trees) [10]. Trees are constructed using rules that determine how well a split at a node (based on the values of a predictor variable) can differentiate observations with respect to the outcome variable. In the following we give an overview of CART principles and motivate the need of an ensemble of CARTs, i.e. the RF in order to better characterize the complexity of our data.

2.1 Classification and Regression Trees

A classification problem consists of a training sample of n observations on a class variable Y assuming values $1, 2, \dots, k$, and l predictor variables, X_1, \dots, X_l . The goal is to find a model for predicting the values of Y from new X values. The most classification tree construction algorithms employ a top-down heuristic search using recursive partitioning of data sample space. Starting from a heterogeneous set in terms of the variation in the outcome variable class label (i.e., case or control) of training samples at root node, each predictor (i.e., a SNP) is evaluated using a statistic to determine how well it classifies the training samples by itself. The best feature is selected to split the training samples to descendant nodes. The whole process is recursively repeated to split the descendant nodes until some pre-specified stopping criteria are met (for example, when all terminal nodes contain only cases or only controls, or when all possible SNPs have been included in a branch). The critical step in tree growing is to select the best predictor (or feature in the machine learning jargon) to split a node [4]. Most algorithms evaluate the performance of a candidate feature in separating different class labels in the training samples. The concept of impurity is usually used. Here we consider the Gini index as measure of impurity within node t :

$$I_G(t) = \sum_{i=1}^k p_i(t)(1 - p_i(t)),$$

where p_1, p_2, \dots, p_k are the proportion of samples in the k classes. Then, in binary trees, a feature and a split are chosen according to the following decrement in impurity:

$$\Delta(s, t) = I(t) - h(t_L)I(t_L) - h(t_R)I(t_R)$$

where s is a split of node t , $h(t_L)$ and $h(t_R)$ are the proportions of the samples in the left and right child nodes of node t , respectively. By recursively using the node-splitting procedure, one usually ends up with an overgrown tree (with too many descendant nodes), which produces a tree that overfits the training samples and is prone to random variations in the data. Two commonly employed strategies to overcome the overfitting are to interrupt the tree growing by a stop-splitting criterion and to apply a pruning step on the overgrown tree. The stop-splitting criterion could be either based on the node size, the node homogeneity, or elaborate criterion based on statistical testing.

2.2 Ensemble of Trees

Although tree models are easy to interpret, single tree-based analysis has its own limitations in analyzing datasets:

- The topology of a tree is usually unstable. A minor perturbation of the input training sample could result in a totally different tree model. The major reason for this instability is the hierarchical nature of the process: the effect of an error in the top split is propagated down to all of the splits below it.
- For high dimensional data a single parsimonious model is not enough to reflect the complexity in the dataset.
- A single tree may have a relatively lower accuracy in prediction with respect to other machine learning approach (e.g., Support Vector Machine (SVM) or Neural Networks (NN) [1]).

Furthermore, when one faces with small sample sizes, their inherent risk of imprecision and overfitting is higher with single tree models. Rather than using a single classification tree, an high increase in prediction accuracy can result from growing an ensemble of trees and letting them “vote” for the most popular outcome class, given a set of input variable values. This way the ensemble adjusts for the instability of the individual trees induced by small changes in the learning sample. Such ensemble approaches can also be used to provide measures of variable importance. RF is the most widely used ensemble tree approach. A RF consists of hundreds or thousands of unpruned tree each one fitted to a bootstrap sample of the same size (for example, the same number of cases and controls) from the original data. Although an individual tree in the forest is not a good model by itself, the aggregated classification has been shown to achieve much better performance than what a single tree may achieve. To construct an RF with B trees from a training dataset with n observations (cases and controls) with l features or predictors (SNPs), we employ the following steps:

1. A bootstrap sample is drawn from the training sample;
2. A classification tree is grown for the bootstrap sample. At each node, the split is selected on the basis of a randomly selected subset (much smaller than l) features. The tree is grown to full size without pruning;
3. Steps 1 and 2 are repeated B times to form a forest. The ensemble classification label is made by a majority vote of all trees in the ensemble.

Out-f-Bag Samples, Classification Accuracy and Variable Importance.

The RF algorithm does not use all training samples in the construction of the individual tree, rather it leaves a set of out-of-bag (oob) samples, which can be used to measure the forest classification accuracy [6] with no need of cross-validation or test sets. The out-of-bag observations can also be used to estimate variable importance by a permutation importance index, namely, to measure a specific feature's importance in the tree. When the b -th tree is grown ($b = 1, \dots, B$), the oob samples are passed down the tree, and the prediction accuracy is recorded. The values for the j -th feature are randomly permuted in the oob samples, and the accuracy is again computed. In this way by randomly permuting a feature values, its original association with the outcome variable is broken. When the permuted feature, together with the remaining unpermuted predictor variables, is used to predict the response, the prediction accuracy (i.e., the number of observation classified correctly) decrease substantially if the original variable was associated with the response. Thus, the larger the decrease of the prediction accuracy, the more important the variable. This process is repeated for each feature.

3 Experiments

To detect the genetic loci and their interactions that influence a phenotypic outcome (case or control), the FCH data sample has been analyzed by RF which produces, in each of its constituent tree, a graphical structure (see Fig. 1) that resembles an upside-down tree that maps the possible values of certain predictor variables (e.g., SNP genotypes) to a final expected outcome (e.g., disease status). Each node of the tree represents a predictor variable and there are edges from each node leading down to "child" nodes, in which edge correspond to a different possible value that could be taken by the variable in the parent node. A path through the tree represents a particular combination of values taken by the predictor variables that are present within the path. In this way, we found a subset of SNPs localized in different genes and their association with FCH.

We also performed classical statistical tests for the frequency comparison. Differences between groups were assessed by χ^2 test and a p -value less than 0.05 was considered significant.

3.1 The Data

After genomic DNA extraction from peripheral blood samples, the *TaqMan* assay was performed for the SNP typing. Real-Time PCR reactions were performed on

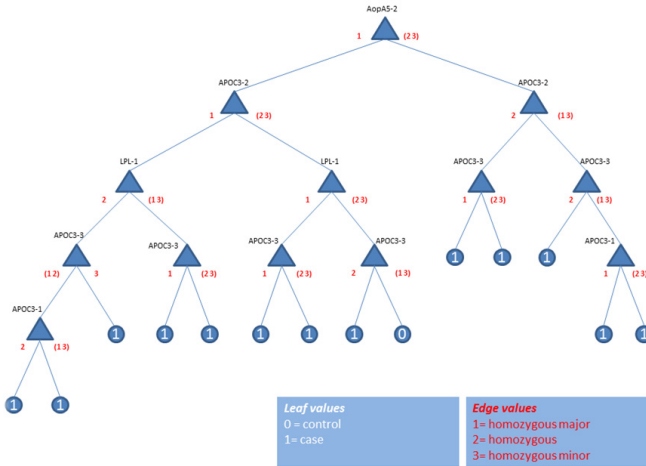


Fig. 1. An example of a single decision tree on case-control data

a ABI Prism 7900–*HT* instrument and subsequent data analysis were done with the Sequence Detection System 2.3 Software (Applied Biosystem). We enrolled 123 patients with a clinical diagnosis of FCH (after exclusion of misdiagnosis of familial hypercholesterolemia) and 142 healthy controls. Therefore, the case-control data set is constituted by 265 observations and 21 SNPs. The homozygous major, heterozygous, homozygous minor alleles have been coded by 1, 2, and 3, respectively. Since our data sample contains few missing values for few SNPs (see Fig 2), the missing genotypes were imputed as the single most likely values on the basis of the genotype frequencies.

3.2 RF Analysis

Starting from the original case-control FCH sample, RF formed by 500 to 1000 trees have been trained in order to assess the SNP importance by the permutation importance measure, as Fig. 3 depicts. It is worth noting that, while in Fig. 3 the results for a RF constituted by 1000 trees are reported, a RF with 500 trees would suffice to obtain a similar and stable results.

Successively, the RF was used jointly with a sequential backward selection procedure [1] as a wrapper feature selection method [15]. Specifically, the sequential backward selection starts from a subset of eight SNPs (determined by imposing, empirically, a threshold at 0.1 to the permutation importance measure) and sequentially removes a feature from the full candidate set until the removal of further features increase the RF misclassification rate. The RF prediction accuracy as well as the sequential backward selection are computed on the oob samples.




Fig. 2. Fraction of missing values for each SNP in the case-control data set




Fig. 3. SNP importance values from permutation importance measure

3.3 Results

The identified optimal feature subset contains 5 SNPs and lead to a RF prediction accuracy of 68.68%: *ApoA5* (*S19W*), *ApoA5* (-1131 T> C), *PCSK9* (25467958 T> C), *GJA4* (P319S), *USF1* (9996 G>A).

The sensitivity and specificity were 66.67% and 71.83%, respectively. Fig. 4 depicts the the oob error with respect to the number of trees of the RF and clearly shows that a smaller RF size is enough to reach comparable results.




Fig. 4. RF oob error vs number of trees (left) and RF ROC curve (right)

Fig. 4 shows the ROC curve obtained by the RF. We also performed standard frequency comparisons to confirm results and we found that genotypes of the above-mentioned SNPs have a different frequency between FCH and controls.

4 Conclusions

RF is an effective methodology for machine learning applications. It is widely used in the research community both for classification and feature selection purposes, due to its some nice properties such as the oob samples which can be exploited to assess the RF performance with no needs of a testing set or a cross-validation procedure, and to evaluate the feature importance through the

permutation importance index. Moreover, an appealing characteristic is the reduced computational costs when compared to other calculus intensive machine learning approach (such as neural networks or support vector machines). Several works in GA literature employ RF as an analysis tool ([14], [16]). We applied a RF analysis to identify a set of SNPs related to Familial Combined Hyperlipidemia: it has been used for a first evaluation of SNP importance through the permutation importance index and then, in conjunction with sequential backward elimination, as a wrapper feature selection procedure that can identify an optimal set on SNPs able to correctly classify 68.68% of subjects in the FCH sample. It is worth stressing that no other employed, for comparison purposes, machine learning approach such as NN and SVM could obtain a prediction accuracy above 65% on the same feature subset. Although the analysis is limited to twenty-one SNPs in ten genes already associated with FCH, the literature reports discordant data about these associations. This motivates our work whose aim was to test the associations with a method of analysis different from the routinely used in case-control studies to confirm the data. So doing, we confirmed the role of the variants *S19W* and $-1131T > C$ in *ApoA5* gene and $9996G > A$ in *USF1* gene as genetic markers of FCH [9], [12], [13]. Moreover, we reported for the first time association of *GJA4* gene and its variant *P319S* with FCH. Definitely, while obtained on a relative small sample size, this result could be promising and encourages us to further investigate on the role of this gene in the development of FCH phenotype. Indeed, FCH has polygenic and multifactorial basis not yet completely clarified [13]. Previous linkage and association studies indicated various genes associated to altered lipid or lipoprotein phenotype, most of them being involved in lipid metabolism [7]. Since the polygenic basis of the FCH disease, the combined study of different genetic variations allow to better characterize patients and to identify high-risk subjects. Obviously, the discovery of new variants requires the use of different genetic screening such as wide-genome scan or linkage analysis rather than the small-scale exploratory investigation presented in this work.

Acknowledgements. CEINGE Convenzione Regione Campania, DGRC 1901/2009 and IRCCS Fondazione SDN.

References

1. Bishop, C.M.: Pattern Recognition and Machine Learning. Springer (2006)
2. Breiman, L.: Random Forests. *Machine Learning* 45(1) (2001)
3. Calcagno, G., Staiano, A., Fortunato, G., Brescia-Morra, V., Salvatore, E., Liguori, R., Capone, S., Filla, A., Longo, G., Sacchetti, L.: A Multilayer Perceptron Neural Network-based Approach for the Identification of Responsiveness to Interferon Therapy in Multiple Sclerosis Patients. *Information Sciences* 180(21), 4153–4163 (2010)
4. Chen, X., Wang, M., Zhang, H.: The use of Classification Trees for Bioinformatics. *WIREs Data Mining Knowl. Discov.* 1, 55–63 (2011)

5. Cordell, H.J.: Detecting Gene-gene Interactions that Underlie Human Diseases. *Nat. Revs. Gen.* 10, 392–404 (2009)
6. Hastie, J., Tibshirani, R., Friedman, J.: *The Elements of Statistical Learning - Data Mining, Inference and Prediction*. Springer (2009)
7. Hegele, R.A.: Plasma Lipoproteins: Genetic Influences and Clinical Implications. *Nat. Revs. Gen.* 10(2), 109–121 (2009)
8. Kruglyek, L., Nickerson, D.A.: Variation in the Spice of Life. *Nat. Genet.* 27, 234–236 (2001)
9. Liu, Z.-K., Hu, M., Baumb, L., Thomas, G.N., Tomlinson, B.: Associations of Polymorphisms in the Apolipoprotein a1/c3/a4/a5 Gene Cluster with Familial Combined Hyperlipidaemia in Hong Kong Chinese. *Atherosclerosis* 208, 427–432 (2010)
10. Loh, W.-Y.: Classification and Regression Trees. *WIREs Data Mining Knowl. Discov.* 1, 14–23 (2011)
11. Long, N., Gianola, D., Rosa, G.J.M., Weigel, K.A., Avendano, S.: Comparison of Classification Methods for Detecting Associations between SNPs and Chick Mortality. *Genet. Sel. Evol.* 41(18) (2009)
12. Mar, R., Pajukanta, P., Allayee, H., Groenendijk, M., Dallinga-Thie, G., Krauss, R.M., Sinsheimer, J.S., Cantor, R.M., de Bruin, T.W.A., Lusis, A.J.: Association of the Apolipoprotein a1/c3/a4/a5 Gene Cluster with Triglyceride Levels and Ldl Particle Size in Familial Combined Hyperlipidemia. *Circulation Reserch* 94(7), 993–999 (2004)
13. Naukkarinen, J., Ehnholm, C., Peltonen, L.: Genetics of Familial Combined Hyperlipidemia. *Current Opinion in Lipidology* 17, 285–290 (2006)
14. Obulkasim, A., Meijer, G.A., van de Wiel, M.A.: Stepwise Classification of Cancer Samples Using Clinical and Molecular Data. *BMC Bioinformatics* 12(422) (2011)
15. Saeys, Y., Inza, I., Larranaga, P.: A Review of Feature Selection Techniques in Bioinformatics. *Bioinformatics* 23(19), 2507–2517 (2007)
16. Yoshida, M., Koike, A.: SNPinterforest: A New Method for Detecting Epistatic Interactions. *BMC Bioinformatics* 12(469) (2011)

A Neural Procedure for Gene Function Prediction

Marco Frasca, Alberto Bertoni, and Andrea Sion

Dipartimento di Scienze dell'Informazione
Università degli Studi di Milano
via Comelico 39 20135 Milano Italy
{frasca,bertoni}@dsi.unimi.it, andrea.sion@unimi.it

Abstract. The graph classification problem consists, given a weighted graph and a partial node labeling, in extending the labels to all nodes. In many real-world context, such as Gene Function Prediction, the partial labeling is unbalanced: positive labels are much less than negatives. In this paper we present a new neural algorithm for predicting labels in presence of label imbalance. This algorithm is based on a family of Hopfield networks, described by 2 continuous parameters and 1 discrete parameter, and it consists of two main steps: 1) the network parameters are learnt through a cost-sensitive optimization procedure based on local search; 2) a suitable Hopfield network restricted to unlabeled nodes is considered and simulated. The reached equilibrium point induces the classification of unlabeled nodes. An experimental analysis on real-world unbalanced data in the context of genome-wide prediction of gene functions show the effectiveness of the proposed approach.

Keywords: Neural Network, Hopfield Network, Gene Function Prediction.

1 Introduction

Label learning in graphs requires, given a graph with a partial classification of the nodes, to extend the classification to all nodes. Methods for solving this problem are useful in application domains where data are naturally represented as connected nodes, i.e. biological networks [16], social networks [4] and World-Wide-Web [6].

Several methods have been proposed for node classification. First algorithms rely on the *guilt-by-association* principle, which classify unlabeled nodes according to the majority of the labels in their direct neighborhoods [11]. Furthermore, nodes can propagate labels to their neighbors with an iterative process until convergence [17]. Markov Random Walks have been applied to tune the amount of propagation we allow in the graph, by setting the length of the walk across the graph [13]. Other approaches are based on graph regularization [1], on global graph consistency [9], on Markov [5] and Gaussian Random Fields [14].

Unfortunately, these methods suffer a decay in the quality of solutions when input data are unbalanced, that is positive examples are significantly less than those negative. This issue is particularly relevant in Gene Function Prediction (GFP), where the imbalance in data requires to adopt cost-sensitive strategies [7].

For their common characteristics, many of the described approaches can be cast into a common framework where a quadratic cost objective function is minimized [2]. From this point of view, it seems natural a neural approach based on Hopfield networks, that are local optimizers of quadratic functions [9].

In [9] the neural algorithm GAIN is applied to GFP. Neurons represent genes, the connection weights the “similarities” between genes, the activation values are 1, -1 and the thresholds are 0 for each neuron. Fixed a functional class, only a subset neurons are classified (positive or negative), while the classification of the other is unknown. For classifying the unlabelled neurons, an initial state x is given by setting 1 the positive neurons, -1 the negative neurons and 0 those still unclassified. The dynamics of the network is applied to this state until the equilibrium point \hat{x} is reached; a gene k is classified as “positive” iff $\hat{x}_k = 1$.

From a biological standpoint, this approach is motivated by the fact that minimizing the overall energy means maximizing the weighted sum of edges connecting neurons with the same activation value. Nevertheless, this algorithm is affected by the *imbalance problem* in functional classes. Since weights are non negative and thresholds are 0, when the positive examples are less than the negative, the network is likely to converge to a trivial state $(-1, -1, \dots, -1)$. Observe that, in biological taxonomies, for most of the functional classes only a small number of positive examples is available.

In [3] another neural algorithm, called COSNet, has been proposed for solving the GFP problem on unbalanced data. As in the previous approach, neurons represent genes and connection weights represent the similarities between genes. However, here a class of networks with 2 parameters is considered: each neuron has activation values $\sin \alpha$ and $-\cos \alpha$ and threshold γ . Firstly, the algorithm learns the optimal values of the parameters α and γ , then it runs the subnetwork restricted neurons with unknown classification, that are classified according to the reached equilibrium state.

We point out that in both previous algorithms all the neurons of the network have the same activation values. Since, in principle, each neuron in a Hopfield network might have different activation values, in this work we investigate this case by partitioning the neurons in two classes and assigning to each class different activation values.

Accordingly, in Sect. 3 a family of parametrized Hopfield networks is introduced, whose parameters are the the possible partitions of neurons in 2 classes and the corresponding activation values. In Sect. 5 it is derived an algorithm that firstly learns the optimal values of the 2 continuous parameters (the different activation values) and the discrete parameter (the neuron partition). Then the algorithm runs the subnetwork restricted to neurons with unknown classification, that are classified according to the reached equilibrium state. Finally, in Sect. 6 we describe the experimental procedure adopted to validate the algorithm on the genome-wide prediction of gene functions in a model organism, including around 200 functional classes of the FunCat taxonomy [12], and using 3 different types of biomolecular data.

2 Gene Function Prediction (GFP)

In our setting, GFP is formalized as the problem of label learning in graphs [2]. Genes are represented by a set of nodes $V = \{1, 2, \dots, n\}$ and relationships between genes are

encoded through a symmetric $n \times n$ real weight matrix W , whose elements w_{ij} represent similarities between genes i and j .

For a given functional class c , the nodes V are labeled with $\{+, -\}$, leading to the subsets P and N of positive and negative vertices for class c . For most model organisms, usually the functional labeling is known only for a subset $S \subset V$, while is unknown for $U = V \setminus S$. Let be $S^+ = S \cap P$ and $S^- = S \cap N$: we can refer to S^+ , S^- and W as the "prior information" of the GFP problem.

The *Gene Function Prediction problem* consists in finding a bipartition (U^+, U^-) of genes in U on the basis of the prior information. Genes in U^+ are then considered candidates for the class $P \cap U$. From this standpoint, GFP is set as a semi-supervised learning problem on graphs, since gene functions can be predicted by exploiting both labeled and unlabeled nodes/genes and the weighted connections between them.

3 Hopfield Networks for GFP

In this Section we consider a family of Hopfield networks [8] with binary neurons partitioned in two classes G_1 and G_2 . The activation values are $\{\sin \alpha_1, -\cos \alpha_1\}$ for neurons in G_1 and $\{\sin \alpha_2, -\cos \alpha_2\}$ for neurons in G_2 ; the thresholds are set to 0.

Formally, in our setting, a *Hopfield network* H with neurons $V = \{1, 2, \dots, n\}$ is a quadruple $H = \langle W, b, \alpha_1, \alpha_2 \rangle$, where:

- $W = (w_{ij})$ is a $n \times n$ symmetric matrix with null diagonal, whose elements $w_{ij} \in \mathbb{R}$ represent the connection strength between neurons i and j
- $b \in \{0, 1\}^n$ is a binary vector partitioning neurons in two classes:

$$G_1 = \{k | b_k = 1\}, G_2 = \{k | b_k = 0\}$$
- α_1, α_2 are (possibly distinct) real values denoting the neuron activation values: $\{\sin \alpha_1, -\cos \alpha_1\}$ (resp. $\{\sin \alpha_2, -\cos \alpha_2\}$) for neurons k such that $b_k = 1$ (resp. $b_k = 0$)

The dynamics of the network is described as follows:

1. At time 0 an initial value $x_i(0)$ is given for each neuron i
2. At time $t + 1$ each neuron is updated asynchronously (up to a permutation) by the following activation rule

$$x_i(t + 1) = \begin{cases} b_i \sin \alpha_1 + (1 - b_i) \sin \alpha_2 & \text{if } \sum_{j=1}^{i-1} w_{ij} x_j(t + 1) + \sum_{k=i+1}^n w_{ik} x_k(t) > 0 \\ -b_i \cos \alpha_1 - (1 - b_i) \cos \alpha_2 & \text{if } \sum_{j=1}^{i-1} w_{ij} x_j(t + 1) + \sum_{k=i+1}^n w_{ik} x_k(t) \leq 0 \end{cases} \quad (1)$$

The state of the network at time t is $x = (x_1(t), x_2(t), \dots, x_n(t))$. The main feature of a Hopfield network is that it admits a Lyapunov function of the dynamics. In particular, consider the following quadratic state function (*energy function*):

$$E(x) = -\frac{1}{2} x^T W x \quad (2)$$

During the dynamics this function is not increasing; this guarantees that the dynamics converges to an equilibrium state $\hat{x} = (\hat{x}_1, \hat{x}_2, \dots, \hat{x}_n)$, which corresponds to a local minimum of the energy function [8].

4 Subnetwork Property

Let be $H = \langle W, b, \alpha_1, \alpha_2 \rangle$ a Hopfield network. Fixed $U = \{1, 2, \dots, h\}$ and $S = \{h+1, h+2, \dots, n\}$, each network state x can be decomposed in $x = (u, s)$, where u and s are respectively the states of neurons in U and in S . The energy function of H can be written by separating the contributions due to u and s :

$$\begin{aligned} E(u, s) &= -\frac{1}{2} (u^T W_{uu} u + s^T W_{ss} s + u^T W_{us} s + s^T W_{us}^T u) \\ &= -\frac{1}{2} u^T W_{uu} u + u^T (-W_{us} s) + C \end{aligned} \quad (3)$$

where $W = \begin{pmatrix} W_{uu} & W_{us} \\ W_{us}^T & W_{ss} \end{pmatrix}$ is the weight matrix W decomposed in its submatrices W_{uu} connecting nodes in U , W_{ss} connecting nodes in S , W_{us} connecting each node in U with each node in S , and W_{us}^T its transpose. $C = -\frac{1}{2} s^T W_{ss} s$ is a term constant w.r.t. u .

Suppose now that a state \tilde{s} of neurons in S is given. We are interested in the dynamics obtained by allowing the update just of neurons in U , without updating neurons in S . We denote with $H_{U|\tilde{s}}$ the Hopfield network with neurons U which realizes this dynamics and $E_{|\tilde{s}}$ the corresponding energy; from equation (3) it holds:

Theorem 1. $H_{U|\tilde{s}} = \langle W_{uu}, b^u, \alpha_1, \alpha_2 \rangle$, with thresholds $-W_{us}\tilde{s}$ and where b^u is the subvector of b restricted to neurons in U .

Given a state \tilde{s} of neurons in S , we say that \tilde{s} is part of global minimum of the energy E of H if there is a state u of neurons in U s.t. (u, \tilde{s}) is a global minimum of E . The introduction of the network $H_{U|\tilde{s}}$, is motivated by the following property:

Theorem 2. (Subnetwork property) *If \tilde{s} is part of a energy global minimum of H , and \tilde{u} is a global minimum of the energy $E_{|\tilde{s}}(u)$, then (\tilde{u}, \tilde{s}) is a energy global minimum of H .*

In our setting, we associate the given bipartition (S^+, S^-) of S with the state $\tilde{s} = x(S^+, S^-)$:

$$x_i(S^+, S^-) = \begin{cases} b_i \sin \alpha_1 + (1 - b_i) \sin \alpha_2 & \text{if } i \in S^+ \\ -b_i \cos \alpha_1 - (1 - b_i) \cos \alpha_2 & \text{if } i \in S^- \end{cases}$$

for each $i \in S$. Suppose, for suitable b, α_1, α_2 , that $x(S^+, S^-)$ is part of a energy global minimum of $H = \langle W, b, \alpha_1, \alpha_2 \rangle$; then, by the subnetwork property, we can predict the hidden part relative to neurons U by minimizing the energy of $H_{U|x(S^+, S^-)}$.

5 Algorithm for GFP

In this Section we exhibit a procedure based on Hopfield networks for dealing with the GFP problem.

For a given similarity matrix W , we consider the class of networks $H = \langle W, b, \alpha_1, \alpha_2 \rangle$ on neurons $V = \{1, 2, \dots, n\}$, where α_1, α_2 are real parameters and $b \in \{0, 1\}^n$ is a discrete parameter.

Fixed a functional class, an instance of GFP problem is given by the matrix W and the sets S^+ and S^- of positive and negative examples. We hypothesize that there exist a triple $(\hat{b}, \hat{\alpha}_1, \hat{\alpha}_2)$ such that:

1. The solution of the problem corresponds to an energy global minimum of $H = \langle W, \hat{b}, \hat{\alpha}_1, \hat{\alpha}_2 \rangle$
2. $x(S^+, S^-)$ is part of an energy global minimum of H

Then, by Theorem 2 we can discover the hidden part \hat{u} of the global minimum by minimizing the energy of the network $H_{U|x(S^+, S^-)}$. Accordingly, the procedure for solving the GFP problem can be factorized into two main steps:

Step 1. Determine the parameters $(\hat{b}, \hat{\alpha}_1, \hat{\alpha}_2)$ such that the state $x(S^+, S^-)$ is approximately part of a global minimum by finding the parameters (b, α_1, α_2) for which $x(S^+, S^-)$ is "as close as possible" to a part of an equilibrium state of H .

Step 2. Minimize the energy function of the network $H_{U|x(S^+, S^-)}$ with the estimated parameters $(\hat{b}, \hat{\alpha}_1, \hat{\alpha}_2)$ by reaching an equilibrium state \hat{u} in a dynamics generated by a suitable initial state.

Finally, the solution (U^+, U^-) of GFP is:

$$\begin{aligned} U^+ &= \{i \in U \mid \hat{u}_i > 0\} \\ U^- &= \{i \in U \mid \hat{u}_i \leq 0\}. \end{aligned}$$

In the following we discuss in more details Step 1 (Section 5.1) and Step 2 (Section 5.2) of the algorithm.

5.1 Finding the Optimal Parameters

The main goal of this step is to find the values of the parameters b , α_1 and α_2 such that the state $x(S^+, S^-)$ is "as close as possible" to an equilibrium state.

To this end, we consider the parametrized subnetwork restricted to neurons in S , i.e. $H_S = \langle W_{ss}, b^s, \alpha_1, \alpha_2 \rangle$, where b^s , α_1 , α_2 are the parameter to be learned.

In the following we describe the objective function adopted for learning the network parameters and the relative optimization procedure.

Objective Function. First of all, we fix b^s , α_1 , α_2 . Every neuron i has an "internal energy" A_i , where:

$$\begin{aligned} A_i &= \sin \alpha_1 \sum_{k \in S} w_{ik} P_k b_k^s + \sin \alpha_2 \sum_{k \in S} w_{ik} P_k (1 - b_k^s) \\ &\quad - \cos \alpha_1 \sum_{k \in S} w_{ik} (1 - P_k) b_k^s - \cos \alpha_2 \sum_{k \in S} w_{ik} (1 - P_k) (1 - b_k^s) \end{aligned} \quad (4)$$

where P is the characteristic vector of S^+ (i.e. $P_k = 1$ iff $k \in S^+$). By means of A_i , we are able in computing the number of *true positive* TP, *false negative* FN and *false positive* FP:

- $TP(b^s, \alpha_1, \alpha_2) = \sum_{i \in S} P_i \text{HS}(A_i)$, i.e. the number of positive examples with positive internal energy (true positive)

- $FN(b^s, \alpha_1, \alpha_2) = \sum_{i \in S} P_i (1 - \text{HS}(A_i))$, i.e. the number of positive examples with negative internal energy (positive misclassification)

- $FP(b^s, \alpha_1, \alpha_2) = \sum_{i \in S} (1 - P_i) \text{HS}(A_i)$, i.e. is the number of negative examples with positive internal energy (negative misclassification)

Here HS denotes the Heaviside function ($\text{HS}(x) = 1$ if $x \geq 0$, 0 otherwise).

The function we want to maximize is the so called F_{score} :

$$F_{score}(b^s, \alpha_1, \alpha_2) = \frac{2TP}{2TP + FP + FN}$$

By observing that $0 \leq F_{score} \leq 1$, this criterion is justified by the following:

Theorem 3. $F_{score}(b^s, \alpha_1, \alpha_2) = 1$ iff $x(S^+, S^-)$ is an equilibrium state of the sub-network H_S .

Optimization Procedure. The values of parameters that maximize the F_{score} criterion are:

$$(\hat{b}^s, \hat{\alpha}_1, \hat{\alpha}_2) = \underset{b^s \in \{0,1\}^{|S|}, \alpha_1, \alpha_2}{\text{argmax}} F_{score}(b^s, \alpha_1, \alpha_2). \quad (5)$$

For every $b^s \in \{0,1\}^{|S|}$, we define $F(b^s) = \max_{\alpha_1, \alpha_2} F_{score}(b^s, \alpha_1, \alpha_2)$. Given b^s , an approximation of $F(b^s)$ can be found by applying a standard continuous optimization procedure.

In order to maximize $F(b^s)$, we adopt a simple local search on hypercube $\{0,1\}^{|S|}$, where the neighborhood of b^s is $\{\bar{b}^s \mid d_H(b^s, \bar{b}^s) = 1\}$, and d_H is the Hamming distance. Once obtained the local optimum \hat{b}^s , we determine the optimal values for α_1 and α_2 as $(\hat{\alpha}_1, \hat{\alpha}_2) = \underset{\alpha_1, \alpha_2}{\text{argmax}} F_{score}(\hat{b}^s, \alpha_1, \alpha_2)$.

Having the optimal values $(\hat{b}^s, \hat{\alpha}_1, \hat{\alpha}_2)$, we want to extend the vector \hat{b}^s to $\hat{b} = (\hat{b}^u, \hat{b}^s)$, where the indices of \hat{b}^u are the elements of U .

With regard to this, compute for all k :

$$\Delta_k^+ = \sum_{i \in S^+} w_{ki} \quad ; \quad \Delta_k^- = \sum_{i \in S^-} w_{ki}$$

In this way, we associate with each neuron k a point $P_k = (\Delta_k^+, \Delta_k^-)$ in the plane. Consider now the subsets of points C_1 and C_2 , where:

$$C_1 = \{P_k : \hat{b}_k^s = 1\} \quad ; \quad C_2 = \{P_k : \hat{b}_k^s = 0\}$$

By using C_1, C_2 we learn two bivariate normal distributions $N_2(\mu_1, \Sigma_1), N_2(\mu_2, \Sigma_2)$ where, for $j = 1, 2$, μ_j and Σ_j are respectively the sample mean and the sample covariance of C_j .

Finally, if $k \in U$, we set $\hat{b}_k^u = 1$ if and only if the probability of P_k , according to $N_2(\mu_1, \Sigma_1)$, is greater than the probability of P_k , according to $N_2(\mu_2, \Sigma_2)$.

5.2 Finding the Unknown Labels by Network Dynamics

After the computation of the optimal parameters $(\hat{b}, \hat{\alpha}_1, \hat{\alpha}_2)$, we consider the sub-network $H_{U|x(S^+, S^-)}$:

$$H_{U|x(S^+, S^-)} = \langle W_{uu}, \hat{b}^u, \hat{\alpha}_1, \hat{\alpha}_2 \rangle \quad (6)$$

with thresholds $-W_{su}^T x(S^+, S^-)$.

Fixed an initial state $u_i = 0$ for each $i \in \{1, 2, \dots, h\}$, we run the sub-network $H_{U|x(S^+, S^-)}$ to learn the unknown labels of neurons U .

If \hat{u} is the stable state reached by this dynamics, we obtain the final solution (U^+, U^-) by setting:

$$U^+ = \{k | \hat{u}_k > 0\}, \quad U^- = \{k | \hat{u}_k \leq 0\}$$

6 Algorithm Validation

In this Section we describe the procedure for experimentally evaluate our algorithm and we discuss the results of the comparison of the algorithm with other state-of-the-art methods.

6.1 Experimental Setting

We performed predictions of gene functions at genome-wide level in the *S.cerevisiae* organism (yeast), using the whole FunCat ontology [12]. We predicted functions of genes belonging to three different biomolecular data sets previously adopted in [3]:

- *Pfam* is an enriched representation of Pfam domains by replacing the binary scoring with log E-values obtained with the HMMER software toolkit. This dataset contains 3528 genes and 5724 features.
- *Expr* data contains 250 gene expression measures of 4523 genes
- *SP-sim* is a data set containing pairwise similarities between 3527 yeast genes represented by Smith and Waterman log-E values between all pairs of yeast sequences

As validation procedure we adopt the 10-folds cross validation: genes are randomly divided into 10 equal-sized subsets, and each time the labels for genes in a fold are hidden and predicted using as training data the other nine folds.

6.2 Results

First of all, we compared our method with semi-supervised and supervised machine learning methods proposed in the literature for the Gene Function Prediction problem. We consider: 1) the *GAIN* algorithm [9]; 2) *Zhu-LP*, a popular semi-supervised label propagation learning algorithm based on Gaussian random fields and its class mass

¹ We used the funcat-2.1 scheme with the annotation data funcat-2.1_data_20070316, available from: ftp://ftpmips.gsf.de/yeast/catalogues/funecat/funecat-2.1_data_20070316




Fig. 1. F-score comparison in terms of averaged F_{score}

normalized version *Zhu-LP-CMN* [17]; 3) Support Vector Machines with linear (*SVM-l*) and Gaussian (*SVM-g*) kernels [10].

In order to take into account the imbalance in positive and negative labels characterizing the GFP context, we adopt the *F-score* performance measure (Sect. 5.1). Figure 1 shows for each dataset and for all the considered methods, the average F-score across all the functional classes.

Our algorithm highly outperforms in terms of average F-score all the other compared methods, and the difference is always significant at 10^{-6} significance level, according to the Wilcoxon signed-ranks test [15]. We think this results are due to the inherent cost-sensitive nature of the algorithm, which is able in automatically finding the parameters that better “re-equilibrate” the imbalance in labels.

Moreover, in order to better analyze the performance of our algorithm, we evaluate also the precision of the algorithm, $precision = \frac{TP}{TP+FP}$, which informally is the probability that a positive prediction corresponds to a true positive. We point out that in GFP context the automatic positive predictions of unknown genes need to be confirmed by expensive experimental laboratory procedures; accordingly, achieving a high precision in predicting functions of unknown genes is central, and provides reliable clues to experimentally check the membership of a gene to a functional class.

In Table 1 we show the averaged *precision* and F_{score} of our algorithm and of another cost-sensitive algorithm, COSNet, proposed in [3].

We can observe that the two algorithms achieve close values of F_{score} ; on the other hand, the present algorithm obtains a significant improvement in precision at $\alpha = 5 * 10^{-10}$ significance level (Wilcoxon signed-ranks test).

Table 1. Precision and F-score of COSNet and of our algorithm averaged across the Funcat functional classes

Pfam		
Method	Precision	F
COSNet	0.445	0.375
Our proposal	0.509	0.370
Expr		
Method	Precision	F
COSNet	0.057	0.085
Our proposal	0.147	0.105
SP-sim		
Method	Precision	F
COSNet	0.445	0.376
Our proposal	0.489	0.368

Conclusions

In this paper we propose a new algorithm for predicting node labels in graph in presence of label imbalance. The algorithm is based on a family of Hopfield networks with 2 real parameters and 1 discrete parameter. The parameters are learned by means of a cost-sensitive procedure, which allows to manage the imbalance in data. Then the sub-network of unlabeled nodes is simulated and the reached equilibrium state provides the classification of unlabeled nodes.

The algorithm has been experimentally validated on the problem of predicting the functions of genes in a model organism; the results, compared with those of the state-of-the-art methods, show the effectiveness of this approach.

In this paper, neurons are bi-partitioned, but in principle we could consider k -partitions, increasing the number of parameters. It should be interesting to evaluate the impact that the number of parameters has on the predicting capabilities of the algorithm, and to define the optimal number of parameters (which ensures to avoid overfitting) through model selection techniques.

References

1. Belkin, M., Matveeva, I., Niyogi, P.: Regularization and Semi-supervised Learning on Large Graphs. In: Shawe-Taylor, J., Singer, Y. (eds.) COLT 2004. LNCS (LNAI), vol. 3120, pp. 624–638. Springer, Heidelberg (2004)
2. Bengio, Y., Delalleau, O., Le Roux, N.: Label Propagation and Quadratic Criterion. In: Chapelle, O., Scholkopf, B., Zien, A. (eds.) Semi-Supervised Learning, pp. 193–216. MIT Press (2006)

3. Bertoni, A., Frasca, M., Valentini, G.: *COSNet*: A Cost Sensitive Neural Network for Semi-supervised Learning in Graphs. In: Gunopulos, D., Hofmann, T., Malerba, D., Vazirgiannis, M. (eds.) ECML PKDD 2011, Part I. LNCS, vol. 6911, pp. 219–234. Springer, Heidelberg (2011)
4. Borgatti, S., Mehra, A., Brass, D., Labianca, G.: Network Analysis in the Social Sciences. *Science* 322, 892–895 (2009)
5. Deng, M., Chen, T., Sun, F.: An integrated probabilistic model for functional prediction of proteins. *J. Comput. Biol.* 11, 463–475 (2004)
6. Dorogovtsev, S., Mendes, J.: *Evolution of networks: From biological nets to the Internet and WWW*. Oxford University Press, Oxford (2003)
7. Elkan, C.: The foundations of cost-sensitive learning. In: Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence, pp. 973–978 (2001)
8. Hopfield, J.: Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA* 79, 2554–2558 (1982)
9. Karaoz, U., et al.: Whole-genome annotation by using evidence integration in functional-linkage networks. *Proc. Natl. Acad. Sci. USA* 101, 2888–2893 (2004)
10. Lin, H.T., Lin, C.J., Weng, R.: A note on Platt’s probabilistic outputs for support vector machines. *Machine Learning* 68(3), 267–276 (2007)
11. Marcotte, E., Pellegrini, M., Thompson, M., Yeates, T., Eisenberg, D.: A combined algorithm for genome-wide prediction of protein function. *Nature* 402, 83–86 (1999)
12. Ruepp, A., et al.: The FunCat, a functional annotation scheme for systematic classification of proteins from whole genomes. *Nucleic Acids Research* 32(18), 5539–5545 (2004)
13. Szummer, M., Jaakkola, T.: Partially labeled classification with Markov random walks. In: *Advances in Neural Information Processing Systems (NIPS)*, vol. 14, pp. 945–952. MIT Press (2001)
14. Tsuda, K., Shin, H., Scholkopf, B.: Fast protein classification with multiple networks. *Bioinformatics* 21(suppl. 2), ii59–ii65 (2005)
15. Wilcoxon, F.: Individual comparisons by ranking methods. *Biometrics* 1, 80–83 (1945)
16. Wuchty, S., Ravasz, E., Barabasi, A.L.: The architecture of biological networks. *Complex Systems in Biomedicine* 5259, 165–181 (2003)
17. Zhu, X., Ghahramani, Z., Lafferty, J.: Semi-supervised learning using gaussian fields and harmonic functions. In: *ICML*, pp. 912–919 (2003)

Handwritten Digits Recognition by Bio-inspired Hierarchical Networks

Antonio G. Zippo, Giuliana Gelsomino, Sara Nencini, and Gabriele E.M. Biella

Institute of Molecular Bioimaging and Physiology,
Consiglio Nazionale delle Ricerche,
Via Fratelli Cervi 21, 20090 Segrate Milan, Italy
{antonio.zippo, giuliana.gelsomino,
sara.nencini, gabriele.biella}@ibfm.cnr.it
<http://www.ibfm.cnr.it>

Abstract. The human brain processes information showing learning and prediction abilities but the underlying neuronal mechanisms still remain unknown. Recently, many studies prove that neuronal networks are able of both generalizations and associations of sensory inputs.

In this paper, following a set of neurophysiological evidences, we propose a learning framework with a strong biological plausibility that mimics prominent functions of cortical circuitries. We developed the Inductive Conceptual Network (ICN), that is a hierarchical bio-inspired network, able to learn invariant patterns by Variable-order Markov Models implemented in its nodes. The outputs of the top-most node of ICN hierarchy, representing the highest input generalization, allow for automatic classification of inputs. We found that the ICN clustered MNIST images with an error of 5.73% and USPS images with an error of 12.56%.

Keywords: pattern recognition, handwritten digits, abstraction process, hierarchical network.

1 Introduction

The brain is a computational device for information processing and its flexible and adaptive behaviors emerge from a system of interacting neurons depicting very complex networks [1]. Many biological evidences suggest that the neocortex implements a common set of algorithms to perform “intelligent” behaviors like learning and prediction. In particular, two important related aspects seem to represent the crucial core for learning in biological neural networks: the hierarchical information processing and the abstraction process [2]. The hierarchical architecture emerges from anatomical considerations and is fundamental for associative learning (e.g. multisensory integration). The abstraction instead leads the inference of concepts from senses and perceptions (Fig. 1D).

Specifically, information from sensory receptors (eyes, skin, ears, etc.) travels into the human cortical circuits following subsequent abstraction processes. For

instance, elementary sound features (e.g. frequency, intensity, etc.) are first processed in the primary stages of human auditory system (choclea). Subsequently sound information gets all the stages of the auditory pathway up to the cortex where higher level features are extracted (Fig. 1E-F). In this way information passes from raw data to objects, following an abstraction process in a hierarchical layout. Thus, biological neural networks perform generalization and association of sensory information. For instance, we can associate sounds, images or other sensory objects that present together as it happens in many natural and experimental settings like during Pavlovian conditioning. Biological networks process these inputs following a hierarchical order. In a first stations inputs from distinct senses are separately processed accomplishing data abstraction. This process is repeated in each subsequent higher hierarchical layer. Doing so, in some hierarchical layer, inputs from several senses converge showing associations among sensory inputs.

Recent findings indicate that neurons can perform invariant recognitions of their input activity patterns producing specific modulations of their synaptic releases [3-7]. Although the comprehension of such neuronal mechanisms is still elusive, these hints can drive the development of algorithms closer to biology than spiking networks or other brain-inspired models appear to be.

In this work, we propose a learning framework based on these biological considerations, called Inductive Conceptual Network (ICN), and we tested the accuracy of this network on the MNIST and USPS datasets. The ICN represents a general biological plausible model of the learning mechanisms in neuronal networks. The invariant pattern recognition that occurs in the hierarchy nodes is achieved by modeling node inputs by Variable-order Markov Models (VMMs) [8, 9].

2 Methods

The methods of this work are based on a set of considerations extracted primarily from the Memory-Prediction framework proposed by Jeff Hawkins in his book *On Intelligence*. Therefore in this section we first present crucial aspects of brain information processing.

2.1 Background about Learning and the Memory-Prediction Framework

As preliminary step we introduce few theoretical concepts about learning and memory experiences in nervous systems. The human brain massively elaborates sensory information. Through some elusive mechanism, the brain builds models (formal representation) from observations. In such models, pattern recognition and abstraction play a crucial role [10]. The former allows for the capture of patterns from observations, the latter allows for transforming raw observations into abstract concepts. For instance, listening to sequence of unknown songs from an unknown singer we perform both pattern recognition and abstraction,

respectively when we identify sound features (e.g. beats per minute) and when we infer abstract information concerning the new singer (e.g. he/she plays jazz). Key features of these brain processes can be translated in algorithms [11]. Jeff Hawkins et al. recently proposed a new learning framework (Memory-Prediction [10]) based on abstraction processes and pattern recognitions. This paradigm claims that abstraction represents one of the most important tasks underlying learning in the brain and that occurs through the recognition of invariances. Moreover, he suggested that sensory inputs are processed hierarchically: each layer propagates to the next layer the invariant recognized patterns. In propagating only invariances and discarding everything else, data are compressed with size decreasing at every next layer. This is finely promoted by a pyramidal shape. Hawkins et al. implemented the Memory-Prediction framework into a set of software libraries specialized in image processing (Hierarchical Temporal Memory, HTM [11]) which exhibits invariant recognition by a complex hierarchy of node implementing the Hidden Markov Model algorithm [12, 13].

2.2 The Inductive Conceptual Network

We propose a different realization of the Memory-Prediction framework, called Inductive Conceptual Network (ICN), where biological neurons are individually identified by nodes (see Figure 1A-B) and invariant recognition is performed by modeling inputs with Variable-order Markov Models (VMM) [8, 9, 14]. The former assumption allowed us to pin down the ICN model into adequate biological background and to evaluate not only its learning ability but also its neurophysiological matching with neuronal population dynamics. The latter assumption addresses the problem of invariant recognition in a powerful and computational efficient way [9].

The Inductive Conceptual Network is a hierarchical spiking network working as online unsupervised learning algorithm. In common with HTM, the ICN displays a tree-shaped hierarchy of nodes (Figure 1D-F). Formally, ICN is a triplet (T, M, k) where $T = \{l_1, l_2, \dots, l_L\}$ is the vector that contains the number of nodes in each layer such that $l_1 > l_2 > \dots > l_L = 1$. Let $q = \sum_{i=1}^L l_i$ be the total number of nodes, and M is the $q \times q$ adjacency matrix representing the connections between nodes and k is the maximum Markov order, an indicator of the memory power of each node. For the construction of $M = \{m_{i,j} | i, j = 1, \dots, q\}$ that is initially set to $m_{i,j} = 0$, we proceeded iteratively following these two steps for nodes in each layer x :

1. a set of deterministic assignments: $\{m_{i,i+l_x} = 1, \dots, m_{i+p,i+l_x} = 1\}$ with $p = \lceil \frac{l_x}{l_{x+1}} \rceil + 1$ and $\forall i \in \{\sum_{k=1}^x l_x, \dots, (\sum_{k=1}^{x+1} l_x) - \lceil \frac{l_x}{l_{x+1}} \rceil\}$;
2. a set of random assignments: $\{m_{i+r,i+l_x} = 1 | r \sim U(1, l_{x+1})\}$

where l_x is the number of nodes in the generic layer x and U is the discrete uniform distribution. Layers handle inputs from the immediately preceding layer (layer below) except for the first that handles the raw input data. The matrix M is semi-randomly assigned respecting the multilayer architecture: each node

receives the downstairs-layer input both from their neighbour nodes and from a small set of randomly chosen ones (Figure 11C).

Nodes read inputs from their *dendrites* (Figs. 11A-B) and an algorithm estimates the joint probability distribution of the observed inputs (see below, VMM). Whether the observed input is the most expected (or is very close to) the node produces a 1 (representing a spike) towards their output nodes otherwise does nothing. The ICN is a general purpose learning framework, and although it has not been tested on non-visual tasks it can however be used for other sensory information processing.

2.3 Variable-Order Markov Models

The learning of spatiotemporal patterns is the subject of study of Sequential data learning that usually involves very common methods, like Hidden Markov Models (HMM). In fact, HMM are able to model complex symbolic sequences assuming hidden states that control the system dynamics. However, HMM training suffers from local optima and their accuracy performance has been overcome by VMMs. Other techniques like N -gram models (or N -order Markov Models) compute the frequency of each N long subsequence but in these models the number of possible model states grows exponentially with N . Therefore, both computational space and time issues arise.

In this perspective, the observed symbolic (binary) sequence is assumed to be generated by a stationary unknown symbol source $S = \langle \Sigma, P \rangle$ where Σ is the symbol alphabet and P is the probability distribution of symbols. A VMM (also known as Variable length Markov Chains), given the maximum order D of conditional dependencies and a training sequence s generated by S , returns a model for the source S that's an estimation \hat{P} of probability distribution P . Applying VMMs, instead of N -gram models, takes several advantages. A VMM estimation algorithm builds efficiently a model for S . In fact, only the occurred D -grams are stored and their conditional probabilities $p(\sigma|s)$, $\sigma \in \Sigma$ and $s \in \Sigma^{d \leq D}$ are estimated. This trick saves lots of memory and computational time and makes feasible to model sequences with very long dependencies ($D \in [1, 10^3]$) on current personal computers.

2.4 The Node Behavior and Invariance Recognition

We consider the inputs from dendrites that each neuron (node) sees as binary symbols emitted by a discrete source which releases outcomes following an unknown non-stationary probability distribution P . The aim of each node is to learn its source as best as possible so that it can recognize correctly recurrent patterns assigning to them highest probabilities. The VMMs are typically used for this task being able to model dependencies among symbols up to an arbitrary order. VMMs can be estimated by many algorithms. We took into consideration a famous efficient lossless compression algorithm, the Prediction by Partial Matching (PPM) [15, 16], implemented in an open-source Java software library [17].

Formally, a node reads a binary input (at each step) $s = (s_1, \dots, s_n)$ of length n that represents the all-or-none dendritic activity. Let $k < n$ be the maximum dependency allowed among input symbols, then each node builds its probability model feeding k -tuples of the received n -ary input s into the PPM algorithm. Each node has its own instance of the PPM algorithm. After this first learning phase, the node passes into the prediction mode and looks if it observes in s the most expected pattern (pattern that has the highest probability assignment). If it happens, the node produces 1 as output in correspondence to the salient patterns thus preserving the spatial structural organization of inputs. We introduce the further condition that a 1 is produced in correspondence of patterns having Hamming distance [18] very close to the most expected one. We make this choice to introduce a sort of noise tolerance in the pattern recognition process. In other words, during the coding (and second) stage, a node processes its input by k symbols at time. If the current k -tuple pattern is the highest probable (or is very close to, by Hamming distance) a 1 is inserted into the output code, otherwise it marks a 0.

For instance, let be $k = 3$ and 101 the most expected pattern. Let 110000101-100101 be the current input that updates the probability distribution P . Finally, the node produces the output sequence 00101 where 1 corresponds to the two occurrences of the most expected pattern (101). The pseudo-code of the algorithms governing respectively nodes and the hierarchy are the following:

Algorithm for nodes

```

Node()
  read input s = (s_1,...,s_n);
  for each k-tuple in s:
    update P by PPM(s_i,...,s_(i+k));
    if HammingDistance((s_i,...,s_(i+k)),bestPattern) < gamma:
      output(1);
    else
      output(0);
    update bestPattern;
  end
end

```

where the function PPM() updates the probabilistic model \hat{P} with the new input $s_i, \dots, s_{(i+k)}$. The function HammingDistance(\cdot, \cdot) computes the Hamming distance between two binary strings and the function PPM_best() returns the current most probable pattern.

Algorithm for ICN

```

ICN()
  for each image in dataset:
    bw = binarizeImage(image);
    assignInputToFirstLayer(bw);
    foreach layer in ICN:

```

```

        setInput(bw);
        learn();
        bw = predict();
    end
    collect(bw)
end
end

```

where in the `learn()` function the distribution \hat{P} of each node are updated and in the `predict()` function the spiking activity of the current layer is returned. Before evaluating the performance of ICN on handwritten digit datasets, we evaluated the learning capabilities of a single node by a simple experiment. We provided a sequence of 1000 binary 5-tuples as input to a node with 5 dendrites and $k = 5$ (Figs. 2A-B). The input sequence of 5-tuple is generated randomly inserting at each time with probability 0.25 a fixed pattern (equal to 10010). Simulating a Hebbian setup where at each of the five dendrites is associated a weight increased in case of positive temporal correlation between pre- and post-synaptic spikes and decreased otherwise [19], we make sure that the end of proposed sequence weights of first and fourth dendrites are strengthened to the detriment of other ones (Figure 2A).

2.5 Learning of Handwritten Digits

In the current form, ICN can perform unsupervised learning. To evaluate the learning capabilities of such framework, we gave as input to the first layer, the images of the MNIST (or USPS) dataset (handwritten digits, 0 to 9). Here we use the MNIST test set which contains 10000 examples and the whole 11000 sample of the USPS. The chosen instantiation of ICN was composed by 4 layers with respectively 50, 20, 5 and 1 node in each layer. The maximum Markov order k was set to 5 for all ICN nodes. All parameters in this section have been chosen empirically to best match the right classification of the digits. We expected that digit images were correctly grouped with respect to the represented number. MNIST images are represented by 28x28 (784) pixels of 8-bit gray level matrix. Instead, USPS images, are represented by 16x16 (256) pixels. Images were binarized setting a threshold on the 8-bit gray-level values to 80. As explained above, nodes produce bits and the result of this unsupervised learning is valuable in the outputs of the top-most node. In fact, this node retains the most abstract information regarding the observed images. Namely, something likes the concept of number. After some empirical tuning of parameters (number of nodes, layer and maximum Markov order), ICN was able to discriminate digits by the top-most node output code. For instance in some experiments, giving an image of digit 0, the ICN emitted the binary code 1000. In the same experiments, the code 0101 was reserved to the digit 1 and so on. Obviously, the ICN made errors and

digit-to-code associations were not unique, e.g. some seven digits can be incorrectly classified with code 0101 reserved for the 1. To estimate the learning error, we chose a representative code for each digit class. The representatives were selected as most frequent for each class. Thus, the learning error was computed by counting mismatches between labels and representative codes.

3 Experimental Results

The ICN algorithm has been developed following strict and recently found biological criteria from the neurophysiology of neuronal networks. Once ascertained that ICN nodes perform a sort of Hebbian plasticity (see section 2.4) we challenged the ICN with the MNIST dataset (handwritten digit images). The MNIST dataset represents a sort of *casting-out-nines* for learning systems; in fact, new proposed algorithms are tested on this dataset to check their attitude to learn.

The learning capabilities of ICN were tested by its clustering efficiency over the MNIST dataset. Before submitted to ICN every digit image was binarized by applying a threshold. Subsequently each image was fed into the first layer nodes. Invariant recognized patterns are then propagated, layer-by-layer up to the highest, following the execution of Algorithm-2 (see Methods). As a whole, an input image elicits a bit (spike) flux in the bottom layer, a code transmitted to the upper layer. The top-most layer, composed by only one node, finally generates its binary codes each corresponding to a digit (class) of the image input. We ascertained that at the best tuning of parameters the ICN model got an average error of 5.73%, an acceptable score in an unsupervised environment, remarkably not requiring any preprocessing stages such as image alignment, centering or dimensionality reduction. For the USPS dataset, however harder to learn, the best achieved error was of 12.56%.

Eventually, we further investigated the influence of dataset size in the learning performance. For this reason, we repeated the same experiments randomly subsampling both datasets to 1000 and 5000 samples. For both datasets, performance improved increasing the dataset size as shown in table 4.

4 Discussion

Even convolutional neural networks (CNNs) [21, 22] are biologically inspired by the pioneer works of Hubel et al. on the retinotopies of the cat's visual cortex [20]. Indeed, CNNs exploit the fact that nearby pixels are more tightly correlated than more distant ones. Furthermore by using a cascading structure of convolutional and subsampling layers, these networks show successfully invariant recognition of handwritten digits subjected to certain transformations (scaling, rotation or elastic deformation). Although CNNs are bio-inspired by the local

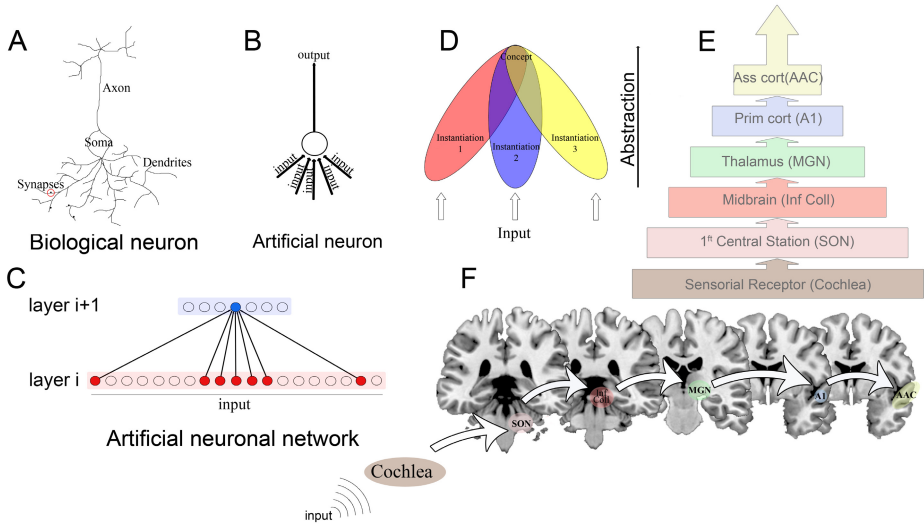


Fig. 1. Background and preliminary concepts on Inductive Conceptual Network (ICN). (A)-(B) Comparison between biological and artificial neurons. Biological signals conducted by each dendrite on soma can be represented by artificial inputs; after the input elaboration, axon conducts the signal output that in artificial neuron is computed by estimation of probability distribution of the observed inputs. (C) Graphically, neurons are represented by nodes (circle) which are organized in layers. They are linked by inter layer connections (edges) following a proximity criterion admitting exceptions. (D) Representation of the hierarchical abstraction framework that occurs getting from input representing raw data (concept instances, e.g. one, seven) to concepts (number). (E) An example of biological correspondence between the ICN and the auditory sensory system in human. Auditory input elaborated from cochlea, through sensory pathway reaches auditory associative cortex. (F) The auditory sensory pathways seen in the coronal MRI template slices. Abbreviations: SON, Superior Olivary Complex; Inf Coll, Inferior Colliculi; MGN, Medial Geniculate Nucleus; A1, Primary auditory cortex, AAC, Auditory Associative Cortex.

receptive fields which constitute the local features, the learning mechanism of the whole network does not appear to have a biological counterpart. Vice versa, the proposed network (ICN) implements invariant recognition exhibiting a spiking behavior in each node which represents a clear correspondence with biological networks. Furthermore, the algorithm governing nodes is the same in the whole network. Since the electrophysiological properties of neurons are quite similar, our network appears to be more plausible than CNNs where a set of special layers (and nodes) exclusively perform the invariant recognition.

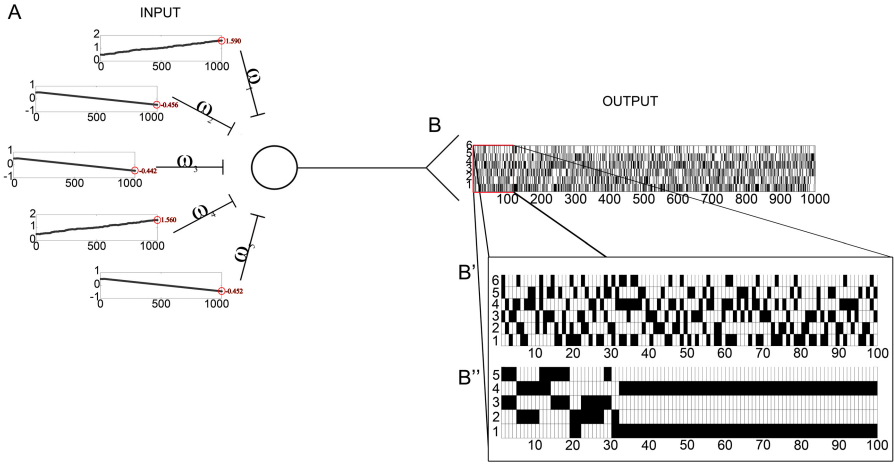


Fig. 2. Learning in a node of ICN assuming synaptic weights (ω) and the Hebb's rule. (A) Starting with equal weights (0.5) and assigning a positive increment (0.01) whether pre-synaptic spike precedes the post-synaptic spike in 2 timesteps at most. Otherwise the synaptic weight incurs in a negative reward (-0.01). The sequence of input patterns is composed by randomly generated binary inputs (with probability 0.75) plus a fixed input equal to 10010 (with probability 0.25). The simulation lasts 1000 timesteps where, at the end, the recurrent pattern 10010 was recognized assigning strong weights to the first and fourth synapses depressing the other ones. (B) In detail, the raster plot of the simulation where the activity of nodes 1-5 matches the activity of the 5 presynaptic inputs and the activity of node 6 is the output of node in examination. (B') An enlargement of the first 100 timesteps. (B'') The evolution of the most expected pattern according to the PPM estimation in the node. After the first 31 timesteps, the fixed pattern 10010 becomes the most expected.

The performance of each node is based on the PPM algorithm that requires $O(n)$ during learning and $O(n^2)$ during prediction as computational time complexities [9]. Although the quadratic complexity, each node receives only small fractions of inputs keeping n within small values. Thus the overall time complexity for each processed image raises to $O(m \cdot n^2)$, where m is the number of nodes. Interestingly, the node executions within each layer can be computed in parallel. Even the space complexity is dictated by the complexity of the PPM algorithm that is $O(k \cdot n)$, where k is the chosen Markov order, in the worst case. Therefore, the ICN algorithm requires $O(m \cdot k \cdot n)$ in space complexity.

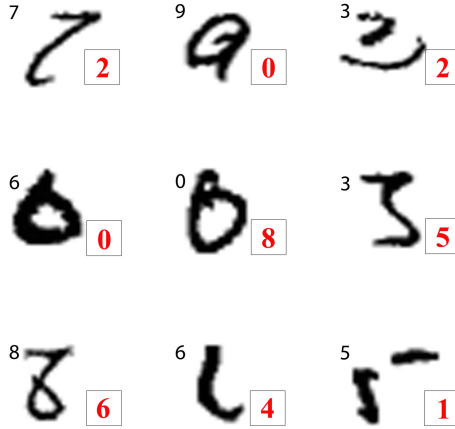


Fig. 3. Sample of common incorrect classifications on MNIST dataset. Numbers in the upper left of boxes indicate the correct representation. Numbers in the lower right of boxes indicate the incorrect classifications.




Fig. 4. Learning performances computed on 100 trials for each subsampling of the original datasets

5 Conclusion

The MNIST dataset is a standard test to evaluate learning accuracy for both linear and non-linear classifiers. We show here that ICN is apt to carry out unsupervised learning tasks with an error rate of 5.73% for MNIST and 12.56% for USPS at most. The percentage may appear weaker, in comparison with other learning methods, seemingly showing better error rates thanks, however to training and preprocessing (check for instance the performance of convolutional nets scoring down to 0.35% error rate). Furthermore, in comparison with other clustering techniques, our method does not fall into the *curse of dimensionality* [23]. Any classical unsupervised learning techniques, such as k-means, Expectation-Maximization or Support Vector Machines generally require an *ad hoc* dimensionality reduction (e.g. by Independent or Principal Component Analysis), a procedure that reduces the algorithm general purposiveness [24]. However, these networks do not acknowledge biological modeling, where ICN is instead adequately biologically oriented.

In conclusion, the proposed model achieves interesting preliminary results. Nevertheless further experiments with other machine learning datasets are required to strengthen its validity. Moreover, future developments can allow for effective multi-input integrations: for instance, two different sources of input (like sounds and images) could be associated by similar output codes even in presence of inputs from a single source.

References

1. Bassett, D.S., Greenfield, D.L., Meyer-Lindenberg, A., Weinberger, D.R., Moore, S.W., Bullmore, E.T.: Efficient physical embedding of topologically complex information processing networks in brains and computer circuits. *PLoS Computational Biology* 6(4), e1000748 (2010); Friston, K.J. (ed.), Public Library of Science, doi:10.1371/journal.pcbi.1000748
2. DiCarlo, J., Zoccolan, D., Rust, N.: How Does the Brain Solve Visual Object Recognition? *Neuron* 73, 415–434 (2012)
3. Takahashi, N., Kitamura, K., Matsuo, N., Mayford, M., Kano, M., Matsuki, N., Ikegaya, Y.: Locally Synchronized Synaptic Inputs. *Science* 335, 353–356 (2012)
4. Kleindienst, T., Winnubst, J., Roth-Alpermann, C., Bonhoeffer, T., Lohmann, C.: Activity-Dependent Clustering of Functional Synaptic Inputs on Developing Hippocampal Dendrites. *Neuron* 72(6), 1012–1024 (2011)
5. Makino, H., Malinow, R.: Compartmentalized versus Global Synaptic Plasticity on Dendrites Controlled by Experience. *Neuron* 72(6), 1001–1011 (2011)
6. Meyers, E.M., Qi, X.L., Constantinidis, C.: Incorporation of new information into prefrontal cortical activity after learning working memory tasks. *Proc. Natl. Acad. Sci. USA* (2012), doi:10.1073/pnas.1201022109
7. Branco, T., Clark, B.A., Hausser, M.: Dendritic Discrimination of Temporal Input Sequences in Cortical Neurons. *Science* 329(5999), 1671–1675 (2010)
8. Bühlmann, P., Wyner, A.J.: Variable Length Markov Chains. *The Annals of Statistics* 27(2), 480–513 (1999)
9. Begleiter, R., El-Yaniv, R., Yona, G.: On Prediction Using Variable Order Markov Models. *Journal of Artificial Intelligence Research* 22, 385–421 (2004)

10. Hawkins, J., Blackeslee, S.: *On Intelligence*. Times Books, New York (2004)
11. Dileep, G., Hawkins, J.: Towards a Mathematical Theory of Cortical Micro-circuits. *PLoS Comput. Biol.* 5(10), e1000532 (2009)
12. Baum, L.E., Petrie, T.: Statistical Inference for Probabilistic Functions of Finite State Markov Chains. *The Annals of Mathematical Statistics* 37(6), 1554–1563 (1966)
13. Baum, L.E., Petrie, T., Soules, G., Weiss, N.: A Maximization Technique Occurring in the Statistical Analysis of Probabilistic Functions of Markov Chains. *The Annals of Mathematical Statistics* 41, 164 (1970)
14. Rissanen, J.: Universal coding, information, prediction, and estimation. *IEEE Transactions on Information Theory* 30(4), 629–636 (1984)
15. Cleary, J., Witten, I.: Data Compression Using Adaptive Coding and Partial String Matching. *IEEE Transactions on Communications* 32(4), 396–402 (1984)
16. Teahan, W.: Probability estimation for PPM. In: *Proceedings of the New Zealand Computer Science Research Students' Conference*. University of Waikato, Hamilton (1995)
17. http://www.cs.technion.ac.il/~ronbeg/vmm/code_index.html
18. Hamming, R.W.: Error detecting and error correcting codes. *Bell System Technical Journal* 29(2), 147–160 (1950)
19. Hebb, D.O.: *The organization of behavior; a neuropsychological theory*. Wiley, New York (1949)
20. Hubel, D., Wiesel, T.: Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology* 195, 215–243 (1968)
21. LeCun, Y., Boser, B., Denker, J.S., et al.: Handwritten Digit Recognition with a Back-Propagation Network. In: Touretzky, D. (ed.) *Advances in Neural Information Processing Systems 2 (NIPS 1989)*. Morgan Kaufman, Denver (1990)
22. LeCun, Y., Bengio, Y.: Convolutional networks for images, speech, and time-series. In: Arbib, M.A. (ed.) *The Handbook of Brain Theory and Neural Networks*. MIT Press (1995)
23. Powell, W.B.: *Approximate Dynamic Programming: Solving the Curse of Dimensionality*. Wiley, New York (2007)
24. Kotsiantis, S., Kanellopoulos, D., Pintelas, P.: Data Preprocessing for Supervised Learning. *International Journal of Computer Science* 1(2), 111–117 (2006)

Forecasting Net Migration by Functional Demographic Model

Valeria D'Amato¹, Gabriella Piscopo², and Maria Russolillo¹

¹ Department of Economics and Statistics,
University of Salerno, Campus Fisciano, Italy

² {vdamato,mrussolillo}@unisa.it

Department of Mathematics for Decision,
University of Florence, Via Lombroso 6, Italy
gabriella.piscopo@unifi.it

Abstract. Net migration is the net total of migrants during the period, that is, the total number of immigrants less the annual number of emigrants, including both citizens and noncitizens. To derive estimates of net migration, the United Nations Population Division takes into account the past migration history of a country or area, the migration policy of a country, and the influx of refugees in recent periods. The data to calculate these official estimates come from a variety of sources, including border statistics, administrative records, surveys, and censuses. When no official estimates can be made because of insufficient data, net migration is derived through the balance equation, which is the difference between overall population growth and the natural increase during the intercensal period. In this contribution, we apply the functional data model to Italian data, for forecasting net migration numbers.

Keywords: Lee-Carter Models, Functional Demographic Model, Net Migration Forecasting, Simulation.

1 Introduction

Recently, the need for high-quality forecasts of international migration is becoming increasingly significant. Population movements are gaining in importance, given the diminishing impact of natural change on population dynamics. In demographic terms, natural change is the difference between the number of births and deaths and measures the contribution of vital events to the dynamics of the population. This phenomenon is crucial especially in the developed regions of the world, such as Europe, which are already facing zero or negative natural population growth (e.g., [12]; [8]). In particular, if we refer to annual net migration, i.e. the balance between immigration and emigration over a given time period, it has substantially increased since the beginning of the 1990s, exceeding natural change as a driver of Italian demographic trends in all years. The net migration rate indicates the contribution of migration to the overall level of population change. High levels of migration can cause problems such as increasing unemployment and potential ethnic strife (if people are coming in) or a reduction in

the labor force, perhaps in certain key sectors (if people are leaving). If we look at the statistics of the Italian population ([14]), we can see that the foreign who lived in Italy at the beginning of January 2011 were 4.570.317, 335.000 more than the year before (+7,9%). Moreover, we can observe that the population increase is slightly lower than 2009 (343.000 units). The role of migration is really important because it is not limited to demography, as it also affects many other areas of social life, including economy, labour relations, politics, and culture. Moreover, in a globalising world, the migratory processes are becoming more and more dynamic and complex ([7]). In consequence, migration forecasts are associated with a high level of uncertainty and usually bear very high ex-post errors ([10]). Nevertheless, as noted by [2], "one of the major purposes of statistical analysis is to make forecasts about the future and to offer suitable measures of uncertainty associated with unknown events or quantities". In this framework, aim of this paper is to develop a migration forecasting methodology in order to produce accurate projections. The structure of the paper is the following: when no official estimates can be made because of insufficient data, net migration is derived through the balance equation, which we describe in Section 2. In Section 3, we apply the method to Italian data in order to simulate future sample paths and probabilistic forecasts of net migration.

2 The Growth Balance Equation

As abovementioned, in this paper we aim to forecasting net international migration for use in national population forecasting through the application of functional data models and time series methods. Usually, to calculate official estimates of net migration a variety of sources, including border statistics, administrative records, surveys, and censuses have to be taken into account. When no official estimates can be made because of the lack of complete and reliable data for international migration, net migration is derived through the balance equation as in [6]. In their paper, the authors specify the demographic growth-balance equation, given by the difference between overall population growth and the natural increase, in order to estimate net migration. Let G_t be the net migration, estimated using the following demographic growth- balance equation:

$$G_t(x, x + 1) = P_{t+1}(x + 1) - P_t(x) + D_t(x, x + 1), \quad x = 0, 1, 2, \dots, p - 2 \quad (1)$$

$$G_t(p - 1^+, p^+) = P_{t+1}(p^+) - P_t(p^+) + P_t(p - 1) + D_t(p - 1^+, p^+) \quad (2)$$

$$G_t(B, 0) = P_{t+1}(0) - B_t + D_t(B, 0) \quad (3)$$

where $G_t(x, x + 1)$ refers to net migration in calendar year t of persons aged x at the beginning of year t , $G_t(p - 1^+, p^+)$ refers to net migration in calendar year t of persons aged $p - 1$ and older at the beginning of year t , and $G_t(B, 0)$ refers to net migration in calendar year t of births during year t . $P_t(x)$ indicates the population of age x at 1 January of year t and $D_t(x, x + 1)$ refers to deaths in calendar year t of persons aged x at the beginning of year t . To obtain forecasts of net migration component, we first develop functional time series model for the component $G_t(x, x + 1)$. We follow [5] model introduced to forecast age-specific mortality and fertility rates, based on the combination of functional data analysis, nonparametric smoothing and robust statistics. The model is adapted to forecast net migration by [6] as follows: let denote the net migration density for age x in year t , $f_t(x)$ the underlying smooth function of x , $\{x, y_t(x)\}$ the functional time series, where

$$y_t(x) = f_t(x) + \sigma_t(x)\epsilon_{t,x} \tag{4}$$

with $\{\epsilon_{t,x}\}$ a sequence of iid standard normal random variable and $\sigma_t(x)$ allowing for the amount of noise to vary with x . The smoothing is carried out with a non parametric method, based on weighted penalized regression splines. Then, the fitted curves is decomposed via a basis function expansion:

$$f_t(x) = \mu(x) + \sum_{k=1}^K \beta_{t,k} \phi_k(x) + e_t(x) \tag{5}$$

where $\mu(x)$ is the mean of $f_t(x)$ across years, $\phi_k(x)$ is a set of orthogonal basis functions calculated using a principal components decomposition and the modeling error $e_t(x) \sim N(0, var(x))$, is given by the difference between the smoothed curves and the fitted curves from the model. The observational variance, $\sigma_t^2(x)$, depends on the nature of the data. For migration data, [6] make no distributional assumptions, and estimate $\sigma_t^2(x)$ using a nonparametric regression of $[y_t(x) - f_t(x)]^2$ against x . As they observe that the model proposed by [5] is good at producing point forecasts, but not quite so good at estimating forecast variance, they propose an adjustment to the forecast variance implied by the model. In order to forecast $y_t(x)$, univariate time series models are fitted to each coefficients $\{\beta_{t,k}\}$, $k = 1, \dots, K$. Using them, the coefficients $\{\beta_{t,k}\}$, $k = 1, \dots, K$ are forecasted for $t = 1n + 1, \dots, n + h$, where h are the forecasted periods ahead. Finally, the previously obtained coefficients are implemented to get the $f_t(x)$ as in formula (5) and the $y_t(x)$ are projected. The estimated variance of error terms in (5) is used to compute confidence intervals for the net migration forecast.

3 Numerical Application

The numerical application involves demographic data, given by the Italian male, female and total mortality experience ranging from 1952 up to 2004. These data were downloaded in single years for 0-110 from the [4]. The Italian female fertility rates for single year of age from 1952 to 2004 for women aged between 13 and 50 years old were downloaded from demo.Istat. In particular, the fertility rates, defined as the number of live births during the calendar year, according to the age of the mother, per 1,000 of the female resident population of the same age at a certain date, are obtained as in [5].

In Fig. 1, we can observe the female and male Italian net migration patterns from 1952 to 2004. Migration is the difference between the number of persons entering and leaving a country during the year per 1,000 persons (based on midyear population). An excess of persons entering the country is referred to as net immigration; an excess of persons leaving the country as net emigration (e.g., -9.26 migrants/1,000 population). The net migration rate indicates the contribution of migration to the overall level of population change ([15]). From Figure 1, we can observe that the female and male patterns are quite similar, with some small variation for males, due to male-dominated emigration rather than of female-dominated immigration.

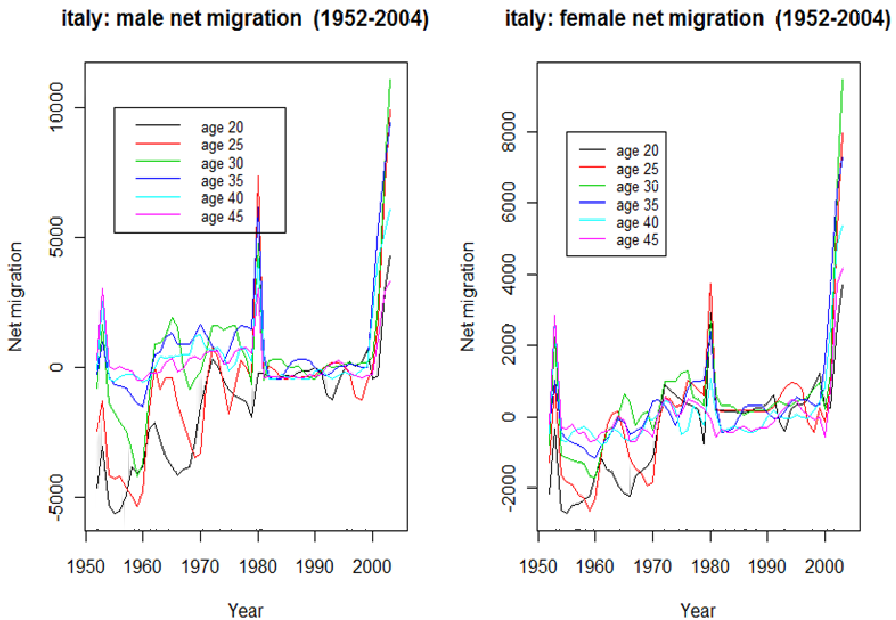


Fig. 1. Italian male and female net migration from 1952 to 2004

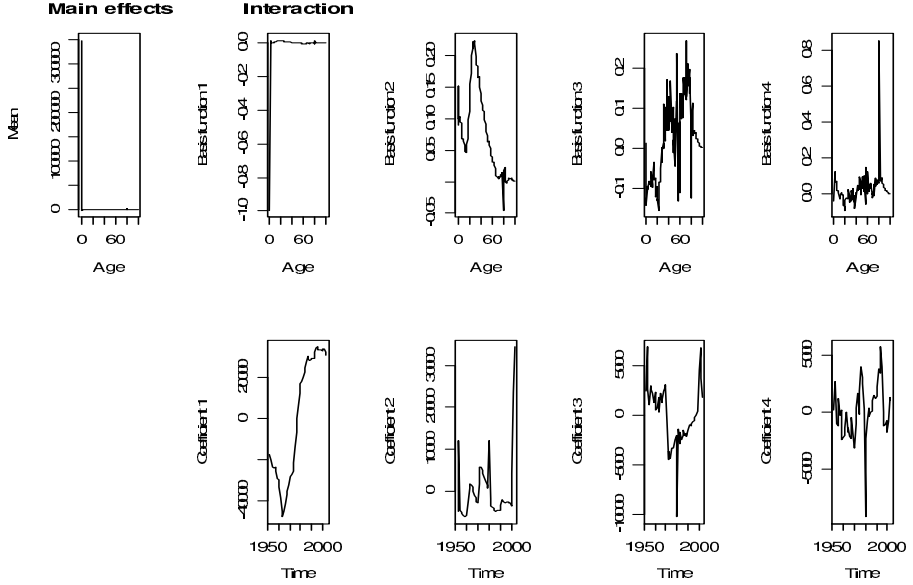


Fig. 2. FDM fitting on Italian female data

Table 1. Goodness of fitting measures

Averages across ages			
ME	MSE	MPE	MAPE
5.775600e-01	2.422651e+07	1.296480e+00	3.946150e+00
Averages across years			
IE	ISE	IPE	IAPE
6.977183e+01	8.182101e+08	1.312435e+02	3.946150e+00

We follow up analyzing the net migration component, by fitting the functional data model to these data. Annual net migration is estimated using the growth-balance equation as in [6]. As observed from the authors, the method is insensitive to the choice of K, provided that K is large enough. For our model, we choose K=4 basis functions. As an example, we report the fitted basis functions and the associated coefficients for Italian female net migration in Fig. 2. We find that the percentage variation due to basis functions is 96.0% 2.4% 0.4% 0.3%, which means that the first term of the basis functions accounts for at least 96% of the variation for female net migration.

To provide general intuition about the goodness of the fitting, we resort to the traditional diagnostics for detecting the structure of the residuals and we report the related error measures in Table 1. Although this estimate includes errors in estimated age-specific deaths, and errors in the estimated number of births, the resulting series is superior to the available migration data in its coverage of years, events and single years of age.

In Fig. 3 we show female net migration forecasts calculated for 20 years ahead, from 2005 to 2024 and the related forecasts error. It can be seen that between ages 20 and 40 there is a peak in net migration, probably due to labor migration and spouses.




Fig. 3. Female net migration forecasts from 2005 to 2024 and related forecasts error

To conclude the application of functional data model to net migration, we implement 10,000 bootstrap simulations on the 2025 female net migration forecasts for ages $x=25$, 35 and 45 in order to derive, for each projection, confidence intervals at level of 95%. We can observe that the prediction intervals decrease with the increasing of the age, as illustrated by the obtained confidence intervals: $2679.021 < x < 17479.45$, for $x=25$; $2631.45 < x < 15672.91$ for $x=35$ and $1371.916 < x < 10435.88$ for $x=45$.

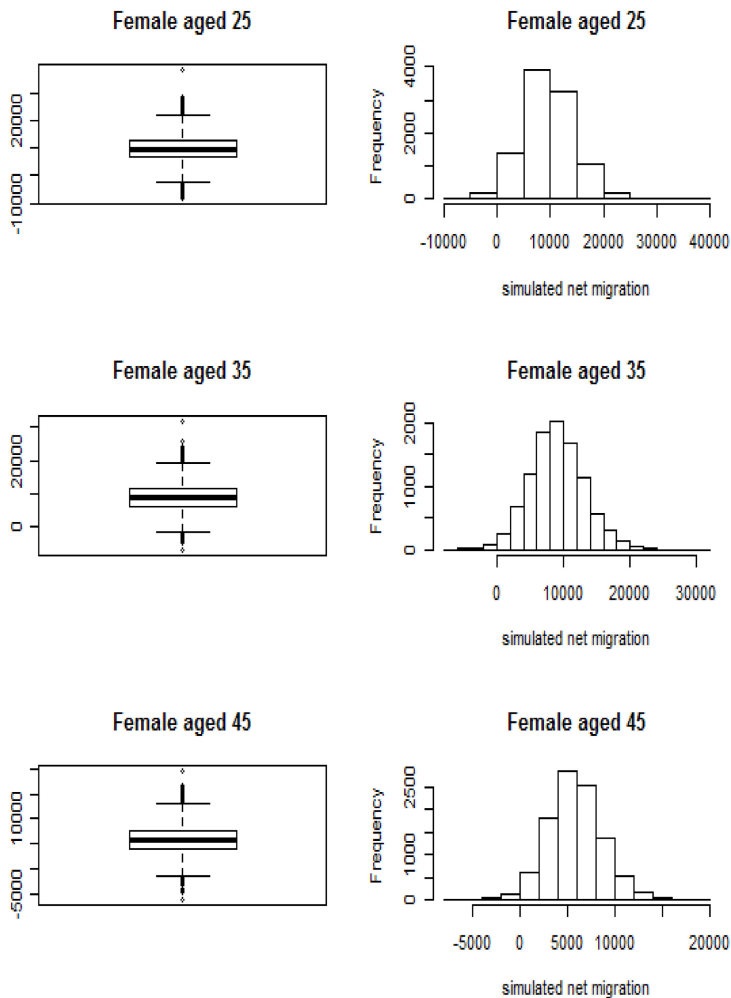


Fig. 4. Simulated female net migration, $x=25, 35$ and 45 years old

References

1. D'Amato, V., Piscopo, G., Russolillo, M.: Iterative Algorithms for detecting mortality trends in the family of Lee Carter Models. In: *Frontiers in Artificial Intelligence and Applications, Neural Nets WIRN 2010*, vol. 226, pp. 69–76. IOS Press – KBIES book series (2011)
2. Dawid, A.P.: Statistical Theory: The Prequential Approach. *Journal of the Royal Statistical Society A* 147(2), 278–292 (1984)
3. Efron, B., Tibshirani, R.J.: *An introduction to the bootstrap*. Chapman & Hall, New York (1993)

4. Human Mortality Database. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany), www.mortality.org or www.humanmortality.de (referred to the period 1952 to 2004)
5. Hyndman, R.J., Ullah, S.: Robust forecasting of mortality and fertility rates: a functional data approach. *Computational Statistics and Data Analysis* 51, 4942–4956 (2007)
6. Hyndman, R.J., Booth, H.: Stochastic population forecasts using functional data models for mortality, fertility and migration. *International Journal of Forecasting* 24(3), 323–342 (2008)
7. Koryś, P., Okólski, M.: Czas globalnych migracji. Mobilność i międzynarodowa perspektywa globalizacji (Time of global migration. International mobility in the perspective of globalisation). ISS Working Papers “Prace Migracyjne”, 55, Center for Migration Research, Warsaw (2004)
8. Kupiszewski, M.: Modelowanie dynamiki przemian ludności w warunkach wzrostu znaczenia migracji międzynarodowych (The role of international migration in the modelling of population dynamics). Institute of Geography and Spatial Organisation, Polish Academy of Sciences, Warsaw (2002b)
9. Lee, R.D., Carter, L.R.: Modelling and Forecasting U.S. Mortality. *Journal of the American Statistical Association* 87, 659–671 (1992)
10. NRC [National Research Council]: *Beyond Six Billion. Forecasting the World's Population*. National Academies Press, Washington, D.C. (2000)
11. Renshaw, A.E., Haberman, S.: Lee-Carter mortality forecasting with age specific enhancement. *Insurance: Mathematics and Economics* 33, 255–272 (2003c)
12. van der Gaag, N., van Wissen, L.: *Analysis and Forecasting of International Migration by Major Groups (Part II)*. Eurostat Working Paper 3(9), Eurostat, Luxembourg (1999)
13. <http://data.worldbank.org/indicator/SM.POP.NETM>
14. Istituto nazionale di statistica, <http://www.istat.it/en/>
15. <http://www.indexmundi.com/>

Simulation Framework in Fertility Projections

Valeria D'Amato¹, Gabriella Piscopo², and Maria Russolillo¹

¹ Department of Economics and Statistics,
University of Salerno, Campus Fisciano, Italy
vdamato,mrussolillo@unisa.it

² Department of Economics,
University of Genoa, Via Vivaldi, Genoa, Italy
gabriella.piscopo@unifi.it

Abstract. By 1951, average fertility had fallen to just over two children per woman, and only five percent of children would die in their first ten years of life. A similar pattern of declining fertility and mortality rates, collectively known as the demographic transition, has been observed in every industrializing country. Financial projections for Social Security systems depend on many demographic, economic and social factors as well as the reduction of fertility rates and the ageing of a population. In order to address the need to develop reliable projections, it is unavoidable to detect appropriate measures to represent the future trends of the quantities of interest. The aim of the paper is apply to Italian data a mathematical scheme suitable for projecting the fertility rates and for measuring the uncertainty around these estimates. Finally a numerical application is provided.

Keywords: Lee-Carter Models, Functional Demographic Model, Fertility Forecasting, Bootstrap.

1 Introduction

The fertility rate study is central if referred to demographic problems, but especially in social and economic field. For instance, the low rate of mortality has caused the reduction of the members in a family and contextually the aging of the population. The phenomena under consideration have brought strong imbalance in the composition per range of ages, in the distribution of the resources and in the welfare services. Furthermore the demographic changes affect also the sustainability of the pension cost. The observed trend in the industrialized countries, such as Italy, is the dramatic reduction of fertility. It causes a great stir also from health care point of view. In other words, the fertility forecasting have a key-influence in guiding the decision for the allocation of the future resources. In light of the previous considerations, it is necessary to detect measures methodologically appropriate and coherent to produce accurate projections. To this aim, several approaches to fertility forecasting have been proposed. Originally, time series methods have been used to forecast the total number of births

([10]). Afterwards, the attention has been given to the total fertility and independent age-specific rates, also forecast by time series methods. Recently, Ortega et al. [11] have analyzed common trends of the total fertility rates for clusters of countries using a country-specific dynamic factor model. Others ([3], [6], [7]) have proposed parameterized models using the gamma, beta and Hadwiger distributions, and the multi-exponential model. In this case, over-parameterization could invalidate the forecast. Lee [8] has forecasted fertility using the Lee-Carter method, pre-specifying the long-term mean value of total fertility because of structural change. Hyndman and Ullah [5] have applied a functional data approach to overcome this problem and capture the structural changes for different ages. In the context of the Italian literature, few contributions have been proposed on this topic. Caputo [2] have used a memory function with perturbation to forecast fertility rate. With respect to the state of art, our paper aims at take advantage of the use of functional models and measure the uncertainty in fertility forecasting for Italian data, by using a simulation scheme. The layout of the paper is the following: Section 2 introduce the simulation framework in the Lee Carter context, section 3 provide numerical evidences and comments by graphical analysis.

2 Backward Bootstrap on FDM

The well-known Lee-Carter model [9] is widely used for forecasting future death rates. In 1993 Lee has applied the model under consideration to the fertility rate to obtain projections also in this field. Hyndman and Ullah [5] have introduced a extension to this methodology to forecast age-specific mortality rates and fertility too, based on the combination of functional data analysis, nonparametric smoothing and robust statistics. The model is summarized in the following.

Let $p_t(x_i)$ be the fertility rate for x_i in year t and $y_t(x_i) = \log(p_t(x_i))$ be the log of the observed fertility rate for age x and year t , $f_t(x)$ the underlying smooth function, $\{x_i, y_t(x_i)\}, t = 1, \dots, n, i = 1, \dots, p$ the functional time series, where

$$y_t(x_i) = f_t(x_i) + \sigma_t(x_i)\varepsilon_{t,i} \quad (1)$$

where $\varepsilon_{t,i}$ is an iid standard normal random variable and the $\sigma_t(x_i)$ allows for the amount of noise to vary with x . The smoothing is carried out with a nonparametric method, based on weighted penalized regression splines, with a concave constraint, that is reasonable for fertility data. Then, the fitted curves is decomposed via a basis function expansion:

$$f_t(x) = \mu(x) + \sum_{k=1}^K \beta_{t,k} \varphi_k(x) + e_t(x) \quad (2)$$

where $\mu(x)$ is a measure of location of $f_t(x)$, $\{\varphi_k(x)\}$ is a set of orthonormal basis functions and $e_t(x) \sim (N(0, \text{var}(x)))$. The error term $e_t(x)$ given by the difference between the smoothed curves and the fitted curves from the model, is the modelling error.

In order to forecast $y_t(x)$, univariate time series models are fitted to each coefficients $\{\beta_{t,k}\}, k = 1, \dots, K$. Using them, the coefficients $\{\beta_{t,k}\}, k = 1, \dots, K$ are forecasted for $t = n + 1, \dots, n + h$. Finally, the previously obtained coefficients are implemented to get the $f_t(x)$ as in formula(2) and the $y_t(x)$ are projected. The estimated variance of error terms in (2) is used to compute prediction intervals for the forecast. Simulation has become one of the most widely applied tools of operations research.

In order to derive confidence intervals for the fertility rates, we resort to a non parametric bootstrap on $p_t(x_i)$ for creating artificial datasets upon fitting the aforementioned FDM model. Thus we obtain a large number of future sample paths of age-sex-specific population and vital event numbers which can be used to estimate the asymptotic properties of the fertility projections.

3 Numerical Application

We have run the application by considering the Italian female fertility rates for single year of age from 1952 to 2004 for women aged between 13 and 50 years old (the data are downloaded from demo.Istat). In particular, the fertility rates, defined as the number of live births during the calendar year, according to the age of the mother, per 1,000 of the female resident population of the same age at a certain date, are obtained as in [5]. Fig. 1 reflects the changing social conditions affecting fertility and shows the data as separate time series. To have an idea of the changing, we can notice there is an increase in fertility in lower ages around the 1960s and a rapid decrease in fertility after the 1970s in all ages. Moreover, looking at Fig. 1 is also evident a drop in fertility at lower ages against an increase in fertility at higher ages in more recent years, probably caused by a delay in child-bearing while careers are established.

The first step of the application consists in fitting the FDM version of the LC model to the data under consideration; the estimated parameters are shown in Figure 2 and the related error measures are reported in Table 1. Fig. 2 shows the fitted basis functions at the top and associated coefficients at the bottom. As described in [5], the basis functions shown in Fig. 2 are obtained by modelling the fertility rates of mothers in different age ranges, where a decomposition of order 4 has been used: the first basis function models young mothers in their 20s, the second one models late-mothers in their 40s, the third one models the mothers in their 30s and the last one models late-mothers in their 50s. The coefficients associated with each basis function outline the changing social conditions affecting fertility noted in Fig. 1: see in particular the increasing in coefficients around 1960 and their decrease during 1970s. Moreover, we note that the basis functions explain 79.3%, 17.1%, 1.3%, 0.8% of the variation respectively, leaving only 1.5% unexplained.

To provide general intuition about the goodness of the fitting, we resort to the traditional diagnostic for detecting the structure of the residuals. The findings of the error measures are shown in Table 1.




Fig. 1. Italian log fertility rates from 1952 to 2004

We keep running the application by calculating the forecasted fertility rates from 2005 to 2024. From the results shown in Fig. 3, it is clear that the forecasts result in a decrement in fertility rates for lower ages and in an increase in higher ages.

Finally, we implement a non parametric bootstrap (4) on the 2024 projected fertility rates for ages $x = 25, 35, 45$ in order to derive, for each projection, confidence intervals at level of 90%; the bootstrap uses the re-sampled error from the fitted model. We generate $b = 100$ bootstrap samples. The obtained confidence intervals are: $0.0127991 < x < 0.0429024$ for $x = 25$; $0.09007308 < x < 0.2210804$, for $x = 35$ and $0.001887952 < x < 0.02110152$ for $x = 45$. Graphically, these results are shown in Fig. 4, Fig. 5, Fig. 6.

4 Final Remarks

Many approaches to fertility forecasting have been developed. Functional demographic models are based on the methodology of extrapolation of time series. They have the advantage of providing a flexible framework. On the opposite, the main limit of extrapolative methods is that the trends of the past will be continued into the future. However, this assumption has proved to be a better basis for forecasting than either structural modeling involving exogenous variables or methods based on expectation (11). For this reason, we have decided to apply the functional demographic model to Italian data. There are just few contributions on this topic in Italy (12). As regards the methodology used, Caputo uses

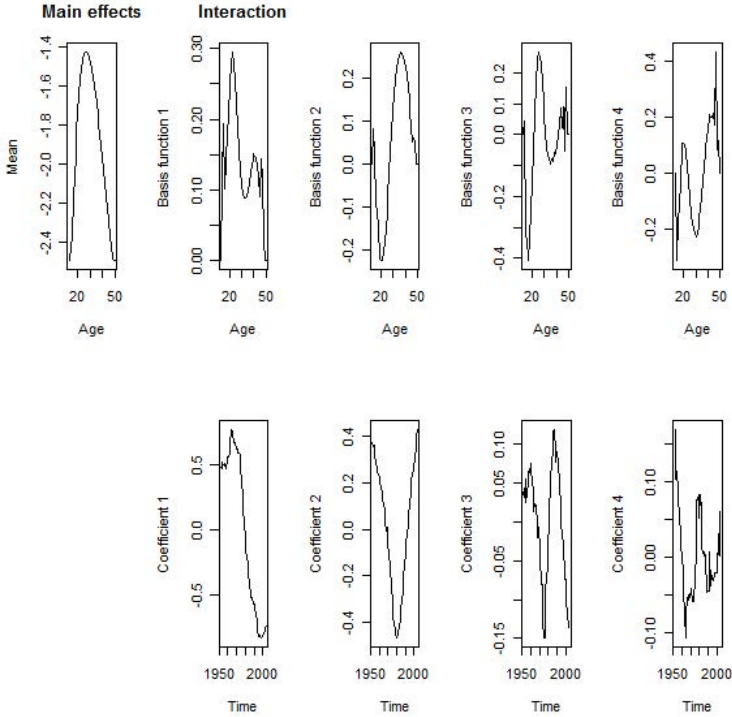


Fig. 2. Basis functions and associated coefficients for the data shown in Figures 1

Table 1. Error Measures: Mean Error (ME), Mean Square Error (MSE), Mean Percentage Error (MPE), Mean Absolute Percentage Error (MAPE), Integrated Error (IE), Integrated Square Error (ISE), Integrated Percentage Error (IPE), Integrated Absolute Percentage Error (IAPE)

Average across ages				
	ME	MSE	MPE	MAPE
	-0.00001	0.00019	0.00007	0.00404
Average across years				
	IE	ISE	IPE	IAPE
	-0.00032	0.00647	0.00259	0.15478

a memory function with perturbation; differently, we use several orthogonormal components to capture the structural changes in the data. The use of several components is more important for fertility, because the first component explains a smaller proportion of the total variation of the data and the additional components may serve to incorporate recent changes in patterns for different ages. We are confident that the different structural changes in correspondence with




Fig. 3. Forecasts of Italian fertility rates from 2005 to 2024 (the curves are shown from left to right)




Fig. 4. Confidence intervals at 90%, $x=25$ years old

different ages have to be modeled with different functions to obtain more accurate estimates and forecasts; the functional demographic model allows us to do this. As regards the results, we have found a reduction in fertility for lower ages in line with the results of [2]. An innovative result is given by the implementation of bootstrap procedure, which offers the probabilistic distribution of prediction intervals of fertility forecast. Similar results have been found for other countries by Hyndman and Ullah ([5]).




Fig. 5. Confidence intervals at 90%, $x=35$ years old




Fig. 6. Confidence intervals at 90%, $x=45$ years old

References

1. Booth, H.: Demographic Forecasting: 1980 to 2005 in review. *International Journal of Forecasting* 22(3), 547–581 (2005)
2. Caputo, M., Nicotra, N., Gloria-Bottini, E.: Fertility transition: forecast for demography. *Human Biology* 80(4), 359–376 (2008)
3. Chandola, T., Coleman, D.A., Hiorns, R.W.: Recent European Fertility Patterns: Fitting Curves to Distorted Distributions. *Population Studies* 53(3), 317–329 (1999)
4. Efron, B., Tibshirani, R.J.: An introduction to the bootstrap. Chapman & Hall, New York (1993)
5. Hyndman, R.J., Ullah, M.S.: Robust forecasting of mortality and fertility rates: a functional data approach. *Computational Statistics & Data Analysis* 51, 4942–4956 (2007)
6. Keilman, N., Pham, D.Q.: Predictive intervals for age-specific fertility. *European Journal of Population* 16, 41–66 (2000)
7. Kostaki, A., Peristera, P.: Modeling fertility in modern populations. *Demographic Research* 16, 141–194 (2007)
8. Lee, R.D.: Modeling and forecasting the time series of U.S. fertility: age distribution, range, and ultimate level. *International Journal of Forecasting* 9, 187–202 (1993)
9. Lee, R.D., Carter, L.R.: Modelling and Forecasting U.S. Mortality. *Journal of the American Statistical Association* 87, 659–671 (1992)
10. McDonald, J.: A time series approach to forecasting Australian total live births. *Demography* 16, 575–601 (1979)
11. Ortega, J.A., Poncela, P.: Joint forecasts of southern European fertility rates with nonstationary dynamic factor models. *International Journal of Forecasting* 21(3), 539–550 (2005)

Building a Global Performance Indicator to Evaluate Academic Activity Using Fuzzy Measures

Marta Cardin¹, Marco Corazza^{1,3}, Stefania Funari², and Silvio Giove¹

¹ Department of Economics, Ca' Foscari University of Venice

² Department of Management, Ca' Foscari University of Venice

³ Advanced School of Economics of Venice

Cannaregio 873, 30121 Venice, Italy

{mcardin, corazza, funari, sgiove}@unive.it

Abstract. The aim of this note is to provide a global performance index that allows to evaluate the performance of each faculty member and which is able to consider the multidimensional nature of the academic activity in terms of research, teaching and other activities that academics should/might exercise. In order to model also the case in which there could be synergic and redundant connections among the different areas of the academic activity, we propose to use fuzzy measures and the Choquet integral as an aggregator of the different components.

Keywords: Fuzzy measures, Choquet integral, performance evaluation.

«Work. Finish. Publish.»

Michael Faraday, 22 September 1791 – 25 August 1867

1 Introduction

The evaluation of the performance of the academic activity has gained a growing interest in recent years and many Universities undertake periodic assessments of their faculty members which act as a first step for a wider performance evaluation process involving, at different levels, departments, faculties and Universities.

Regarding in particular the research activity, a vast literature is interested in the measurement of the research output of scholars. Several approaches have indeed been proposed in order to assess the scientific production.

We remind the use of bibliometric indicators, the well-known h -index proposed by Hirsh [6] and other indices that complement and extend it, such as the g -index, the m -index, the $h^{(2)}$ -index, the h_α -index, the A -index, the R -index, the normalized h -index, the dynamic h -type index and many others indicators related to the h -index (see [1] for a review). This great variety of research output measures, even if with some differences, intends to capture both the quantitative and the qualitative dimension of the research, summarized in the number of publications and the number of citations that each publication receives, by focusing in many cases on *the most productive core* of a researcher ([3]).

The bibliometric indices allow to obtain a ranking of the authors; however, having each indicator some advantages but also some drawbacks, none of them can be considered the best measure of the research outputs. Without emphasizing a particular advantage or drawback, Marchant [8] proposed a general theory of bibliometric ranking, which explores the properties of the scoring rules that permit to rank the authors by taking into account all their publications, although weighted by some partial scores.

There is therefore a vast debate on the performance measurement of the research production; on the other hand, turning towards the teaching area, we may note that the discussion on the assessment of the teaching activity is as much active.

At present, many university departments have activated some teacher evaluation systems, but also in this case there is no a uniform system of indicators fully recognized. A reasonable solution is to consider not only the classroom teaching and thus the student ratings as a measure of teaching effectiveness, rather to take account of all educational activities of a faculty member, including the development of new courses, the restyle of the old ones, the use of innovative teaching strategies, the offer of seminars, workshops, and so on. Therefore, experts in the field of teacher evaluation tend to prefer the combined use of multiple approaches to evaluate faculty teaching (see for instance [7], [5]).

Although both research and teaching are important components of the university production, the literature on the university performance mainly focused separately either on the evaluation of the research dimension or on the teaching component, obtaining a partial view of the goodness of the academic activity. In order to provide an overall assessment of the performance of each faculty member, so that to capture the multidimensional nature of the evaluation problem, it should be more desirable to simultaneously consider the different faces of the academic activity, shared mainly among research, teaching and service.

Based on this thought, Bana e Costa and Oliveira ([2]) have recently proposed a faculty evaluation model that takes into consideration the whole range of academic activities pursued by each faculty member. They used a multicriteria decision analysis based on value functions, in order to associate to each faculty member an overall score, according to which each member of the faculty can be finally assigned to one of the four rating categories (inadequate, sufficient, relevant and excellent). The authors used value functions which transform individual performance into value for the university and used a hierarchical aggregation procedure which involves first an aggregation of the criteria (in terms of value) within each area of activity and then an aggregation of values across areas, leading finally to an overall score for each faculty member.

However, when different evaluation criteria within each area are simultaneously considered, some synergic effects may arise. For example, referring to the research area, Tucci, Fontani and Ferrini ([12]) noted that the joined activity of the researcher in terms of publishing and of accomplishment of other activities (as for example organizing meetings and conferences, attending to seminars, coordinating research groups, being editor and being referee), should be awarded in the computation of the *R*-Factor index that measures the research output. The idea is that the involvement in several activities of the researcher requires an additional effort that should be considered in the researcher's performance evaluation.

A similar synergic behaviour can also occur among the different areas of the academic activity. For example, why not reward (in terms of performance score) an academic that performs well both in research and teaching or even in all the areas considered? In this case a sort of premium on performance could be considered in addition to the partial performance evaluation scores, which concurs to increase the final overall performance score. Otherwise, academics working on management activity but who do not directly contribute to publishing activity, are penalized.

In decision making problems, a natural approach that allows to model the interaction among criteria is the use of the Choquet integral with respect to a fuzzy measure, which is proved to be an adequate aggregator operator that extends the classical weighted arithmetic mean when the criteria adopted interact ([9]).

In this contribution, we propose to use fuzzy measures and the Choquet integral in order to compute a global performance index which allows to evaluate the performance of faculty members by taking into account multiple areas of the academic activity, multiple evaluation criteria within each area and also their possible connections. In particular, we adopt a two-step Choquet integral which permits to model a hierarchy in the aggregation process.

The rest of the paper is structured as follows. Section 2 reminds the concepts of Choquet integral and fuzzy measures and that of multi-step Choquet integral. Section 3 first shows how the Choquet integral can be used in the performance evaluation process of the academic activity and then presents a numerical example. Section 4 reports some final remarks.

2 Non Additive Measures for Aggregation

Non additive measures are up to now a commonly used method to to represent interactions between the elements of a set. In cooperation with integral aggregation functions, they are a well-founded framework able to aggregate information from several sources. Usually the monotonicity property is required in most of practical applications.

In particular in this paper we consider non-additive monotonic measures, also called fuzzy measures and Choquet integral for a n -dimensional vector as in [9] for example. Let N be a finite index set, $N = \{1, \dots, n\}$.

Definition 1. A set function $v : 2^N \rightarrow [0, 1]$, with $v(\emptyset) = 0$, $v(N) = 1$ and $A \subseteq B \subseteq N$ implies that $v(A) \leq v(B)$, is called a non-additive measure on N .

We note that if $S \subseteq N$, $v(S)$ can be viewed as the importance of the set of elements S . Note also that non-additive measures encompass probability measures, belief functions and capacities.

Let now introduce the discrete Choquet integral on N viewed as an aggregation function that generalizes the weighted arithmetic means.

Definition 2. Let v a non-additive measure $v : 2^N \rightarrow [0, 1]$ and $x \in \mathbb{R}^n$. Let σ be a permutation of N such that $\{x_{\sigma(n)}\}$ forms a non decreasing sequence and we define $x_{\sigma(0)} = 0$. Then the Choquet integral of x is:

$$C_v(x) = \int x dv = \sum_{i=1}^n (x_i - x_{i-1})v(\{\sigma(i), \dots, \sigma(n)\}) . \quad (1)$$

This formula can be interpreted as an expectation operator with respect to a generalized measure. Notice that in the case in which the measure $v(\{\cdot, \dots, \cdot\})$ is additive it is immediate to prove that (1) is the well known arithmetic weighted mean.

Now we give some basic definitions and on multi-step Choquet integral. The two-step Choquet integral has been investigated mainly in [10] and [11]. Let us now give a formal definition of a multi-step Choquet integral.

Definition 3. Let $\Gamma \subseteq \mathbb{R}^n$. For any $i \in N$, the projection $\mathbf{x} \mapsto x_i$ is a 0-step Choquet integral. Let us consider $F_i : \Gamma \rightarrow \mathbb{R}$, $i \in M := \{1, \dots, m\}$, being k_i -step Choquet integrals, and a non-additive measure v on M . Then

$$F(\mathbf{x}) := C_v(F_1(\mathbf{x}), \dots, F_m(\mathbf{x}))$$

is a k -step Choquet integral, with $k := \max\{k_1, \dots, k_m\} + 1$. A multi-step Choquet integral is a k -step Choquet integral for some integer $k > 1$.

3 Performance Evaluation of Academic Activity: A Numerical Example

The aim of this section is twofold:

- to illustrate, in a synthetic way, the main academic activities, the related criteria, their hierarchical organization and their interactions that, in our opinion, have to take into account in order to evaluate an academic subject (a scholar, a department, a faculty, ...);
- to provide an exemplification about the “mechanics” of our two-step evaluation procedure.

As far as the first point is concerned, we propose an extension of the Research Factor index (on following: R), recently introduced in [12] for the evaluation of the research quality of an academic subject. In short, the R is obtained by a two-level process. In the first level this index is determined by a (simple) aggregation of two macro-criteria and their interaction; in the second level each of these macro-criteria is achieved by (simple) aggregation of various criteria and their interactions.

Given these bases, in order to globally evaluate the activities of an academic subject, we propose a Global index (on following indicated by C_G) obtained by a two-step Choquet integral based aggregation procedure. In particular, going backward, in the second step the C_G is achieved by the aggregation of three indexes and their interactions. These indexes are respectively related to the research (from which the x_R index), to the teaching (from which the x_T index) and to services¹ (from which the x_S index). We recall

¹ By “services” we mean, for instance, consultancy activities, membership on non-academic committees and so on.




Fig. 1. Two-steps evaluation process

that we are interested in the evaluation of the overall quality of an academic subject. This is the main reason for which we consider all the various typologies of academic activities.

Formally, we can represent the output of this second step as follows:

$$C_G = C_v(x_R, x_T, x_S),$$

where v indicates the non-additive measure associated to the three indexes.

In the first step each of these indexes is determined by the aggregation of various criteria and the interactions among them.

Figure 1 schematizes the two-steps evaluation process.

Of course, also each of these criteria could be obtained by the aggregations of suitable sub-criteria, so leading to a three-step Choquet integral based aggregation process (the generalization to an n -step evaluation procedure is obvious).

About the choice of the criteria for each of the considered indexes, there exists a wide enough literature, and our personal expertises, from which drawing inspiration (see, for instance, [2], [5], [6], [8] and [12]). As the specification of such criteria does not constitute a goal of this note, the list of them we take into account has to be simply considered as exemplificative and not prescriptive. In any case, as general recommendation, we suggest that the number of criteria chosen for each index is not particularly large given

the “curse of dimensionality” implied in the specification of the non-additive measures respectively related to the various indexes themselves (see, for details, definition 2).

Now we can present the list of these criteria. Notice that they are generally relating to a prefixed time period (for instance, the last five academic years). With reference to the research activities, we consider the following $N_R = 3$ criteria:

- $y_{R,1}$: an indicator of the quantity and the quality of the publications;
- $y_{R,2}$: an indicator of the numbers of the self-citations and the others-citations received by those publications;
- $y_{R,3}$: an indicator of the quantity and the quality of the coauthors of those publications;

(see, for more details on these criteria, [4]). With reference to the teaching activities, we consider the following $N_T = 3$ criteria:

- $y_{T,1}$: an indicator of the quantity and the quality of the degree courses;
- $y_{T,2}$: an indicator of the quantity and the quality of the master and the Ph.D. courses;
- $y_{T,3}$: an indicator of the student evaluations.

Finally, with reference to the service activities, we consider the following $N_S = 2$ criteria:

- $y_{S,1}$: an indicator of the quantity and the quality of the consultancy activities;
- $y_{S,2}$: an indicator of the numbers of membership on non-academic committees.

Formally, we can represent the outputs of this first step as follows:

$$x_j = C_{v_j} (y_{j,1}, \dots, y_{j,N_j}), \text{ for all } j \in \{R, T, S\},$$

where v_j s are the non-additive measures respectively associated to the various subset of multicriteria.

At this point we can pass to deal with the second of the two starting points, that is to give a numerical example of our approach. Since now we premise that in this application we do not use real data. However, the considered data are generated in such a manner to embody our views about the interactions among the three indexes (second step) and among the criteria respectively relating to each of these indexes (first step).

Now we first need to specify the transformations, defined in section 2, of the non-additive measures associated to the three indexes (see Table 1) and to the various subset of multicriteria (see Table 2).

We emphasize that our main views incorporated in $\pi(\cdot)$, $\pi_R(\cdot)$, $\pi_T(\cdot)$ and $\pi_S(\cdot)$ are:

- “*to do more is better than to do less*”, in fact all the considered measures are strictly increasing with respect to their respective arguments (for instance, $\pi_T(y_{T,2}) + \pi_T(y_{T,3}) < \pi_T(y_{T,2}, y_{T,3})$);
- “*academic activities are better than non-academic activities*”, in fact, among all the considered measures, the only case of subadditivity is related to the one associated to the three indexes with respect to the service inputs (for instance, $\max\{\pi(x_T), \pi(x_S)\} < \pi(x_T, x_S) < \pi(x_T) + \pi(x_S)$).

Table 1. Transformations of the measures associated to the three indexes

$\pi(\cdot)$
$\pi(\emptyset) = 0.00$
$\pi(x_R) = 0.35$
$\pi(x_T) = 0.30$
$\pi(x_S) = 0.15$
$\pi(x_R, x_T) = 0.75$
$\pi(x_R, x_S) = 0.45$
$\pi(x_T, x_S) = 0.40$
$\pi(Universe) = 1.00$

Table 2. Transformations of the measures associated to the various subset of multicriteria

$\pi_R(\cdot)$	$\pi_{TFi}(\cdot)$	$\pi_S(\cdot)$
$\pi_R(\emptyset) = 0.00$	$\pi_T(\emptyset) = 0.00$	$\pi_S(\emptyset) = 0.00$
$\pi_R(y_{R,1}) = 0.50$	$\pi_T(y_{T,1}) = 0.45$	$\pi_S(y_{S,1}) = 0.55$
$\pi_R(y_{R,2}) = 0.35$	$\pi_T(y_{T,2}) = 0.30$	$\pi_S(y_{S,2}) = 0.45$
$\pi_R(y_{R,3}) = 0.20$	$\pi_T(y_{T,3}) = 0.25$	$\pi_S(Universe) = 1.00$
$\pi_R(y_{R,1}, y_{R,2}) = 0.95$	$\pi_T(y_{T,1}, y_{T,2}) = 0.85$	-
$\pi_R(y_{R,1}, y_{R,3}) = 0.80$	$\pi_T(y_{T,1}, y_{T,3}) = 0.70$	-
$\pi_R(y_{R,2}, y_{R,3}) = 0.65$	$\pi_T(y_{T,2}, y_{T,3}) = 0.65$	-
$\pi_R(Universe) = 1.00$	$\pi_T(Universe) = 1.00$	-

Then, in order to make directly comparable the various criteria among them, we need to suitably normalize the criteria themselves in a prefixed interval, let us say $[0, 1]$. Also this aspect does not constitute a goal of this note. Anyway, a suggestion for a possible normalization function can be a piecewise linear function of the following type:

$$y = \begin{cases} 0 & \text{if } x \leq x_{min} \\ a + bx & \text{if } x_{min} < x \leq x_{max}, \text{ with } x_{min} < x_{max}, \\ 1 & \text{if } x > x_{max} \end{cases}$$

where y indicates the normalized criterion, x indicates the non-normalized criterion, $a = 1 / (x_{max} - x_{min}) > 0$ and $b = -x_{min} / (x_{max} - x_{min}) < 0$.

Notice that both the specification of the non-additive measures and the definition of the normalized functions used in our approach should generally be performed by teams of experts or by focus groups.

Now, we can provide the numerical exemplification. Let us consider two academic subjects, respectively characterized by the following values of the criteria:

- $y_{R,1}^1 = 0.6, y_{R,2}^1 = 0.4, y_{R,3}^1 = 0.2, y_{T,1}^1 = 0.5, y_{T,2}^1 = 0.1, y_{T,3}^1 = 0.3, y_{S,1}^1 = 0.2, y_{S,2}^1 = 0.4;$
- $y_{R,1}^2 = 0.5, y_{R,2}^2 = 0.1, y_{R,3}^2 = 0.3, y_{T,1}^2 = 0.6, y_{T,2}^2 = 0.4, y_{T,3}^2 = 0.2, y_{S,1}^2 = 0.2, y_{S,2}^2 = 0.4.$

where the superscript identifies the academic subject.

Notice that the difference between the two academic subjects consists in the “exchange” of the values associated to the research criteria ($y_{R,1}$, $y_{R,2}$ and $y_{R,3}$) with the values associated to the teaching criteria ($y_{T,1}$, $y_{T,2}$ and $y_{T,3}$). In the following Table 3 we report in details for both the academic subjects the calculations related to our approach.

Table 3. Calculations related to our two-step Choquet integral based aggregation procedure

Step	Choquet integral
First	$x_R^1 = (0.20 - 0.00)1.00 + (0.40 - 0.20)0.95 + (0.60 - 0.40)0.50 = 0.4900$
	$x_T^1 = (0.10 - 0.00)1.00 + (0.30 - 0.10)0.70 + (0.50 - 0.30)0.45 = 0.3300$
	$x_S^1 = (0.20 - 0.00)1.00 + (0.40 - 0.20)0.45 = 0.2900$
	$x_R^2 = (0.10 - 0.00)1.00 + (0.30 - 0.10)0.80 + (0.50 - 0.30)0.50 = 0.3600$
	$x_T^2 = (0.20 - 0.00)1.00 + (0.40 - 0.20)0.85 + (0.60 - 0.40)0.45 = 0.4600$
Second	$C_G^1 = (0.29 - 0.00)1.00 + (0.33 - 0.29)0.75 + (0.49 - 0.33)0.35 = 0.3760$
	$C_G^2 = (0.29 - 0.00)1.00 + (0.36 - 0.29)0.75 + (0.46 - 0.36)0.30 = 0.3725$

As $C_G^1 > C_G^2$, academic subject 1 is “better” than academic subject 2.

4 Final Remarks

In this paper we propose a Choquet integral based method to evaluate academic performance. With respect to the previous literature one of the main contribution of our proposal concerns the computation of a numerical index which takes into account not only the scientific productivity but also all the other activities linked to academic functions. Again the non linear computation method permits to consider interactions among the criteria in a formal way, and, at the same time, to realize an easy approach to understand, implement and apply. We fill that this Choquet based linear aggregation algorithm will constitute an innovative way to approach this kind of evaluation method in contrast with most of the methods actual in use.

References

1. Alonso, S., Cabrerizo, F.J., Herrera-Viedma, E., Herrera, F.: h-Index: A review focused in its variants, computation and standardization for different scientific fields. *Journal of Informetrics* 3, 273–289 (2009)
2. Bana e Costa, C.A., Oliveira, M.D.: A multicriteria decision analysis model for faculty evaluation. *Omega* 40, 424–436 (2012)
3. Burrell, Q.: On the h index, the size of the Hirsh core and Jin’s A index. *Journal of Informetrics* 1, 170–177 (2007)
4. Cardin, M., Corazza, M., Funari, S., Giove, S.: A fuzzy-based scoring rule for author ranking. An alternative of h index. In: Apolloni, B., Bassis, S., Esposito, A., Morabito, C.F. (eds.) *Neural Nets WIRN 2011. Proceedings of the 21st Italian Workshop on Neural Nets. Frontiers in Artificial Intelligence and Applications*, vol. 234, pp. 36–45. IOS Press (2011)

5. Felder, R.M., Brent, R.: How to evaluate teaching. *Chemical Engineering Education* 38, 200–202 (2004)
6. Hirsch, J.: An index to quantify an individual's scientific research output. *Proceedings of the National Academy of Sciences* 102, 16569–16572 (2005)
7. Hoyt, D.P., Pallett, W.H.: Appraising Teaching Effectiveness: Beyond Student Ratings. *IDEA PAPER* 36 (1999)
8. Marchant, T.: Score-based bibliometric rankings of authors. *Journal of the American Society for Information Science and Technology* 60, 1132–1137 (2009)
9. Marichal, J.L.: Aggregation of interacting criteria by means of the discrete Choquet integral. In: Calvo, T., Mayor, G., Mesiar, R. (eds.) *Aggregation Operators: New Trends and Applications*. *STUDFUZZ*, vol. 97, pp. 224–244. Physica-Verlag (2002)
10. Mesiar, R., Vivona, D.: Two-step integral with respect to fuzzy measure. *Tatra Mountains Mathematical Publications* 16, 359–368 (2006)
11. Narukawa, Y., Murofushi, T.: The n -step Choquet integral on finite spaces. In: *Proceedings of the 9th International Conference on Information Processes and Management of Uncertainty in Knowledge-Based Systems*, Annecy, France (2002)
12. Tucci, M.P., Fontani, S., Ferrini, S.: L'R-Factor: un nuovo modo di valutare l'attività di ricerca. *Studi e Note di Economia* 1, 103–140 (2010)

Testing the Weak Form Market Efficiency: Empirical Evidence from the Italian Stock Exchange

Giuseppina Albano, Michele La Rocca, and Cira Perna

Dept. of Economics and Statistics, University of Salerno,
via Ponte don Melillo, 84084 Fisciano (SA), Italy
{pialbano,larocca,perna}@unisa.it

Abstract. This paper investigates the use of feedforward neural networks for testing the weak form market efficiency. In contrast to approaches that compare out-of-sample predictions of non-linear models to those generated by the random walk model, we directly focus on testing for unpredictability by considering the null hypothesis that a given set of past lags has no effect on current returns. To avoid the data-snooping problem the testing procedure is based on the StepM approach in order to control the familiwise error rate. The procedure is used to test for predictive power in FTSE-MIB index of the italian stock market.

Keywords: Random walk, Market efficiency, Multiple testing scheme, Resampling methods.

1 Introduction and Background

Efficiency of financial markets is certainly one of the most controversial issue in finance. Volatility, predictability, speculation and anomalies in financial markets are also related to the efficiency issue and are all interdependent. The efficient market hypothesis is a concept of informational efficiency and refers to markets ability to process information into prices. The idea at the bases of the efficient market hypothesis (EMH) emerges already at 1900 but a formal definition of market efficiency is due to Fama who distinguish three forms of efficiency: weak form, semi-strong form and strong form. Weak form of efficiency assumes that actual price of an asset incorporates the past prices information. Thus, it will be not possible for investors, using past prices, to discover undervalued stocks and develop strategies to systematically earn abnormal returns. The semi-strong form hypothesis assumes that all publicly available information is incorporated in the price of the asset and investors cannot take advantage of this information, winning abnormal returns. Finally, strong form of market efficiency hypothesis is more restrictive than the previous two, maintaining that all information, public or private, is incorporated in the current stock prices and investors cannot systematically earn abnormal returns. From that time on, various extensions

of Fama's definition were proposed to include different levels of information and transaction costs. Anyway, even if there are various critics towards Fama's definition, it is still the most commonly used standard and benchmark for determining market efficiency.

In the following we will focus on the weak form efficiency. Historically, there was a very close link between EMH and the random-walk model and martingale. An *orthogonality* condition where all versions of the random walk and martingales hypotheses are captured is the following ([2]):

$$\text{cov}[f(r_t), g(r_{t+k})] = 0 \quad (1)$$

where $f(\cdot)$ and $g(\cdot)$ are two arbitrary functions, r_t and r_{t+k} are asset's returns at t and $t+k$ ($k \neq 0$). For example, if (1) holds when $f(\cdot)$ and $g(\cdot)$ are linear functions, then returns are serially uncorrelated. Alternatively, if $f(\cdot)$ is unrestricted but $g(\cdot)$ is restricted to be linear, the (1) is equivalent to a martingale hypothesis. Finally, if (1) holds for all $f(\cdot)$ and $g(\cdot)$ the returns are mutually independent.

A large body of literature has accumulated in order to test market efficiency using different approaches based on RW and martingale hypotheses (see, for example, [17] or [12] for a nice survey). Further, results are often puzzling since they contradict the conventional wisdom that all developed markets should be more efficient in incorporating information into prices than those markets from the developing economies. For example, Malkiel in [16] justifies return "anomalies" in the major stock markets as chance results that tend to disappear in the long term with a reasonable change in methodology, hence supporting the view that mature capital markets are generally efficient. In contrast, Shiller in [20] has to distance himself from the "presumption" that financial markets always work well and that price changes always reflect genuine information.

Anyway, evidence against the random walk hypothesis (RWH) for stock returns in the capital markets is often shown (see, for example, [3], [14], [15] and references therein). Failure of models based on linear time series techniques to deliver superior forecasts to the simple random walk model has forced researchers to use various non-linear techniques, such as Engle test, Tsay test, Hinich bispectrum test, Lyapunov exponent test. Also in such a literature the market efficiency confirms to be a changelling issue in finance (see, for example, [1]).

The aim of this paper is to test the weak form market efficient hypothesis for the italian stock exchange. According to the evidence of non-linear patterns in stock markets, we propose a neural network approach for characterizing and analyzing the closing price of the FTSE-MIB index, from 1/12/2003 to 23/03/2012.

In particular, we focus on a multiple testing scheme in which the null hypothesis specifies that information contained in past returns cannot be used to predict the current returns. To avoid the data-snooping problem and in order to control the familiwise error rate, the testing procedure uses a multiple testing scheme. Given the probabilistic complexity of the neural network model, a resampling technique is proposed to calibrate the test. The procedure is used to test for predictive power in FTSE-MIB index of the italian stock market.

The paper is organized as follows. In Section 2 we describe the structure of the data generating process and the neural network model employed. In Section

3 we propose a testing scheme for the weak form market efficiency based on a multiple testing scheme. In Section 4 we discuss the application of the proposed procedure to Italian stock market. Some concluding remarks close the paper.

2 Neural Network Modelling for Financial Returns

Following a standard practice, we construct return time series as $Y_t = \nabla \log S_t$, where S_t is the stock index at time t , in order to avoid potential problems associated with estimation of nonstationary regression functions. The model is defined as:

$$Y_t = g(Y_{t-1}, Y_{t-2}, \dots, Y_{t-d}) + \varepsilon_t \tag{2}$$

where $g(\cdot)$ is a non-linear function and ε_t is zero-mean error term with finite variance. The unknown function $g(\cdot)$ can be estimated by using a feedforward neural network with d input neurons:

$$f(\mathbf{y}, \mathbf{w}) = w_{00} + \sum_{j=1}^r w_{0j} \psi(\tilde{\mathbf{y}}^T \mathbf{w}_{1j}) \tag{3}$$

where $\mathbf{w} \equiv (w_{00}, w_{01}, \dots, w_{0r}, \mathbf{w}_{11}^T, \dots, \mathbf{w}_{1r}^T)^T$ is a $r(d+1) + 1$ vector of network weights, $\mathbf{w} \in \mathbf{W}$ with \mathbf{W} being a compact subset of $\mathbb{R}^{r(d+1)+1}$ and $\tilde{\mathbf{y}} \equiv (1, \mathbf{y}^T)^T$ is the input vector augmented by a bias component 1. The network (3) has d input neurons, r neurons in the hidden layer and the identity function for the output layer. The (fixed) hidden unit activation function ψ is chosen in such a way that $f(\mathbf{x}, \cdot) : \mathbf{W} \rightarrow \mathbf{R}$ is continuous for each x in the support of the explanatory variables and $f(\cdot, \mathbf{w}) : \mathbf{R}^d \rightarrow \mathbf{R}$ is measurable for each \mathbf{w} in \mathbf{W} .

Single hidden layer feedforward neural networks have a very flexible non-linear functional form. The activation functions can be chosen quite arbitrarily and it can be shown that they can arbitrarily closely approximate (in the appropriate corresponding metric) to continuous, or to p -th power integrable, non-linear functions $g(\cdot)$, so long as the activation function ψ is bounded and satisfies the necessary conditions of not being a polynomial [13].

Given a training set of n observations, estimation of the network weights (learning) is obtained by solving the optimization problem

$$\min_{\mathbf{w} \in \mathbf{W}} \frac{1}{n} \sum_{t=1}^n q(Y_t, f(Y_{t-1}, Y_{t-2}, \dots, Y_{t-d}; \mathbf{w})) \tag{4}$$

where $q(\cdot)$ is a proper chosen loss function, usually a quadratic function. From an operational point of view, estimates of the parameter vector \mathbf{w} can be obtained by using non-recursive or recursive estimation methods such as the back-propagation (BP) algorithm and Newton's algorithm.

Under general regularity conditions, both the methods deliver a consistent weight vector estimator. That is, a weight vector $\hat{\mathbf{w}}_n$ solving equation (4) exists and converges almost surely to \mathbf{w}_0 , which solves

$$\min_{\mathbf{w} \in \mathbf{W}} \int q(y, f(\mathbf{x}, \mathbf{w})) d\pi(\mathbf{z}) \tag{5}$$

provided that the integral exists and the optimization problem has a unique solution vector interior to \mathbf{W} . Moreover, under general regularity conditions the weight vector estimator is asymptotically normally distributed.

Therefore, given a consistent estimator of the asymptotic variance covariance matrix of the weights, one could be tempted to test hypotheses about the connection strengths which would be of great help in defining pruning strategies with a strong inferential basis. This approach in any case could be misleading. The parameters (weights) of the neural network model have no clear interpretation and this makes this class of models completely different from classical non-linear parametric models. Moreover, the same output of the network can be obtained with very different configurations of the weights. Finally, as a model selection strategy, variable selection and hidden layer size should follow different schemes: variables have a clear interpretation while hidden layer size has no clear meaning and should be considered a smoothing parameter which is fixed to control the trade-off between bias and variability.

As a consequence, we advocate a model selection strategy where an informative set of input variables is selected by looking at its relevance to the model, in a statistical significance testing framework, while the hidden layer size is selected by looking at the fitting or predictive ability of the network.

In this respect the Predictive Stochastic Complexity (PSC) index proposed by Rissanen and already used by Kuan and Liu ([8]) proved to be an effective tool to select appropriate hidden layer size. However, other tools based on "honest" prediction errors $Y_t - f(Y_{t-1}, \dots, Y_{t-d}; \hat{\mathbf{w}}_t)$ can be used as well, where $\hat{\mathbf{w}}_t$ is computed by using information up to time $t - 1$. The prediction error is labeled as "honest" in the sense that no information at time t or beyond is used to calculate the estimated parameter vector $\hat{\mathbf{w}}_t$.

The general idea behind variable relevance analysis is to compute some measures that can be used to quantify the relevance of explanatory variables with respect to a given model. Following White and Racine ([24]) and La Rocca and Perna ([9],[10],[11]), the hypotheses that the j -th lag, let's say Y_{t-j} , has no effect on Y_t , in model (2), can be formulated as:

$$\frac{\partial g(Y_{t-1}, Y_{t-2}, \dots, Y_{t-d})}{\partial Y_{t-j}} = 0. \quad (6)$$

Of course the function g is unknown but we equivalently investigate the hypotheses

$$f_j(Y_{t-1}, Y_{t-2}, \dots, Y_{t-d}; \mathbf{w}_0) = \frac{\partial f(Y_{t-1}, Y_{t-2}, \dots, Y_{t-d}; \mathbf{w}_0)}{\partial Y_{t-j}} = 0. \quad (7)$$

since f is known and \mathbf{w}_0 can be closely approximated. So, if the j -th lag has no effect on Y we have

$$\theta_j = \mathbb{E} [f_j^2(Y_{t-1}, Y_{t-2}, \dots, Y_{t-d}, \mathbf{w}_0)] = 0 \quad (8)$$

where the square function is used to avoid cancellation effects.

This approach appear to be justified since, with reasonable assumptions on the activation function ψ , a single hidden layer neural network can arbitrarily

closely approximate to $g(\cdot)$ as well its derivatives, up to any given order (provided that they exist), as measured by a proper norm (see [7]). Moreover, a feedforward neural network can achieve an approximation rate that grows linearly in r and thus this class of models is relatively more parsimonious than other non-parametric methods in approximating unknown functions. These two properties make feedforward networks an attractive econometric tool in non-parametric modelling.

3 Testing the Weak form Market Efficiency by Using Neural Networks

In this paper, in contrast to approaches that compare out-of-sample predictions of non-linear models to those generated by the random walk model, we employ the approach proposed in La Rocca and Perna ([11]) for exchange rates and focus on directly testing for unpredictability by conducting tests of significance on models inspired by technical trading rules. In this perspective, the hypothesis that a given set of lags has no effect on Y can be formulated in a multiple testing framework as

$$H_j : \theta_j = 0 \quad vs \quad H'_j : \theta_j > 0, \quad j = 1, 2, \dots, d. \quad (9)$$

and each null H_j in (9) can be tested by using the statistic,

$$\hat{T}_{n,j} = n\hat{\theta}_{n,j} \quad (10)$$

where

$$\hat{\theta}_{n,j} = n^{-1} \sum_{t=1}^n f_j^2(Y_{t-1}, Y_{t-2}, \dots, Y_{t-d}; \hat{\mathbf{w}}_n) \quad (11)$$

and the vector $\hat{\mathbf{w}}_n$ is a consistent estimator of the unknown parameter vector \mathbf{w}_0 . Clearly, large values of the test statistics indicate evidence against the hypothesis H_j .

Thus the problem here is how to decide which hypotheses reject, taking into account the multitude of tests. In such a context, several approaches have been proposed to control the familywise error rate (FWE), defined as the probability of rejecting at least one of the true null hypotheses. The most familiar multiple testing methods for controlling the FWE are the Bonferroni method and the stepwise procedure proposed by Holm ([6]). In any case, both the procedures are conservative since they do not take into account the dependence structure of the individual p -values. These drawbacks can be successfully avoided by using a proposal by Romano and Wolf ([18],[19]), suitable for joint comparison of multiple (possibly misspecified) models.

The algorithm runs as follows. Relabel the hypothesis from H_{r_1} to H_{r_d} in redescending order with respect to the value of the test statistics $\hat{T}_{n,j}$, that is $\hat{T}_{n,r_1} \geq \hat{T}_{n,r_2} \geq \dots \geq \hat{T}_{n,r_d}$. In the first stage, the stepdown procedure tests the joint null hypothesis that all hypotheses H_j are true. This hypothesis is rejected if

\hat{T}_{n,r_1} (the maximum over all the d test statistics) is large, otherwise all hypotheses are accepted. In other words, in the first step the procedure constructs a rectangular joint confidence region for the vector $(\theta_{r_1}, \dots, \theta_{r_d})^T$, with nominal joint coverage probability $1 - \alpha$, of the form $[\hat{T}_{n,r_1} - c_1, \infty) \times \dots \times [\hat{T}_{n,r_d} - c_1, \infty)$. The common value c_1 is chosen to ensure the proper joint (asymptotic) coverage probability. If a particular individual confidence interval $[\hat{T}_{n,r_j} - c_1, \infty)$ does not contain zero, the corresponding null hypothesis H_{r_s} is rejected. Once a hypothesis is rejected, it is removed and the remaining hypotheses are tested by rejecting for large values of the maximum of the remaining test statistics. If the first R_1 re-labelled hypotheses are rejected in the first step, then $d - R_1$ hypotheses remain, corresponding to the labels r_{R_1+1}, \dots, r_d . In the second step, a rectangular joint confidence region for the vector $(\theta_{R_1+1}, \dots, \theta_{r_d})^T$ is constructed with, again, nominal joint coverage probability $1 - \alpha$. The new confidence region is of the form $[\hat{T}_{n,r_{R_1+1}} - c_2, \infty) \times \dots \times [\hat{T}_{n,r_d} - c_2, \infty)$, where the common constant c_2 is chosen to ensure the proper joint (asymptotic) coverage probability. Again, if a particular individual confidence interval $[\hat{T}_{n,r_j} - c_2, \infty)$ does not contain zero, the corresponding null hypothesis H_{r_j} is rejected. The stepwise process is repeated until no further hypotheses are rejected.

Given the probabilistic complexity of the neural network model which makes the use of analytic procedures very difficult, estimation of the quantile of order $1 - \alpha$ is obtained by using resampling techniques. Here we will refer to the maximum entropy bootstrap proposed by Vinod [21,22] e used by Guegan and de Peretti [4] and Heracleous et al. [5].

Given the time series $\{Y_1, Y_2, \dots, Y_n\}$ the resampling algorithm runs as follows (see Vinod and Lopez-de-Lacalle [23]).

1. Sort $\{Y_1, Y_2, \dots, Y_n\}$ by ascending order and let $\{Y_{(1)}, Y_{(2)}, \dots, Y_{(n)}\}$ be the sorted series. Denote $\{i_1, i_2, \dots, i_n\}$ the index series containing the ordering of the original series.
2. Compute intermediate points $z_t = (Y_{(t)} + Y_{(t+1)})/2, t = 1, 2, \dots, n - 1$.
3. Define intervals I_1, I_2, \dots, I_n with equiprobability as well as bounds for the first and last intervals. Then compute the trimmed mean of ΔY_t and compute the lower bound by removing this mean to z_1 and the upper bound by adding this mean to z_n .
4. On each interval compute the desired means defined as $m_1 = 0.75Y_{(1)} + 0.25Y_{(2)}$ for I_1 , $m_n = 0.25Y_{(n-1)} + 0.75Y_{(n)}$ for I_n and $m_j = 0.25Y_{(j-1)} + 0.5Y_{(j)} + 0.25Y_{(j+1)}$ for intermediate values.
5. Draw uniform realizations on $[0, 1]$ and compute the associated quantiles for $\{Y_{(1)}, Y_{(2)}, \dots, Y_{(n)}\}$ by linear interpolation.
6. Adjust the quantiles using z_t to preserve the means, and reorder the adjusted quantiles according to i_t . This returns a bootstrap realization for $\{Y_1, Y_2, \dots, Y_n\}$.
7. Repeat steps from 1 to 6, B of times, with B denoting the number of desired runs.

4 Empirical Evidence from the Italian Stock Market

To test the weak form of market efficiency for the Italian stock market we considered the closing price of the FTSE-MIB index from 1/12/2003 to 23/03/2012. To avoid spurious effects due to local behaviors we also considered a variety of periods of analysis. We first consider the full data set, and we then move to shorter periods in order to examine periods of analysis which may be more homogeneous than the entire sample. From a visual inspection in figure 1, panel (a), it is clear that up to 15/01/2008 (the first 1070 observations) the data series shows relatively low volatility while, in the remaining sample, volatility appear to be much higher. In all the periods considered, as expected, data show strong kurtosis and the Jarque Bera test clearly rejects the hypothesis of normally distributed returns (see Table 1).

Table 1. Descriptive statistics for the FTSE-MIB, daily returns from 1/12/2003 to 23/03/2012. Low volatility period from 1/12/2003 to 15/01/2008, high volatility period from 16/01/2008 to 23/03/2012. P-values are in parenthesis.

Statistics	Whole sample	Low volatility	High volatility
Min.	-0.0860	-0.0384	-0.0860
1st Quartile	-0.0063	-0.0034	-0.0107
Median	0.0007	0.0008	0.0003
Mean	-0.0002	0.0003	-0.0008
3rd Quartile	0.0065	0.0049	0.0111
Max.	0.1087	0.0229	0.1087
Skewness	-0.0884	-0.6619	0.0260
Kurtosys	6.8252	1.8478	3.4026
Jarque Bera test	4159.9 (0.0000)	232.1 (0.0000)	518.0 (0.0000)

In order to test market efficiency hypothesis in its weak form, as described in the previous section, we directly focus on testing for unpredictability by testing the null hypothesis that a given set of past lags has no effect on current returns. More precisely, we consider the following model for the involved stock price at time t , P_t :

$$P_t = g(P_{t-1}, \dots, P_{t-5}) + \epsilon_t \quad (12)$$

and we estimate the unknown function g by using a neural network depending on the considered period (whole: from 1/12/2003 to 23/03/2012, characterized by low volatility: from 1/12/2003 to 15/01/2008, characterized by high volatility: from 16/01/2008 to 23/03/2012). So the problem becomes to test for the significance of P_{t-i} ($i = 1, \dots, 5$) in the model (12).

To avoid the data-snooping problem the testing procedure has been based on the StepM approach in order to take under control the familiwise error rate. The results of the test procedure are reported in figure 1. For the low volatility period we estimated a neural network with $d = 5$ input neurons and 3 neurons

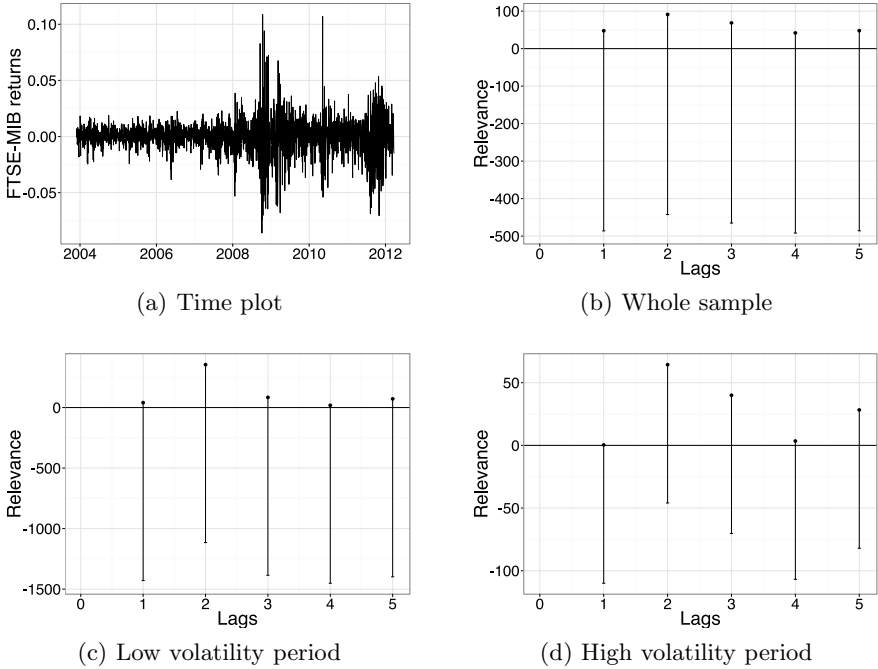


Fig. 1. Tests for weak form market efficiency for FTSE-MIB returns on the whole sample and on low and high volatility periods. All tests have been based on feedforward neural networks with $d = 5$ input neurons and calibrated by resampling with 4999 runs. Hidden layer size is equal to 2 for the whole period and, respectively, equal to 3 and 1 for the low and high volatility period.

in the hidden layer. Clearly all the first five lags appear to be not significant at nominal level $\alpha = 0.05$ and we cannot reject the hypothesis for unpredictability (see panel (c)). Same conclusions can be drawn if we consider the high volatility period (see panel (d)). In this latter case the StepM test has been conducted on the base of neural networks with $d = 5$ input neurons and 1 neuron in the hidden layer. Finally, for the whole period we estimated a neural network with $d = 5$ input neurons and 2 neurons in the hidden layer. Again, at the nominal level of $\alpha = 0.05$ we cannot reject the hypothesis of unpredictability (see panel (b)). In all cases hidden layer size has been fixed by using the Rissanen PSC.

Eventually, we can conclude that when we apply our method to test for predictive power in FTSE-MIB index of the italian stock market we find that it does not appear to contain information that is exploitable for enhanced point prediction.

5 Concluding Remarks

In this paper we have investigated the use of feedforward neural networks for testing market efficiency in its weak form. In contrast to approaches that compare out-of-sample predictions of non-linear models to those generated by the random walk model, we have focused on checking for unpredictability by testing the null hypothesis that a given set of past lags has no effect on current returns. To avoid the data-snooping problem the testing procedure has been based on the StepM approach in order to take under control the familywise error rate. We have applied our method to test for predictive power in FTSE-MIB index of the Italian stock market finding that it does not appear to contain information that is exploitable for enhanced point prediction. Our results suggest that bootstrap-based inference and multiple testing can be a valuable addition to modeling non-linear phenomena with feedforward neural networks.

Acknowledgments. The authors gratefully acknowledge support from the University of Salerno grant program "Sistema di calcolo ad alte prestazioni per l'analisi economica, finanziaria e statistica (High Performance Computing - HPC)- prot. ASSA098434, 2009.

References

1. Antoniou, A., Ergul, N., Holmes, P.: Market Efficiency, Thin Trading and Non-linear Behaviour: Evidence from an Emerging Market. *Eur. Financ. Manag.* 3(2), 175–190 (1997)
2. Campbell, J.Y., Lo, A.W., MacKinlay, A.C.: *The Econometrics of Financial Markets*. Princeton University Press, Princeton (1997)
3. Fama, E.: Market efficiency, long-term returns, and behavioral finance. *J. Financ. Econ.* 49, 283–306 (1998)
4. Guegan, D., de Peretti, P.: An Omnibus Test to Detect Time-Heterogeneity in Time Series. *CES Working papers* 2010.98, 1–22 (2010)
5. Heracleous, M.S., Koutris, A., Spanos, A.: Testing for Nonstationarity Using Maximum Entropy Resampling: A Misspecification Testing Perspective. *Economet. Rev.* 27, 363–384 (2008)
6. Holm, S.: A simple sequentially rejective multiple test procedure. *Scand. J. Stat.* 6, 65–70 (1979)
7. Hornik, K., Stinchcombe, M., Aue, P.: Degree of approximation results for feedforward networks approximating unknown mappings and Degree of approximation results for feedforward networks approximating unknown mappings and their derivatives. *Neural Comput.* 6, 1262–1275 (1994)
8. Kuan, C., Liu, T.: Forecasting exchange rates using feedforward networks and recurrent neural networks. *J. Appl. Econom.* 10, 347–364 (1995)
9. La Rocca, M., Perna, C.: Variable selection in neural network regression models with dependent data: a subsampling approach. *Comput. Stat. Data An.* 48, 415–429 (2005a)
10. La Rocca, M., Perna, C.: Neural Network Modeling by Subsampling. In: Cabestany, J., Prieto, A., Sandoval, F. (eds.) *IWANN 2005*. LNCS, vol. 3512, pp. 200–207. Springer, Heidelberg (2005)

11. La Rocca, M., Perna, C.: Neural Network Modelling with Applications to Euro Exchange Rates. In: Kontoghiorghes, E.J., Rustem, B., Winker, P. (eds.) *Computational Methods in Financial Engineering*, pp. 163–189 (2009)
12. Yen, G., Lee, C.: Efficient market hypothesis: past, present and future. *Rev. Pacific Basin Financ. Markets and Policies* 11(2), 305–329 (2008)
13. Leshno, M., Lin, V., Pinkus, A., Schocken, S.: Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural Networks* 6, 861–867 (1993)
14. Lim, K.P.: Ranking of efficiency for stock markets: a nonlinear perspective. *Physica A* 376, 445–454 (2007)
15. Lo, A.W., MacKinlay, A.C.: Stock market prices do not follow random walks: evidence from a simple specification test. *Rev. Financ. Stud.* 1(1), 41–66 (1988)
16. Malkiel, B.G.: The efficient market hypothesis and its critics. *J. Econ. Perspect.* 17(1), 59–82 (2003)
17. Park, C., Irwin, S.H.: What do we know about the profitability of technical analysis? *J. of Econ. Surv.* 21(4), 786–826 (2007)
18. Romano, J.P., Wolf, M.: Exact and approximate stepdown methods for multiple hypothesis testing. *Journal of the American Statistical Association* 100, 94–108 (2005a)
19. Romano, J.P., Wolf, M.: Stepwise multiple testing as Stepwise multiple testing as formalized data snooping. *Econometrica* 73, 1237–1282 (2005b)
20. Shiller, R.J.: From efficient markets theory to behavioral finance. *J. Econ. Sur.* 17, 83–104 (2003)
21. Vinod, H.D.: Ranking Mutual Funds Using Unconventional Utility Theory and Stochastic Dominance. *J. of Empir. Financ.* 11, 353–377 (2004)
22. Vinod, H.D.: Maximum Entropy Ensembles for Time Series Inference in Economics. *J. of Asian Econ.* 17, 955–978 (2006)
23. Vinod, H.D., Lopez-de-Lacalle, J.: Maximum Entropy Bootstrap for Time Series: The meboot R Package. *J. Stat. Softw.* 29, 1–19 (2009)
24. White, H., Racine, J.: Statistical Inference, The Bootstrap, and Neural Network Modeling with Application to Foreign Exchange Rates. *IEEE T. Neural Networ.* 12, 657–673 (2001)

Part IV

Special Session on “Smart Grids: New Frontiers and Challenges”

Real Time Techniques and Architectures for Maximizing the Power Produced by a Photovoltaic Array

Giovanni Petrone¹, Francisco Jose Sánchez Pacheco², and Giovanni Spagnuolo¹

¹ University of Salerno

Dipartimento di Ingegneria Elettronica ed Ingegneria Informatica (D.I.E.I.I.),
Via Ponte Don Melillo 84084 Fisciano - Salerno - Italy
gspagnuolo@unisa.it, gpetrone@unisa.it

² Universidad de Málaga

Escuela Politécnica Superior (despacho 3050) Dept. de Tecnología Electrónica
C/ Dr. Ortiz Ramos (campus Teatinos) 29071 Málaga - Spain
fsanchezp@uma.es

Abstract. The inherent low conversion efficiency, from solar to electrical energy, of the photovoltaic cells makes the use of techniques and architectures aimed at maximizing the electrical power a photovoltaic array is able to produce at any weather condition mandatory. In order to understand what are the challenging problems cropping up in some modern applications, an overview of the main techniques for photovoltaic arrays modeling is given first. Afterwards, the control strategies for the maximum power point tracking used in commercial products dedicated to photovoltaic strings and modules are compared and their advantages and drawbacks are put into evidence, with a special emphasis on their efficiency. Some methods presented in literature and based on the use of artificial neural networks are compared with more classical ones. Finally, a brief overview of other applications of artificial neural networks to photovoltaic-related problems is also given.

1 Introduction

Photovoltaic (PV) technology in its modern era is referred to be dated by 1954 [1]. It was described that a p-n semiconductor junction under the effect of sun light could generate electricity. A cheap widespread source of energy is since then used to increasingly satisfy human needs of electric power [2]. Furthermore, this technology is a key issue to satisfy nowadays requirements to reduce CO_2 emission rates due to the use of fossil fuels in electricity production. PV devices can be tailored to supply power electricity in different scale ranges, from single cells for energy harvesting purposes aimed at supplying remote sensors, to small module applications involving a few number of cells, up to the huge scale of grid connected high power plants requiring a large number of parallel connected strings made of series connected panels. Up to now, the largest part of the research efforts and of the commercial products have been devoting to large power applications involving grid connected power plants. In such cases, all the PV panels are of the same type/model and they have the same orientation towards the sun. The plant

is installed in a flat site without any obstacle in the neighborhood, so that the maximum possible electrical power production is ensured from sunrise to sunset.

More recently, the scientific community as well as many companies operating in the field of PV systems have shown a large interest in applications dedicated to small power loads. These range from harvesting systems dedicated to remote sensors and communication systems up to battery chargers for sustainable mobility applications. In such cases, the number of PV cells electrically connected in series is quite low, so that the main issue is to raise up the voltage level up to that one required by the load. Nonetheless, very often in these applications the cells do not have the same orientation towards the sun, also changing dynamically as in applications to mobility, and some of them can be subjected to a time varying shadowing effect. The latter also occurs in Building Integrated PV (BIPV) applications, which are of interest to the aim of making buildings as much autonomous as possible from the point of view of the energy production.

In order to describe the behavior of a PV generator, regardless of its size, thus a cell, a module, a panel, a string or a large field, working in such a non conventional conditions some dedicated models and numerical methods are needed. This task is of fundamental importance in order to understand the mechanisms that affect the power production of a so-called *mismatched* PV generator and for preventing those operating conditions that might lead to a permanent damage of some cells. Furthermore, comprehensive models are the basis for developing strategies that allow a proper control of the PV generator also in presence of mismatched conditions, especially ensuring that the maximum available power is harvested at any time. The latter is not a trivial task: as demonstrated by the huge amount of papers that can be found in literature, Maximum Power Point Tracking (MPPT) is a challenging problem in any PV system, because of the need of extracting the maximum electrical power from the PV generator without having any knowledge about the type/model of cells and without measuring neither the irradiation level nor the temperature at which the cells work. The MPPT operation becomes much more involved if the PV generator works in mismatched conditions, because the MPPT algorithm must be able to distinguish the absolute maximum power point from the relative ones in a multi-modal characteristic.

In this paper an overview of the methods presented in literature and used in commercial products for affording the problems mentioned above is given. First of all, some of the main techniques used for modeling and simulating PV generators working both in uniform and in mismatched conditions are compared. Such approaches are usually based on a proper description of the PV non linearities and on an effective solution of the non linear system of equations that follows. A special attention is devoted to the PV modeling methods that are based on Artificial Neural Networks (ANNs). In fact, ANNs seem to be an effective tool for describing a complex system, which is strongly non linear and depending on a large number of time varying parameters, as in the case of PV generators is. In fact, some parameters like the irradiation level and the operating temperature, but also some others related to the semiconducting material the cells are made of, may assume unpredictable values which are also subjected to some drifts. The ANNs profit from their ability of learning from a training set, which might be made of some experimental data taken from the real PV generator working in different conditions. The second part of this paper is dedicated to the the real time techniques and

to the architectures that are commonly used for maximizing the power produced by a PV array. Advantages and drawbacks are put into evidence, especially by referring to the solutions that are used in the largest part of products available on the market and dedicated to large scale systems as well as to single PV modules. The power electronics solutions and system architectures used for grid connected as well as stand alone systems are overviewed. The role of the MPPT efficiency in the design of a PV power processing system will be emphasized and the factors, especially the effect of noises, affecting it are discussed. Some solutions for achieving a high MPPT efficiency in noisy conditions are mentioned and compared. The role of ANNs in this field is analyzed by referring to some techniques introduced by recent papers appeared in literature.

In a final section, some applications of ANNs to PV-related problems are overviewed and the most recent literature on these topics is referenced. In fact, ANNs can be beneficial in the weather forecast, thus in the prediction of the irradiation level and temperature at which the PV generator will work, in the estimation of the power production and efficiency of PV plants as well as in their sizing, both in stand-alone and in grid-connected applications. Conclusions end the paper.

2 Photovoltaic Source Modeling

A PV generator is made of a parallel connection of strings obtained by connecting in series a number of panels. The number of panels in a string is dictated by the input voltage requirements of the power processing systems, while the number of strings to put in parallel depends on the required power level. Each PV panel is made of a few number of PV modules electrically connected in series and assembled in the same frame. Commercially available PV panels consist in usually two or three modules, each one of them formed by about twenty cells equipped with an anti-parallel diode, as shown in Fig. 1.




Fig. 1. PV module made of a number of series-connected cells and the anti-parallel bypass diode

If the PV generator is made of cells that are exactly equal, in terms of characteristics of the semiconductor material they are made of, and working in exactly the same conditions, especially temperature and irradiation level, its current vs. voltage (I-V) curve is merely a scaled up version of the same curve referred to a single PV cell. The voltage value is scaled up by the number of cells connected in series and the current value is multiplied by the number of strings electrically connected in parallel. This leads to the curves shown in Figg. 2 and 3. The curves are in normalized units and clearly put into evidence the non linear and time dependent nature of the PV generator. The peak

in the power vs. voltage (P-V) curve must be tracked in order to ensure that the PV field delivers the maximum power to the load. The peak position changes due to the irradiance level G the field is subjected to, but it is also affected by the ambient temperature T_a . The former mainly affects the current level, but it is related to the voltage by a logarithmic law. On the contrary, the voltage values are mainly influenced by the working temperature of the cells, because at a higher temperature the open circuit voltage decreases.




Fig. 2. Current vs. voltage characteristic of a PV panel: effect of irradiation G

The behavior of the PV generator, at whichever granularity level (cell, module, panel, string and field), is described by the equivalent circuit shown in Fig. 4 and by the corresponding equation (1).

$$I = I_{ph} - I_{sat} \cdot \left(e^{\frac{V+I \cdot R_s}{\eta V_t}} - 1 \right) - \frac{V + I \cdot R_s}{R_p} \tag{1}$$

In (1) V_t is the thermal voltage, I_{sat} is the saturation current and I_{ph} is the photoinduced current, depending on the type of cell used and on the irradiance level and on the temperature [3]. The resistances R_s and R_p represent the various loss mechanism taking place in the cell and in the whole PV array. As stated above, equation (1) allows to model any type of PV generator working in uniform conditions, provided that the values assumed by the parameters and the variables appearing in it, that are supposed to be known for a single cell composing the generator, are scaled up properly.

Unfortunately, such a simple, although non linear, model is too rough to be used for describing the mismatched operation of a PV array, so that much more sophisticated tools have been presented in literature.




Fig. 3. Power vs. voltage characteristic of a PV panel: effect of irradiation G




Fig. 4. Single diode model accounting for ohmic losses

2.1 Classical Modeling Approaches

A first step in obtaining a compact model of a mismatched PV array, which also have the advantage of ensuring a fast numerical simulation, is the manipulation of the equation (1) that is not able to give explicitly neither the voltage nor the current values as a function of the other electrical variable. In literature, the Lambert W function has been used fruitfully in order to achieve this result. In [4] many details and useful references about the Lambert W function can be found, but the main thing to know is that it is the solution of the equation:

$$f(x) = x \cdot e^x \tag{2}$$

which evidently occurs in (1), and that can be calculated by means of a suitable series expansion [4]. The Lambert W function allows to obtain a compact expression giving the value of the current at whatever voltage value, with a clear dependence on all the parameters depending on the semiconductor material used to realize the cells, as well as on irradiance and temperature:

$$I = \frac{R_p \cdot (I_{ph} + I_{sat}) - V}{R_s + R_p} - \frac{\eta \cdot V_t}{R_s} \cdot \text{LambertW}(\theta) \quad (3)$$

where:

$$\theta = \frac{(R_p // R_s) \cdot I_{sat} \cdot e^{\frac{R_p \cdot R_s \cdot (I_{ph} + I_{sat}) + R_p \cdot V}{\eta \cdot V_t \cdot (R_p + R_s)}}}{\eta \cdot V_t} \quad (4)$$

As shown in [5], in the same way the explicit expression of the current flowing into a PV module, including a string of cells equipped with the anti-parallel bypass diode, takes the following expression, wherein θ assumes the form already given in (4):

$$I = I_{ph} - I_{sat} \cdot \left(e^{\frac{V + I \cdot R_s}{\eta V_t}} - 1 \right) - \frac{V + I \cdot R_s}{R_p} + I_{sat, dby} \cdot \left(e^{-\frac{V}{\eta_{dby} V_t, dby}} - 1 \right) \quad (5)$$

In this expression, the part with the subscript *dby* refers to the bypass diode. This formula allows an effective simulation of one PV panel, consisting of a two or three modules in series, operating with different irradiation, temperature or exhibiting different values of the material parameters for each module. In this way, the level at which a PV can be simulated is reduced at the module, not panel, level, thus is reduced at one half or one third of a panel. Such scaling factors, the half or one third, are cited as examples because of the real numbers used in commercially available panels. For instance, the Suntech STP280 polycrystalline panel has 72 cells, organized into six rows of twelve cells, with three bypass diodes, each one connected in anti-parallel to twenty four cells. The approach proposed in [5] profits from (5) and from the properties of the Lambert W function. In fact, it allows to achieve an explicit expression also for the differential conductance, that is the derivative of the PV module current with respect to its voltage. In this way, the non linear system of equations, that allows to calculate the operating points of all the PV modules in a string at whatever value of the string voltage, can be solved effectively by means of any classical algorithm, e.g. the Newton-Raphson method. The method allows to reconstruct the mismatched I-V curve at the desired level of accuracy, but it can be also useful for the real time simulation of the string in environments like PSIM or PSPICE.

An approach that concentrates the calculation effort across the so-called *inflection points* is introduced in [6]. It allows to calculate in an effective way the voltage values where, starting from the low voltages at which only the module receiving the highest irradiation produces power, the bypass diodes end their conduction and the modules with a lower short circuit current start giving their contribution. According to the authors of [6], the method is able to give an approximate version of the mismatched curve quite quickly, but may lack in the accurate reconstruction of the whole curve.

In mismatched conditions, the PV curve of the array is multimodal, namely it does not look like those ones shown in Fig.3, characterized by a single maximum power point, but it shows multiple peaks. This is due to the bypass diode operation, so that some modules are bypassed and absorb some electrical power. This determines a power loss, which is fortunately low thanks to the low activation voltage of the anti-parallel bypass diode. Such a small power loss is the price to pay in order to avoid that the modules that receive a higher irradiation are penalized because of the draft series connection with those ones producing less power. Fig.5 shows an example of PV curves in uniform

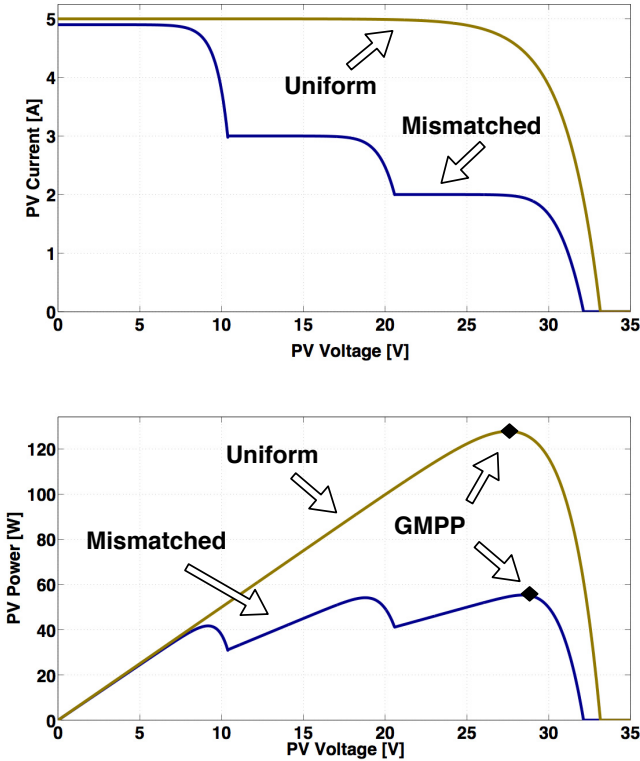


Fig. 5. Photovoltaic characteristics in homogeneous and partial shading conditions. a) Current vs. Voltage curve, b) Power vs. Voltage curve.

and mismatched conditions, by putting into evidence the existence of a Global Maximum Power Point (GMPP). The detrimental effect of the local maxima has been well illustrated in [7] where some experimental studies have been carried out on different commercial products.

The methods mentioned above assume that the smaller entity to be modeled in the PV array is the module, thus considering that all the cells belonging to the module work in the same conditions or, otherwise, that an average behavior of the module can be determined if some mismatching event occurs at a sub-module level. Unfortunately, this granularity level is not satisfying if *hot spot* phenomena must be modeled and if the real behavior of a PV array in which a few number of cells is subject to shadowing must be reproduced. In fact, if a single cell in a module receives an irradiation that is significantly lower than the one at which the others work, its operating voltage can be placed at its own breakdown value, which is deeply negative, with a positive value of the current. Thus, the module works at a low voltage, with the positive one of the fully irradiated cells that is compensated by the negative one of the shadowed cell. In such conditions, the bypass diode does not enter into conduction and the shadowed cell dis-

sipates a significant power that can damage it irreversibly. Such a mechanism, that is explained in [3], is analyzed in some papers (e.g. [8]) and requires a heavy simulation model, which takes into account thermal effects and describes each cell, instead of each module. It is evident that such approach might be non adequate if it has to be used for processing a huge amount of data, i.e. for evaluating the annual PV energy yield. Moreover, by considering the strongly nonlinearity of PV systems, in some cases, it is very difficult or almost impossible to determine an analytical or numerical model for describing the whole PV generator, especially if the extreme variability of the climatic conditions has to be taken into account. In such a case, behavioral models are more suitable to characterize the whole system because they are not focused on the identification of the exact value of a specific variable but rather they have the objective to estimate a value with a given level of confidence.

2.2 Artificial Neural Network Based PV Modeling

In recent literature some applications of the ANNs to PV systems modeling have been presented. All of them put into evidence the advantage given by ANNs of being independent of the complexity of the relationship between the PV array current and its voltage as well as the involved and interleaved dependency on the physical and weather parameters. ANNs process the information in parallel through many simple elements: the neurons. All these neurons are interconnected and every connection has a given weight. Finally, each neuron supplies an output through an activation function. This structure fits with the need, in PV systems modeling, of describing the system on the basis of the relationship between the given inputs (e.g. irradiance, temperature, voltage) and desired outputs (e.g. current). If compared with the mathematical model described above, the ANN does not require the knowledge of internal system parameters, involves less computational effort and offers a compact solution for such a multiple-variable problem.

In [9], a neural network based PV panel model that uses (I) describing the equivalent circuit shown in Fig 4 is considered. The ANN training is done by taking five operating points per panel at some given couples of irradiance and temperature values. On this basis, the ANN is able to determine the values of the five parameters in (I), that are $\{I_{ph}, I_{sat}, \eta, R_s, R_p\}$, at whichever irradiance and temperature values, so that the whole curve is available in any condition. The authors show a very good accuracy of the curve reconstruction at a low irradiation, just where the classical approaches based on the parameters identification in Standard Test Conditions (STC) show the larger inaccuracy with respect to experiments. Unfortunately, such a method does not allow to keep into account mismatching conditions at a panel level.

In [10] a Multi Layer Perceptron (MLP) was trained by using experimental measurements and it shows good results, especially if the reconstruction of the I-V curve is required at low irradiation levels.

In [11], it was pointed out that the approach based on a MLP ANN may be characterized by a slow training process, which can also remain trapped into local minima. Such limitations are avoided by using a radial basis function network, which can be designed by affording a sort of curve fitting problem in a multi dimensional space. This training task is afforded by means of an orthogonal least squares method. In [12], the

optimization of the hidden layers is performed by using a genetic algorithm. In this way, the number of the radial basis functions is not too low, with a poor function approximation performance, and not too high, with an over fitting of the input data taken from experiments on the PV array. In both [11] and [12] the ANN takes the radiation, the ambient temperature and the PV array voltage as inputs and calculates the array current as output. The procedures are good candidates for predicting the path described by the maximum power point along the day, according to the current weather conditions.

In a recent paper [13], the importance of taking into account the spectral distribution of the incident light, especially at a low irradiance level and for cell technologies having a spectral response narrower than mono crystalline silicon, has been put into evidence. In [13], the authors have included the information about the spectral distribution of the light as a further input of an ANN. They have also implemented a non-random selection of the data used as training set, with an improved performance of the network trained with the spectral information. The price to pay for this accurate modeling is in the amount of data representing the spectral information, which have to be given to the ANN as an input, and in the pre-processing of the training set, which is a time consuming procedure. In [13], the reader can also find an up-to-date glance at the most recent papers published in the field.

3 Maximum Power Point Tracking (MPPT)

Due to the time varying environmental condition such as temperature and solar irradiation, the P-V characteristic exhibits a maximum power point (MPP) which is strongly variable in P-V plane. This fact is quite evident by looking at Fig.3 where the variation of the MPP due to the irradiance excursion is documented. Unfortunately, the joined effect of the irradiance and temperature variations leads to a change of the voltage at which the MPP occurs along the day in a wide region, so that the locus of the MPP's is not a line or a curve, but an area. This large variation makes the straightforward connection of the PV array to a constant voltage port wrong. In fact, the PV generator can be used for recharging a battery or can be plugged at a DC bus that is the input port of a DC/AC converter, the latter feeding the grid or an AC load. At a constant voltage value, the PV field does not deliver the maximum power, unless the DC voltage at which it is forced to work is exactly the MPP voltage. A coarse matching between the DC voltage and the PV array might force the latter, at some values of the irradiance and of the ambient temperature, to work at a voltage level that is higher than the open circuit voltage, with a zero power produced.

In order to ensure the optimal utilization of PV arrays, a MPP Tracking (MPPT) is realized by means of a suitably controlled power converter, which can be a DC/DC converter or an inverter. This is helpful for adapting the PV optimal operating point to the load or grid requirements [14].

The MPPT operation must be able to ensure that the PV operating point is as much close as possible to the MPP, both in steady state weather conditions and when irradiance transients, which are faster than temperature ones, occur. As a consequence, a MPPT efficiency can be defined as the ratio between the energy extracted from the PV array and the energy that the same array would have been able to produce by always

working at its MPP. In practice, this efficiency is less than one because the MPPT algorithm is not able to stay in the MPP in steady state conditions and it is not prompt enough for tracking the MPP when the irradiation level changes suddenly (e.g. in presence of clouds moving at high speed due to the wind or in PV applications to sustainable mobility) [15]. Furthermore, noise affecting PV current and voltage worsen the tracking performances of the MPPT algorithm [16]. Example of noise are the switching ripple introduced by the converter that operates the MPPT and the quantization effect introduced by the use of A/D converters and digital controllers implementing the MPPT algorithm. A further noise source is the DC/AC stage in single phase applications: the low frequency voltage oscillation at the DC bus, at a frequency that is the double of the AC voltage one, back propagates up to the PV sources, thus degrading the MPPT performances.

A classical double stage architecture for AC PV applications is shown in Fig. 6: the bulk capacitance, placed between the two conversion stages, is used to manage the fluctuating AC power and consequently a voltage ripple with a frequency that is the double of the grid frequency appears at its terminals.

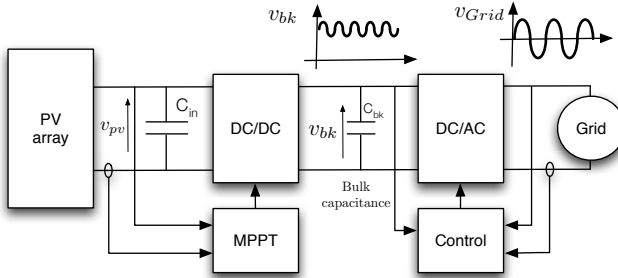


Fig. 6. Double stage grid-connected inverter

The amplitude of the voltage ripple is given by:

$$\Delta V_b = \frac{P_{PV}}{2 \cdot \omega_{grid} \cdot C_{bk} \cdot V_{bk}} \tag{6}$$

where P_{PV} is the DC power extracted by the PV field.

The oscillation affecting the voltage of the bulk capacitor has detrimental effects on both the DC and the AC part of the power processing system, so that a large electrolytic capacitor is almost always used at the DC link. In fact, as (6) reveals, the larger the bulk capacitance the smaller the voltage oscillation. Unfortunately, this component is a weak point of the conversion chain, because of the effects that the operation temperature has on the electrolytic capacitors lifetime. A significant effort is done by PV inverters manufacturers in order to keep the working temperature of the bulk capacitor as close as possible to that one at which the capacitor manufacturer has tested the component for some thousands of hours, so that the Mean Time Between Failures (MTBF) is increased. A reduced value of the bulk capacitance might allow to use film capacitors instead of

electrolytic ones, if suitable control techniques or different topologies must be adopted in order to reduce such an effect. In [16], the way in which the control network of the DC/DC stage can be designed for achieving such objective has been explained.

In the next subsection the main MPPT techniques are overviewed and some control methods aimed at reducing the low frequency disturbances in AC applications are compared.

3.1 Classical MPPT Approaches

The switching converter would be able to ensure the maximization of the power produced by the PV generator provided that the parameter, or the parameters, which allow to change its input voltage/current levels are suitably controlled. In the largest part of PV power processing systems, e.g. that one shown in Fig. 6 the circuit that performs the MPPT operation is a DC/DC converter and the control parameter is the duty cycle, as shown in Fig. 7




Fig. 7. Connection scheme of a dc/dc converter dedicated to the dynamical optimization of a PV generator

The two main techniques that are used in commercial systems for performing the MPPT function are the Perturb and Observe (P&O) and the Incremental Conductance (IC) methods. They are perturbative approaches that change the PV array voltage repeatedly until the PV power is maximized. PV current and voltage measurements are needed in order to determine the PV power produced by the array at each value of the PV voltage settled by the MPPT algorithm through the switching converter. The perturbed variable is usually the converter duty cycle, but in literature there are evidences of the fact that the reference signal in a closed loop switching converter is the best variable to control for achieving the MPPT [16]. This choice allows to improve the steady state and the transient performances of the MPPT algorithm. In fact the converter's dynamics can be improved by a proper design of the feedback compensator, so that the time between two consecutive perturbations can be shortened. Additionally, the perturbation amplitude can be reduced as well, because the closed loop transfer function of the converter can be designed in order to have a very low gain of the output-to-input closed loop transfer function in the noise bandwidth.

The superiority, claimed in some papers, of the IC method with respect to the P&O one in terms of MPPT efficiency [17] does not have a practical confirmation, essentially because of the noise that in practice affects both the PV array voltage and current. The perturbed variable is in any case the PV voltage because of a logarithmic dependence from the irradiation level.

As discussed in the recent literature, the proportionality between the irradiation and the PV current worsens the tracking capability of any current-based MPPT technique, in particular when a significant irradiation drop occurs. This problem has been addressed by using an innovative control technique that matches the sliding mode control with the P&O MPPT technique [18]. The approach guarantees a high tracking promptness, an intrinsic independence of the MPPT technique from the PV array parameters and an inherent rejection of the noises propagating from the output towards the input of the switching converter.

As for the observed variable, whose value has to be maximized by the MPPT algorithm, in the basic systems, that are the majority, it is the PV power, obtained as the product of the digitalized values of the PV current and voltage. Nevertheless, the non linear relation between the PV voltage and the efficiency of the converter can lead to the paradox of having the maximum power produced by the PV array that is processed by a switching converter that does not work at its highest efficiency. The best tradeoff is achieved by tracking the maximum of the power at the converter output. This solution may allow to save the voltage sensing and may require a current sensor only [19] [20].

Many variants of IC and P&O algorithms have been proposed in literature. Some of them adopt a variable perturbation amplitude in order to achieve a high MPPT performance in steady state conditions and a good promptness in presence of a varying irradiance [21]. The algorithms are almost always implemented in a digital way, in order to profit from both the flexibility and the IP protection ensured by modern digital devices at a reasonable cost. Nevertheless, some approaches that can be implemented by means of analog circuitry only are also presented in literature [22] [23] [24]. They are based on the evaluation of the effect that a small PV voltage oscillation has on the PV power: if the operating point is on the left side of the MPP, the forced PV voltage oscillation and the consequent PV power oscillation are in phase. On the contrary, if the operating point is on the right side of the MPP, the two oscillations are in phase opposition. At the MPP, the PV power oscillation falls below a given threshold and the objective is achieved. This method, often referred to as *ripple correlation control* or *extremum seeking control*, improves the steady state MPPT efficiency because the repeated climbing across the MPP required by P&O and IC is avoided. An excellent promptness during irradiance transient is also obtained, but the price to pay is the reduced flexibility if an analog implementation of such methods is adopted.

The MPPT operation becomes more complicated when the PV array works in mismatching conditions, because the whole PV curve is multimodal, namely characterized by multiple peaks, thus the maximum global peak or Global Maximum Power Point (GMPP) is not tracked easily. Conventional MPPT algorithms, which are based on an hill climbing approach, are not able to distinguish the GMPP from the local peaks, so that they are trapped into a voltage range across a power maximum, without knowing if it is the GMPP. In some cases, e.g. depending on the shading pattern, the GMPP can

occur at a voltage that is out of the working range of the inverter, so that the best peak falling in that range must be tracked. In literature some approaches to the tracking of the GMPP are proposed, but they need quite complicated algorithms and require some preliminary assumptions. As a consequence, in practice, the GMPP is detected by means of a periodical, and energy costly, sweep of the P-V curve performed by the MPPT algorithm.

Such a puzzling problem can be avoided by using Distributed MPPT (DMPPT) architectures [25], which employ a small power switching converter dedicated to each PV panel. In this case, each panel is controlled independently from the other ones and a multiple peak P-V curve can appear at a panel level only. Should the panel level MPPT not be able to track the GMPP, but it remains trapped in one of the other few MPP's, the power loss remains limited to the mismatched panel. Two solutions have been proposed in literature and have been becoming products available on the market [26]. The first one uses a DC/DC converter for each panel, the output terminals of the converters being connected in series and plugged to the inverter's input port. This architecture does not require unusual values of the voltage conversion ratio, but it gives rise to significant control problems in presence of a large different in the operating conditions of panels in the same string. The other possible solution employs DC/AC converters that inject the PV power produced by each panel straightforwardly into the grid. The maximum modularity of this architecture is counterbalanced by the need of having a large voltage conversion ratio, with a conversion efficiency that is lower than 95%.

In literature, some hybrid DMPPT solutions, with a distributed power processing and a centralized MPPT algorithm have been presented [20]. They maximize the total output power, thus accounting for the non constant conversion efficiency of the DC/DC converters, and save a number of current sensors with respect to the basic solution.

4 ANN Based MPPT Approaches

Although conventional MPPT algorithms operate very well under uniform irradiation conditions, and in fact they are widely used in commercial products, several works recently appeared in literature have been focused on the use of artificial intelligence techniques for tracking the maximum power point. Truly speaking, many of them, also appeared recently, are focused on the MPPT in wind energy systems [27], instead than on PV ones. To this aim, ANNs are also combined with other soft computing methods: for instance, Evolutionary Algorithms (EA) or the Particle Swarm Optimization (PSO) have been fruitfully applied in [28,43] in conjunction with ANNs.

In [29] a comparison between a Fuzzy Logic Control (FLC) algorithm and an ANN-based approach has been presented. By assuming that the PV field works in uniform conditions, it has been demonstrated that the FLC controller is able to generate up to 99% of the actual maximum power while the ANN controller can reach the 92% value only. However, although the FLC-MPPT tracking is more effective than the ANN method, the former requires extensive processes which include fuzzification, rule base storage, inference mechanism and de-fuzzification operations. Consequently, a compromise has to be made between the tracking speed and the computational cost.

The joined use of ANNs and fuzzy logic is also effective in the MPPT under mismatched operating conditions, where the GMPP position depends on shading patterns (see fig 5) and the voltage at which it occurs may change within a large voltage range. In [30] a ANN is trained by using many different partially shaded conditions to determine the corresponding GMPP voltage of the whole array. The ANN output is the voltage reference for the FLC used to generate the required control signal for the power converter; the proposed configuration is shown in fig 8.

The input signals for the ANN are the irradiance level (G) and the cell temperature (T_c). The neural network predicts the global MPP voltage (V_{dc}^*) and power (P_{dc}^*) and those values are compared with the actual voltage and power expressed in V_{dc} and P_{dc} . The predicted output voltage (V_{dc}^*) is used as a reference signal for the FLC voltage based MPPT controller. The proposed method has been experimentally validated, so that the reliability and the performances of such an MPPT algorithm have been demonstrated to be superior to the conventional P&O method in mismatched conditions.




Fig. 8. ANN-FLC mppt architecture

In [31] an approach that uses a ANN and a FLC is also introduced. In this case, the neural network is used for tuning automatically the membership functions of a FLC that is employed to track the MPP. The method has been compared in simulation with the standard P&O technique and with a manually tuned FLC. Results presented in the paper show that the proposed optimized FLC provides a fast and accurate tracking of the PV maximum power point under various operating conditions, including mismatched ones. Fig 9 shows how in presence of a mismatched condition the operating point moves towards the MPP's. The proposed approach is the only one that is able to drive the operating point towards the GMPP.

In [32] a full-bridge inverter is chosen as an active low-frequency ripple-control circuit (ALFRCC). Fig 10 shows the scheme of the proposed architecture: the additional stage is used to remove the DC-bulk capacitor and operates with an AC-bulk capacitor (C_{bo}). In this configuration the full-bridge inverter works for injecting a suitable




Fig. 9. MPP under partial shading $G_1 = 200W/m^2$ and $G_2 = 1000W/m^2$

compensation current into the high-voltage bus. Although such approach is not new, its novelty is in the adoption of a combination of an adaptive linear network and a sliding mode control for generating the driving command for the ALFRCC that is able to mitigate the low frequency ripple at the PV terminals. The ALFRCC effectiveness has been verified by numerical simulations and experimental results. Its superiority is indicated in comparison with a conventional high-pass filter and a proportionalintegral controller.

The use of the neural network for controlling the active filters is also proposed in [33] and [34] where it is employed to improve the power factor and to reduce the line current harmonics.

5 Further ANN-Based Techniques for PV Power Production Maximization

In [35] a wide overview of the possible contributions ANN methods and techniques can give in maximizing the power produced by a PV plant is given.

In the field of stand alone PV systems and of residential applications, ANNs are helpful because of their capability of foreseeing the PV energy production, which allows to manage the energy flows more effectively. The objective is not limited to the maximization of the PV energy production but extends to the matching with the local energy consumption in order to reduce the contribution coming from the AC grid. The encouragement, especially in residential applications, of the self-consumption is a key point of the future European Union strategies for a better use of the energy [36]. In [37] a control system based on an Active Demand-Side Management (ADSM) for PV residential application has been proposed. The ADSM is a distributed control system made of several ANNs dedicated to the different appliances in the house, so that appliances self-organize their activities and a coordinator corrects their actions in order to enhance self-consumption. The system acts in an almost transparent way to the user and it takes in charge the schedule of the household tasks for the next day on the basis of the foresee of the PV power production, thus leading to an increased energy efficiency.




Fig. 10. Double stage grid-connected inverter with active filter at the bulk interface

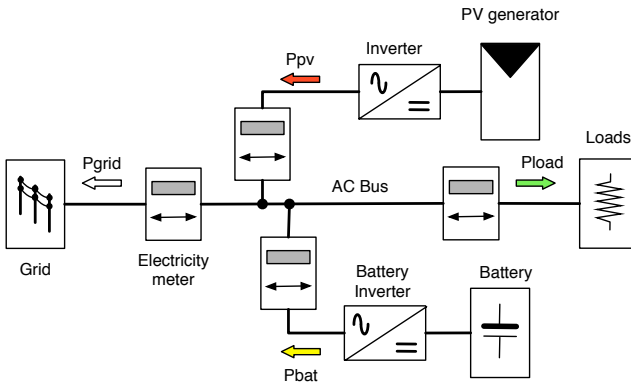


Fig. 11. Residential ADSM architecture

Fig. 11 shows the proposed architecture, including the PV array, a storage element, a connection to the grid and a home automation system.

In [38], an ANN-based control of a stand alone system including a PV array, a diesel generator and a wind turbine has been proposed. The MPPT of the wind generator is performed by means of a ANN that regulates the blades pitch angle. A Radial Basis Function Network (RBFN) performs the PV MPPT function, Simulation results show that an efficient power sharing technique among energy sources is obtained and the voltages and power can be well controlled in presence of environmental variations.

In [39] the MPPT task employs a RBFN with a back-propagation network for predicting the effects of passing clouds on a stand alone PV system equipped with a storage unit. By using the irradiance as input signal, the network models the effects that the random cloud movement has on the electrical variables of the system, thus reducing the problems related to the overload/underload of the power lines due to the PV power variation in the short periods of time when the cloud movement affects the PV plant.

Both in large and in small PV plants, as well as in both grid connected and stand alone systems, a key factor in energy efficiency is the correct sizing of the generator.

In [35] a stand alone system operating under variable climatic conditions is modeled and simulated by means of a ANN. Electrical and weather parameters recorded during several years of operation of the plant have been used for the training and testing of the developed models. In [40], the concept of Loss of Load Probability (LOLP), related to the ability of the system to satisfy load requirements, is used in conjunction with a MLP neural network.

ANNs are also widely used in PV plant productivity estimation. In [41] the historical data concerning irradiance and temperature are used for a MLP ANN. In [42] an ANN trained by a Genetic Swarm Optimization (GSO) algorithm is used to foresee the production of a PV plant.

6 Conclusions

In this paper an overview of the techniques used for modeling, controlling and designing a photovoltaic system has been given. The non linearity of the model describing a photovoltaic generator, especially when it works in mismatched conditions, has been put into evidence. Some analytical models and a number of approaches based on the adoption of artificial neural networks have been mentioned and compared. The maximum power point tracking problem has been described and, also in this case, a glance to advantages and drawbacks of some techniques that are used in commercial systems and of novel ones based on artificial neural networks has been given. Finally, some attractive applications of the artificial neural networks, often used in conjunction with soft computing techniques, have been overviewed. The main ones are in the design of hybrid systems including generators employing renewable energies and backup devices.

Acknowledgements. This paper has been supported by PRIN 2008 and FARB funds.

References

1. Chapin, D., Fuller, C., Pearson, G.J.: A new silicon p-n junction photocell for converting solar radiation into electrical power. *Applied Physics* 25, 676–677 (1954)
2. Luque, A., Hegedus, S.: *Handbook of photovoltaic science and engineering*. John Wiley and Sons (2006) ISBN:978-0-471-49196-5
3. Eicker, U.: *Solar Technologies for Buildings*. Wiley (2003)
4. Weisstein, E.: Lambert w-function (2008)
5. Petrone, G., Spagnuolo, G., Vitelli, M.: Analytical model of mismatched photovoltaic fields by means of Lambert w-function. *Solar Energy Materials and Solar Cells*, 1652–1657 (2007)
6. Petrone, G., Ramos-Paja, C.: Modeling of photovoltaic fields in mismatched conditions for energy yield evaluations. *Electric Power Systems Research*, 1003–1013 (2011)
7. Bruendlinger, R., Bletterie, B., Milde, M., Oldenkamp, H.: Maximum power point tracking performance under partially shaded pv array conditions. In: *Proceedings of the Twenty-first European Photovoltaic Solar Energy Conference*, Dresden, Germany, September 4-8, vol. 34(5), pp. 2157–2160 (2006)
8. Giaffreda, D., Omana, M., Rossi, D., Metra, C.: Model for thermal behavior of shaded photovoltaic cells under hot-spot condition. In: *2011 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, pp. 252–258 (October 2011)

9. Karatepe, E., Boztepe, M., Colak, M.: Neural network based solar cell model. *Energy Conversion and Management* 47(910), 1159–1178 (2006)
10. Almonacid, F., Rus, C., Hontoria, L., Fuentes, M., Nofuentes, G.: Characterisation of silicon crystalline pv modules by artificial neural networks. *Renewable Energy* 34(4), 941–949 (2009)
11. Al-Amoudi, A., Zhang, L.: Application of radial basis function networks for solar-array modelling and maximum power-point prediction. *IEE Proceedings-Generation, Transmission and Distribution* 147(5), 310–316 (2000)
12. Zhang, L., Bai, Y.F.: Genetic algorithm-trained radial basis function neural networks for modelling photovoltaic panels. *Engineering Applications of Artificial Intelligence* 18(7), 833–844 (2005)
13. Piliouquine Rocha, M., Mora Lopez, L., Sidrach de Cardona, M., Elizondo, D.A.: Photovoltaic module simulation by neural networks using solar spectral distribution. *Progress in Photovoltaics: Research and Applications* (accepted for publication)
14. Salas, V., Olias, E., Barrado, A., Lázaro, A.: Review of the maximum power point tracking algorithms for stand-alone photovoltaic systems. *Solar Energy Materials & Solar Cells* 90, 1555–1576 (2006)
15. Femia, N., Petrone, G., Spagnuolo, G., Vitelli, M.: Optimization of perturb and observe maximum power point tracking method. *IEEE Transactions on Power Electronics* 20(4), 963–973 (2005)
16. Femia, N., Petrone, G., Spagnuolo, G., Vitelli, M.: A technique for improving p&o mppt performances of double stage grid-connected photovoltaic systems. *IEEE Transactions on Industrial Electronics* 56(11), 4473–4482 (2009)
17. Hussein, K., Muta, I., Hoshino, T., Osakada, M.: Maximum photovoltaic power tracking: an algorithm for rapidly changing atmospheric conditions. *IEE Proceedings-Generation, Transmission and Distribution* 142(1), 59–64 (1995)
18. Spagnuolo, G., Petrone, G., Vitelli, M., Calvente, J., Ramos-Paja, C., Giral, R., Mamarelis, E., Bianconi, E.: A fast current-based mppt technique employing sliding mode control. *IEEE Transactions on Industrial Electronics* PP(99), 1 (2012)
19. Shmilovitz, D.: On the control of photovoltaic maximum power point tracker via output parameters. *IEE Proceedings-Electric Power Applications* 152(2), 239–248 (2005)
20. Petrone, G., Spagnuolo, G., Vitelli, M.: A multivariable perturb-and-observe maximum power point tracking technique applied to a single-stage photovoltaic inverter. *IEEE Transactions on Industrial Electronics* 58(1), 76–84 (2011)
21. Mei, Q., Shan, M., Liu, L., Guerrero, J.: A novel improved variable step-size incremental-resistance mppt method for pv systems. *IEEE Transactions on Industrial Electronics* 58(6), 2427–2434 (2011)
22. Casadei, D., Grandi, G., Rossi, C.: Single-phase single-stage photovoltaic generation system based on a ripple correlation control maximum power point tracking. *IEEE Transactions on Energy Conversion* 21(2), 562–568 (2006)
23. Kimball, J., Krein, P.: Discrete-time ripple correlation control for maximum power point tracking. *IEEE Transactions on Power Electronics* 23(5), 2353–2362 (2008)
24. Brunton, S., Rowley, C., Kulkarni, S., Clarkson, C.: Maximum power point tracking for photovoltaic optimization using ripple-based extremum seeking control. *IEEE Transactions on Power Electronics* 25(10), 2531–2540 (2010)
25. Femia, N., Lisi, G., Petrone, G., Spagnuolo, G., Vitelli, M.: Distributed maximum power point tracking of photovoltaic arrays: Novel approach and system analysis. *IEEE Transactions on Industrial Electronics* 55(7), 2610–2621 (2008)
26. Petrone, G., Spagnuolo, G., Vitelli, M.: Distributed maximum power point tracking: challenges and commercial solutions. *Automatika Journal for Control, Measurement, Electronics, Computing and Communications* (accepted for publication)

27. Cirrincione, M., Pucci, M., Vitale, G.: Growing neural gas based mppt of variable pitch wind generators with induction machines. *IEEE Transactions on Industry Applications* PP(99), 1 (2012)
28. Ngan, M.S., Tan, C.W.: Multiple peaks tracking algorithm using particle swarm optimization incorporated with artificial neural network. *World Academy of Science, Engineering and Technology* 58, 379–385 (2011)
29. Salah, C.B., Ouali, M.: Comparison of fuzzy logic and neural network in maximum power point tracker for pv systems. *Electric Power Systems Research* 81(1), 43–50 (2011)
30. Syafaruddin, Karatepe, E., Hiyama, T.: Artificial neural network-polar coordinated fuzzy controller based maximum power point tracking control under partially shaded conditions. *IET Renewable Power Generation* 3(2), 239–253 (2009)
31. Mohamed, S.A., Shareef, H.: Hopfield neural network optimized fuzzy logic controller for maximum power point tracking in a photovoltaic system. *International Journal of Photoenergy*, 13 (2012)
32. Wai, R.J., Lin, C.Y.: Active low-frequency ripple control for clean-energy power-conditioning mechanism. *IEEE Transactions on Industrial Electronics* 57, 3780–3792 (2010)
33. Radzi, M.A.M., Rahim, N.A.: Neural network and bandless hysteresis approach to control switched capacitor active power filter for reduction of harmonics. *IEEE Transactions on Industrial Electronics* 56, 1477–1484 (2009)
34. Cirrincione, M., Pucci, M., Vitale, G., Miraoui, A.: Current harmonic compensation by a single-phase shunt active power filter controlled by adaptive neural filtering. *IEEE Transactions on Industrial Electronics* 56, 3128–3143 (2009)
35. Mellit, A., Kalogirou, S.A.: Artificial intelligence techniques for photovoltaic applications: A review. *Progress in Energy and Combustion Science* 34(5), 574–632 (2008)
36. del Rio, P., Gual, M.: The promotion of green electricity in europe: present and future. *European Environment* 14(4), 219–234 (2004)
37. Matallanas, E., Castillo-Cagigal, M., Gutierrez, A., Monasterio-Huelin, F., Caamao-Martinez, E., Masa, D., Jimnez-Leube, J.: Neural network controller for active demand-side management with pv energy in the residential sector. *Applied Energy* 91(1), 90–97 (2012)
38. Lin, W.M., Hong, C.M., Chen, C.H.: Neural-network-based mppt control of a stand-alone hybrid power generation system. *IEEE Transactions on Power Electronics* 26(12), 3571–3581 (2011)
39. Giraud, F., Salameh, Z.: Analysis of the effects of a passing cloud on a grid-interactive photovoltaic system with battery storage using neural networks. *IEEE Transaction on Energy Conversion* 14(4), 1572–1579 (1999)
40. Hontoria, L., Aguilera, J., Zufiria, P.: A new approach for sizing stand alone photovoltaic systems based in neural networks. *Solar Energy* 78(2), 313–319 (2005); ISES Solar World Congress 2003
41. Balzani, M., Reatti, A.: Neural network based model of a pv array for the optimum performance of pv system. In: *Research in Microelectronics and Electronics*, 2005 PhD, vol. 2, pp. 123–126 (July 2005)
42. Caputo, D., Grimaccia, F., Mussetta, M., Zich, R.: Photovoltaic plants predictive model by means of ann trained by a hybrid evolutionary algorithm. In: *The 2010 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–6 (July 2010)
43. Lee, C.-Y., Chen, P.-H., Shen, Y.-X.: Maximum power point tracking (MPPT) system of small wind power generator using RBFNN approach. *Expert Systems with Applications* 38(10), 12058–12065 (2011),

<http://www.sciencedirect.com/science/article/pii/S0957417411002600>

doi:10.1016/j.eswa.2011.02.054, ISSN: 0957-4174

Sustainable Energy Microsystems for a Smart Grid

Maria Carmen Falvo, Luigi Martirano, and Danilo Sbordone

University of Rome “Sapienza”,
Electrical Engineering Department,
Via delle Sette Sale 12b – 00184 Rome, Italy
{luigi.martirano, mariacarmen.falvo}@uniroma1.it,
sbordonedanilo@gmail.com

Abstract. The paper deals with the proposal of a new architecture, called Sustainable Energy Microsystem (SEM), for a smart grid project in urban context. SEM includes energy sub-systems (SS) currently independent, such as high efficiency buildings, sustainable mobility systems (Electrical Vehicles and metro transit-systems), dispersed generation from renewables and Combined Heat and Power units. The present paper includes the description of the main SEM elements and some results of an energy analysis on each subsystem, showing the effective possibilities of integration, aimed to energy saving and environmental sustainability.

Keywords: energy, environment, mobility, power systems, smart grid, sustainability.

1 Introduction

In national and international context, many research projects are conducted to study the possible evolution of the distribution of electricity to the so-called smart grid. The European Technology Platform defines a Smart Grid as “an electricity network that can intelligently integrate the actions of all users connected to it - generators, consumers and those that do both – in order to efficiently deliver sustainable, economic and secure electricity supplies” [1].

In other words the aims of smart grids implementation are to allow:

- the connection and operation of generation power plants with different primary resources, size and technology;
- the consumers play an active role in the operation of the power systems that supply them (demand response as ancillary service);

with the final task of significantly reducing the environmental impact of the whole energy system, enhancing the level of system reliability and improving security of supply.

The great interest of the international scientific community towards the topic of smart grid is proved by a vast and recent literature [2]-[12].

The analysis of the scientific literature shows that smart grids are usually investigated from the point of view of distributors, with a top-down approach, that means starting from the evolution of grid control technologies, then focusing on the consequent effects upon active and passive users.

In [13] some of the authors proposed a new architecture, called Sustainable Energy Microsystem (SEM), for a smart grid project, aimed at promoting an aggregation at the customer level, in order to show how several energy sub-systems, presently operated as independent. In particular a SEM includes energy sub-systems currently independent, present in a urban context, such as high efficiency buildings, sustainable mobility systems (Electrical Vehicles and mass transit-systems), dispersed generation from renewables and Combined Heat and Power units.

In the present paper, Section II summarizes the main figures of the proposed architecture of SEM, presented in details in [13]; Section 3 includes the results of a energy analysis, showing the possibilities of integration of the different sub-systems, aimed to energy saving and environmental sustainability; the conclusions are drawn in Section 4.

2 Sustainable Energy Microsystem Proposal

The innovative idea of the authors is to move beyond the current concept of smart grid, including the possibility of an integrated management of energy flows between different sub-systems (SS) currently independent. The global energy system is defined Sustainable Energy Micro-system (SEM).

In SEM, the centrality of customers is proposed not only in economical and commercial terms, but also with functional implications, that means asking the grid customers to be actively involved in the system operation; just as an example, customers could vary their cumulate load profile and/or the power produced locally in the SEM according to the actual hourly cost of energy furnished by the supplier. This enhanced interaction between distributors and customers could help the system operation, with significant enhancements in cost-effectiveness of investments, energy efficiency, optimal use of grid capability and power quality.

The purpose of SEM is to build sustainable micro energy islands connected to the electricity public distribution grid exchange, with the minimization of the energy to the network and the optimization of energy flows in the exchange.

The energy sub-systems (SS) taken into account in the SEM proposed, as shown in Figure 1, are:

- SS for the urban mobility: metro-transit system and trams (SS1);
- SS for the connection to the electric network and the recharging of plug-in electric vehicles (EV) for the surface mobility (SS2);
- SS of final users and high efficiency buildings (SS3);
- SS of dispersed generation from renewable and Combined Heat and Power (CHP) units (SS4).

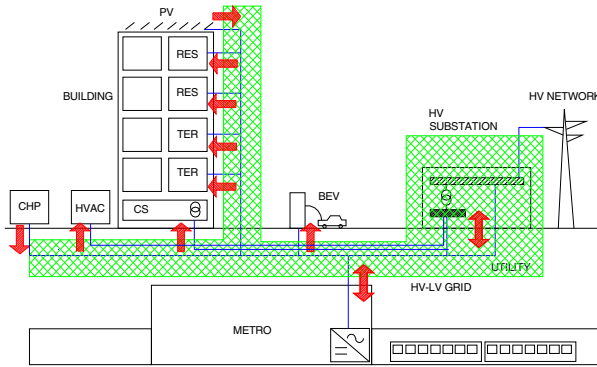


Fig. 1. Energy Subsystems (SS) considered in Sustainable Energy Microsystem (SEM)

In Figure 1 the symbols are:

- RES: residential units,
- TER: tertiary units,
- CS: common service for the building,
- HVAC Heating Ventilation Air Conditioned for the building,
- CHP Combined Heat and Power Generator,
- PV: Photovoltaic generator,
- BEV: Batteries Electric Vehicles System,
- METRO: Metro-transit and tram mobility system;

and the filled area represents the property of the public utility.

The reasons at the basis of the choice of these sub-systems are reported in [16]-[17]. Their integration has been implicitly assigned to the electricity grid to which they are connected, which had to perform the functions of a large flywheel, capable of providing adequate power at any time of integration, as well as to absorb the potential energy of an exuberant sub-system. SS (passive and active users) connected to the network is not sought any form of estimates of their needs or surpluses injected into the network.

Each user is therefore appropriate to a pure traditional limitation of its maximum power of exchange, while maintaining a high degree of autonomy and independence in terms of timing diagram.

The regulatory framework of the major western systems is experiencing a gradual expansion to small users of the logical prediction of timing diagrams of power, so far required only to larger sizes. In an effort to meet the needs of programmable profiles imposed by the new network codes and to observe a behavior perceived as virtuous from the system. So each sub-system is called internally to identify appropriate control variables that allow opportunities for more effective interfacing with the electrical system. But only a single view of the logic design and power management sub-systems neighbors can allow their effective interaction at the local level. The proposed SEM would be able to create a dialogue between the single SS and the

electricity network, in terms of energy trade allowing a full integration, where the sub-systems are aggregated as a single functional block.

The main technological difficulty of realization of SEM, is the ability to identify the optimum energy of the micro-system architectures, since the characteristics of the utilities that constitute the micro-system are complex and complicated and their synchronization.

This complexity extends to the ability of management and continuous optimization of the micro-system to be understood as a real self-operated network to maximize performance, energy efficiency, economic competitiveness and minimize environmental impact. The distributor can then interact in a simplified and smart reference that will think to respond to requests from the distributor through the internal management system.

From traditional layout shown in Figure 1, the proposed integration can be accomplished in various ways (Energy; Distribution Grid and Virtual Power Plant) that have been the subject of research and are deeply described in [13]-[14].

In the present paper just the integration in terms of energy is reminded: it means maintaining the actual distribution structure and introducing a model for management and optimization of energy flows exchanged on the network provider, as shown in Figure 2.

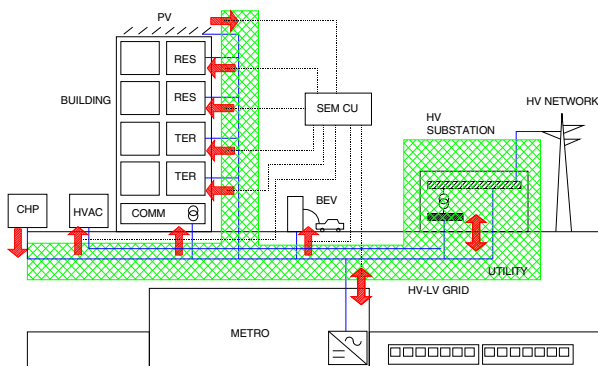


Fig. 2. Architecture for the integration of the SS in terms of energy, with a unique Point of Delivery (POD) by the public utility

The filled area is the same of Figure 1, but a SEM Controller Unit (SEM CU) manages the energetic behaviors of all the SS.

3 Energy Evaluation for the Assessment of Integration Level of SS

In order to understand the potential level of integration among the four SS included in SEM, a preliminary energy analysis on the single SS has been performed.

3.1 Energy Evaluation on a Real SS1

As explained in [15]-[17] a consistent energy saving in metro-transit systems is possible thanks to the recovery of the braking energy of the trains.

With the aim of quantifying the amount of recovered energy in a real metro-transit system, a statistical energy analysis has been performed using the data collected by an experimental survey and some simulation studies performed on lines serving the city of Rome. On these metro-lines trains equipped with drives providing both a regenerative braking and a pneumatic braking system (completely substitutive of the regenerative one) are involved.

The details of the above-mentioned metro-transit system and of the wide analyses are included in other papers of some of the authors [17]-[18].

By experimental and simulation studied, it is possible to evaluate, for different time band of the official daily line timetable, the values of the energy saving percentage ($ES_{\%}$), defined for as:

$$ES_{\%} = \frac{E_{ESS}^{W_REC}}{E_{ESS}^{W/O_REC}} \times 100 \quad (1)$$

where $E_{ESS}^{W_REC}$ is the supplied energy by Electrical Sub-Stations (ESS) in case of recovering of the trains braking energy; E_{ESS}^{W/O_REC} is the supplied energy by ESS without the recovering of the trains braking energy.

This parameter allows to get a measure of impact the recovering of the braking trains energy on the performance of the power system.

Another important parameter is the effective recovered braking energy percentage ($ER_{\%}$) in respect of the recoverable braking energy only in case of braking energy recovering, defined as:

$$ER_{\%} = \frac{E_{TR_REC.ED}}{E_{TR_REC.BLE}} \times 100 \quad (2)$$

where $E_{TR_REC.ED}$ is the effective recovered braking energy; $E_{TR_REC.BLE}$ is the potential recoverable braking energy.

This parameter gives an assessment of the capacity of the line to receive the braking energy by the train: the differential energy is usually dissipated in heat by means of on-board rheostats, but in alternative it could be recovered using storage stationary systems.

ER values assessment is very useful for the energy analysis of SEM: the saved energy in SS1 could be used in SS2 for the recharging of plug-in electric vehicles (EV) for the surface mobility.

ES and ER parameters have to be calculated for different time band, because it is demonstrated that their values are strictly linked to the trains frequency that varies in each time interval. Figure 3 and 4 show this link, and report real values obtained by the simulations of one of Rome metro-line [17].

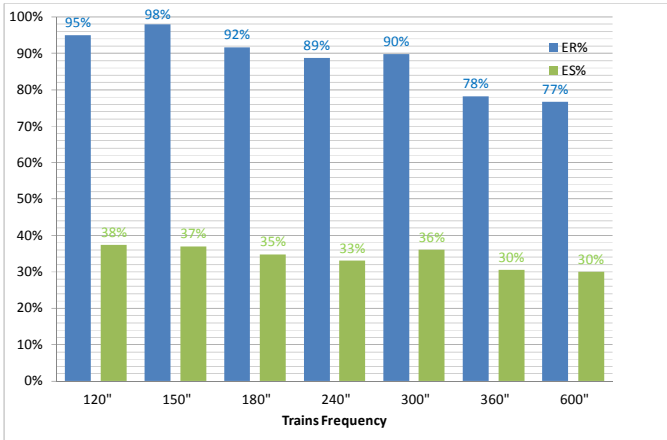


Fig. 3. ES and ER percentage values for different trains frequencies




Fig. 4. Comparison between system potential recoverable and effective recovered braking energy for different trains frequency in 1h of simulation

A more detailed calculation of these magnitudes and a complete correlation analysis with traffic and train figures are reported in [19].

A typical timetable of the metro-transit system of Rome city, that corresponds to a real diagram of the trains frequency in 24 hours service, generally includes:

- 6 hours without service and 18 hours with service;
- 4 hours of service with a trains frequency at 150'';
- 4 hours of service with a trains frequency at 360'';
- 10 hours of service with a trains frequency at 240''.

Referring to the this timetable, the daily ES has been estimated around 30%, and daily difference between the potential recoverable braking energy and the effective recovered braking energy is about 14.5 MWh.

This energy could be stored in batteries located along the traction line in each ESS and used for the energy integration with the other SS of SEM, especially the BEV of SS2.

3.2 Energy Evaluation on Real SS2, SS3 and SS4

Relating to the other SS included in SEM, a real case study has been considered including:

- For SS2: 100 battery electric vehicles BEV (60% high capacity and 40% low capacity) with a peak demand of about 200kW;
- For SS3: 2 buildings with 150 residential and tertiary units, 25 small and medium commercial stores and commons services as parks, gardens and halls, with a peak demand of 400 kW and total energy consumption of about 2000 MWh/year;
- SS4: 1 photovoltaic (PV) generator with a peak power of 100 kW and 1 Combined Heat and Power (CHP) with a peak power of 200 kW.

By simulation models, the authors got the value of power and energy demand/production related to those energy SS. The results are summarized in Table 1.

Table 1. Power and energy demand/rproduction for each SS

SS in SEM		Power Demand/Production [kW]	Annual Energy Demand/Production [MWh/year]
BEV recharging system (SS2)		-200	-800
complex of 2 buildings (SS3)	common services CS (parks, gardens, elevators, hall, etc), (SS3-a)	-50	-400
	residential/tertiary/commercial individual units (SS3-b)	-250	-1100
	heating ventilation air conditioned HVAC common service (SS3-c)	-100	-600
PV power plan(SS4-a)		100	130
CHP system (SS4-b)		200	900

Figures 5 and 6 show the load profile (power demand) simulated for the SS2, SS3-a and SS3-c in the winter (Figure 4) and in the summer season (Figure 5). In fact only the aggregation of those SS in SEM, sharing a unique point of delivery (POD) with the electric distribution, it is possible. The reason is that the current regulation in Europe makes not feasible the aggregation of the SS3-b. Therefore SS3-b has a direct supplying by the distribution grid through individual POD.




Fig. 5. Load profiles in winter season




Fig. 6. SS Load profiles in summer season

Figures 7 and 8 show the power generation profile simulated for the SS4-a and SS4-b, in the winter (Figure 6) and in the summer season (Figure 7).




Fig. 7. Power generation profile in winter season




Fig. 8. Power generation profile in summer season

Figures 9 and 10 show the load, the generation and the balance profiles at SEM POD of SS2, 3 and 4 in summer (Figure 9) and winter (Figure 10) season of the case study.




Fig. 9. Load, generation and balance profiles at SEM POD in summer season




Fig. 10. Load, generation and balance profiles at SEM POD in winter season

The results show that yearly energy balance for the microsystem is acceptable:

- Energy demand: about 1.800 MWh/year;
- Energy generation: about 1.030 MWh/year.

In the summer season the balance of power at SEM POD is acceptable. In the winter season there is an important fraction of power introduced in the distribution network. A portion of this energy could be recovered from SS1.

4 Conclusions

The paper contains a proposal of design of a sustainable energy system (SEM) including the possibility of an integrated management of energy flows among different sub-systems (SS) currently independent, such as high efficiency buildings, sustainable urban mobility, dispersed generation from renewable and Combined Heat and Power (CHP) units. This system is defined Sustainable Energy Microsystem (SEM), and could be considered the basis for a smart grid project in a urban context.

The results of an energy analysis on the single SS, referring to real data as case study, are reported and they point out the effective possibilities of energy integration of the different SEM SS and the advantages that the implementation of the proposed architecture potentially presents, according to the worldwide policies of energy savings and environmental sustainability.

References

1. Strategic Deployment Document for Europe's Electricity Networks of the future. Draft (September 2008)
2. Lo, E.O., Kai, X., Senthil, J., Alaywan, Z., Rothleder, M.: Adjustable inter-Scheduling Coordinator trade modelling in congestion management. In: Proc. Power Engineering Society Summer Meeting 2001, Vancouver, BC (2001)
3. De Tuglie, E., Torelli, F.: Load following control schemes for deregulated energy markets. IEEE Transaction on Power Systems 21(4) (November 2006)
4. Houwing, M., Ilic, M.: The value of IT for Virtual Power Plants with micro cogeneration systems. In: Proc. IEEE International Conference on Systems, Man and Cybernetics, Singapore (2008)
5. Houwing, M., Papaefthymiou, Heijnen, P.W., Ilic, M.: Balancing wind power with virtual power plants of micro-CHPs. In: Proc. Power Tech. 2003, Bucharest (2009)
6. Vukmirovic, S., Erdeljan, A., Kulic, F., Lukovic, S.: Software architecture for Smart Metering systems with Virtual Power Plant. In: Proc. 15th IEEE Mediterranean Electrotechnical Conference, La Valletta (2010)
7. Newman, G., Mutale, J.: Characterising the VPP. In: Proc. 20th International Conference and Exhibition on Electricity Distribution, Cired 2009, Prague (2009)
8. Morais, H., Cardoso, M., Castanheira, L., Vale, Z.: A decision-support simulation tool for virtual power producers. In: Proc. International Conference on Future Power Systems, Amsterdam (2005)

9. Deuse, J.: The UK aggregation experiment combining wind and demand response. In: Proc. 20th International Conference and Exhibition on Electricity Distribution, Cired 2009, Prague (2009)
10. Fernandes, R.A.S., Silva, I.N., Oleskovicz, M.: Identification of residential load profile in the Smart Grid context. In: Proc. PES General Meeting 2010, Minneapolis (2010)
11. Pedrasa, M.A.A., Spooner, T.D., MacGill, I.F.: Coordinated scheduling of residential distributed energy resources to optimize smart home energy services. *IEEE Transactions on Smart Grids* 1(2) (September 2010)
12. Schulke, A., Bauknecht, J., Haussler, J.: Power demand shifting with smart consumers – A platform for power grid friendly consumption control strategies. In: Proc. 1st IEEE International Conference on Smart Grid Communications, Gaithersburg, MD (October 2010)
13. Falvo, M.C., Martirano, L.: From Smart Grids to Sustainable Energy Microsystems. In: Proc. 10th International Conference on Environment and Electrical Engineering, Rome (May 2011)
14. Brenna, M., Falvo, M.C., Foiadelli, F., Martirano, L., Poli, D.: Sustainable Energy Microsystem (SEM): Preliminary Energy Analysis. In: Smart Grid Conference 2012, Washington, D.C. (January 2012)
15. Falvo, M.C., Foiadelli, F.: Preliminary analysis for the design of an energy-efficient and environmental sustainable integrated mobility system. In: Proc. IEEE PES General Meeting 2010, Minneapolis - USA, Luglio, July 25-29 (2010)
16. Falvo, M.C., Lamedica, R., Bartoni, R., Maranzano, G.: Energy Saving in Metro-Transit Systems: Impact of Braking Energy Management. In: Proceedings SPEEDAM 2010 (IEEE), Pisa, Italia, June 14-16 (2010)
17. Falvo, M.C., Lamedica, R., Bartoni, R., Maranzano, G.: Energy management in metro-transit systems: an innovative proposal toward an integrated and sustainable urban mobility system including plug-in electric vehicles. *Elsevier Electric Power System Research Journal* (December 2011)
18. Brenna, M., Foiadelli, F., Roscia, M.C., Zaninelli, D.: Prospective for energy saving by means of Ultracapacitors in electric systems for transportation. In: Fourth International Conference and Exhibition on Ecological Vehicles and Renewable Energies, Monaco, France, March 26-29 (2009)

SVM Methods for Optimal Management of a Virtual Power Plant

Emanuele Crisostomi, Mauro Tucci, and Marco Raugi

Department of Energy and Systems Engineering, University of Pisa, Italy
`emanuele.crisostomi@dsea.unipi.it`

Abstract. The current electrical grid is undergoing a deep renovation that poses new problems in terms of technologies, communication and control. The increasing level of penetration of renewable energy is leading towards the concept of distributed energy production, and it is expected that Virtual Power Plants (VPPs) will play an important role in the future *smart grid*. The stochastic nature of the power flows in the VPP, caused by the fluctuating availability of renewables, by the users' demand and by the energy market price, complicates the task of power balancing for the VPP. This paper proposes the use of simple learning mechanisms to support power scheduling decisions and to improve a correct supply of the connected loads.

Keywords: Support Vector Machine methods, Virtual Power Plants, optimal scheduling.

1 Introduction

1.1 Motivation

The future *smart grid* is expected to facilitate an efficient and sustainable use of natural resources, for instance by integrating renewable sources and by increasing the utilisation of storage systems to cope with the gap between energy demand and energy production. The smart grid refers to a vast infrastructure, composed by many geographically distributed local resources, and as typical complex large-scale networks, poses several problems in terms of management and maintenance. For this reason, it is expected that Virtual Power Plants (VPPs), defined as *clusters of dispersed generator units, controllable loads and storage system, aggregated in order to operate as a unique power plant* [1], will provide a viable answer to the increasing distributed generation [2], by constituting small smart (micro-)grids.

The structure of a VPP is illustrated in Figure 1. The adjective virtual refers to the fact that the VPP is in general a multi-fuel, multi-location and multi-owned power station. However, from a grid operator's perspective, purchasing energy or ancillary services from a VPP is equivalent to purchasing from a conventional station [3]. Figure 1 emphasises the central role of the Energy Management

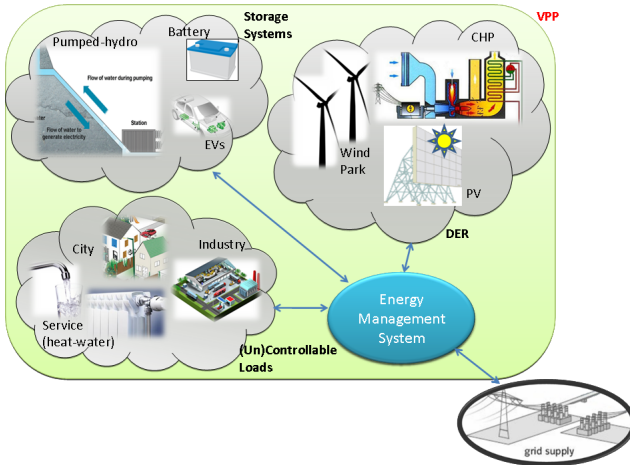


Fig. 1. A Virtual Power Plant (VPP) is generally characterised by Distributed Energy Resources (DERs) (e.g., a Combined Heat and Power (CHP) plant, and wind/solar plants), storage systems, controllable and uncontrollable loads, Electric Vehicles (EVs), and is connected to other VPPs through the grid.

System (EMS) that is responsible of computing the optimal power flows within the VPP. The EMS has to take important decisions in real-time, such as

- how much energy to produce;
- whether to supply or to curtail non-vital controllable loads;
- whether to buy/sell energy to/from other VPPs;
- whether to store/take energy into/from the storage systems.

Decisions will be generally taken in real-time on the basis of instantaneous conditions, i.e., current availability of renewables, current energy price, current energy demand. However, should realistic and accurate predictions of future values of such unknown quantities be available, it makes sense to take decisions also on the basis of future expectations.

1.2 Paper Contribution

This paper extends previous work of some of the authors, where an optimal power scheduling algorithm for the management of VPPs had been proposed [4]. Such a scheduling algorithm was able to compute the optimal power flows that minimised a short-run cost function, on the basis of some instantaneous information of current energy price, and current energy availability. The demand-side load was modelled as a combination of vital loads, and controllable loads with three different priorities. Vital loads, would correspond to the loads that the EMS can never curtail, and should not enough energy be available, then the EMS would take energy from storage systems, or buy energy from the grid to

supply such loads. On the other hand, controllable loads correspond to the loads that the EMS could disconnect under particular conditions (e.g., peak hours, or high energy prices).

Obviously, the EMS has an intrinsic difficulty in automatically deciding when the current energy price is high or low. In paper [4], the energy price was evaluated on the basis of the day-before energy price profile. The objective was to use such a day-before profile to find the price that would allow loads to be connected for exactly 25%, 50% or 75% of the day, depending on their priority. This paper substitutes such a simple one-day-before predictor with a more complicated intelligent mechanism to provide a more accurate prediction of the next day energy price.

Several papers in the literature have investigated algorithms able to predict future electrical energy prices, see for instance [5], [6], [7]. The novelty of the method proposed here lays both in the type of algorithm adopted and in the choice of the utility function used to compare the different algorithms.

This paper is organised as follows: Section 2 reviews the algorithm introduced in [4], and illustrates its application for the optimal management of renewable energy resources. Section 3 describes the innovative modification to such an algorithm, through the adoption of a regression algorithm based on support vector machines (SVM) to predict more accurate energy prices, than just using the day-before energy profile. Section 4 shows the benefits of the proposed method through realistic data taken from [8]. Finally, Section 5 summarises the obtained results and outlines current and future lines of research.

2 EMS Optimal Power Scheduling Algorithm

The optimisation algorithm introduced in [4] computes every 30 minutes the energy (in *MWh*) produced/sold/bought/supplied by each of the components of the VPP, according to the solution of the following linear optimisation problem:

$$\begin{aligned}
 \min \quad & p_w \mathbf{E}_w + p_{PV} \mathbf{E}_{PV} + p_{CHP} \mathbf{E}_{CHP} + \\
 & + p_{grid} \mathbf{E}_{storage}^{fw} + p_{grid} \mathbf{E}_{storage}^{hy} + p_{grid} \mathbf{E}_{grid} + \\
 & - p_{CL}^1 \mathbf{E}_{CL}^1 - p_{CL}^2 \mathbf{E}_{CL}^2 - p_{CL}^3 \mathbf{E}_{CL}^3
 \end{aligned}$$

$$\text{subject to } \left\{ \begin{array}{l}
 \mathbf{E}_w + \mathbf{E}_{PV} + \mathbf{E}_{CHP} + \\
 + \mathbf{E}_{storage}^{fw} + \mathbf{E}_{storage}^{hy} + \mathbf{E}_{grid} = \\
 \mathbf{E}_{UL} + \mathbf{E}_{CL}^1 + \mathbf{E}_{CL}^2 + \mathbf{E}_{CL}^3 \\
 \underline{\mathbf{E}}_w \leq \mathbf{E}_w \leq \overline{\mathbf{E}}_w \\
 \underline{\mathbf{E}}_{PV} \leq \mathbf{E}_{PV} \leq \overline{\mathbf{E}}_{PV} \\
 \underline{\mathbf{E}}_{CHP} \leq \mathbf{E}_{CHP} \leq \overline{\mathbf{E}}_{CHP} \\
 \underline{\mathbf{E}}_{storage}^{fw} \leq \mathbf{E}_{storage}^{fw} \leq \overline{\mathbf{E}}_{storage}^{fw} \\
 \underline{\mathbf{E}}_{storage}^{hy} \leq \mathbf{E}_{storage}^{hy} \leq \overline{\mathbf{E}}_{storage}^{hy} \\
 -NTC/1h \leq \mathbf{E}_{grid} \leq NTC/1h \\
 0 < \mathbf{E}_{CL}^i < \overline{\mathbf{E}}_{CL}^i, \quad i = 1, 2, 3
 \end{array} \right. \tag{1}$$

In the optimisation problem (II), E and p denote energies and prices respectively; although prices and energies change every 30 minutes, the dependence on the time step has been dropped from the notation for simplicity; the *unknowns* have been represented with bold terms to improve readability; in practice, energies are the unknowns, and as time steps are fixed, this is equivalent to saying that power flows are the unknowns; the known terms are however known only at the same time step the optimisation problem is solved (prices are not predicted); the terms w , PV , CHP , CL , UL , fw and hy refer to wind, photovoltaic, Combined Heat and Power, Controllable Loads, Uncontrollable Loads, flywheel and pumped hydro respectively; the utility function is a direct costing function: therefore energy production terms appear with a positive sign; the sign of the storage systems can be positive (energy is released) or negative (energy is stored); the sign of the energy exchanged with the grid (E_{grid}) is positive if energy is bought and negative if sold; controllable loads have a negative sign to make it attractive for the EMS to drive them (if possible). The equality constraint is given by the power balance (energy produced, released from storage systems or bought from the grid must equal to that used to drive all loads plus eventual energy that might be stored or sold to the grid).

The first three *inequality constraints* represent the minimum and maximum production of energy from each distributed resource, upper/lower bounds indicated with upper/lower bars. Bounds are given by the (stochastic) availability of wind/solar energy, by the net capacity of the plant and by the energy production in the last time interval (energy production can not be modulated arbitrarily). Similarly, the bounds of the storage systems are both given by the capacity of the storages, and by whether the storage systems is full or empty at a given time. The energy exchanged with the grid has bounds arising from the Network Transfer Capacity (NTC, expressed in MW). The controllable loads have been prioritised in three categories (the smallest number 1 implying the highest priority), so that the EMS knows which category of controllable loads should be disconnected if not enough energy is produced, and it is too expensive to buy extra energy from the grid to run them.

Prices associated with storing energy and driving or curtailing the controllable loads do not have a direct interpretation, and the following discussion relates to the optimal strategy of estimating the price of controllable loads.

2.1 Optimal Price of Controllable Loads

In many occasions controllable loads are not associated with a price, but with a desired number of hours they should be driven. Following the approach of [4], we assume that controllable loads can be divided in three categories according to their priority, and it is required that they are connected 6, 12 and 18 hours per day. If they are connected for a shorter time, then the quality of service is lower than required by the user; on the other hand, if they are connected for more time than required, then the VPP pays an extra unnecessary cost.

A simple automatic way of finding an appropriate cost, is to choose the cost of controllable loads as the price level such that 25%, 50% or 75% of the time the

energy price is below such a level. In this way, the EMS can decide to buy energy directly from the grid to drive the controllable loads. In [4], the 25, 50 and 75 percentile of the energy prices were estimated from the day before energy price profile; thus, should the energy price of the following day be the same, then the problem would have been optimally solved. In this paper, we predict the 25, 50 and 75 percentile of the energy prices using a SVM regression algorithm that exploits a longer past history of data. Next section illustrates in detail the new mechanism to predict one-day ahead energy prices. Section 4 applies the proposed algorithm to energy prices in the month of February 2012 and shows that the proposed method is more accurate than the day-before assumption in predicting the energy price profile, and in particular the price thresholds such that 25%, 50% and 75% of the time the prices will be below the thresholds.

3 Price Prediction Based on SVM

In order to estimate the 25, 50 and 75 percentile of the price for each day, a prediction algorithm of the price has been developed adopting an SVM regression algorithm. The formulation used in this work is the support vector regression ϵ -SVR illustrated in [9]. A practical reference for applying the previous methodology can be found in [10]. Consider a set of l training points, $\{(\mathbf{x}_1, z_1), \dots, (\mathbf{x}_l, z_l)\}$, where $\mathbf{x}_i \in \mathbb{R}^n$ is a feature input vector, and $z_i \in \mathbb{R}$ is the target output. Then the ϵ -SVR regressor learns the pairs relationships $\mathbb{R}^n \rightarrow \mathbb{R}$, and after the training it can be used to predict an output z_i given a new input \mathbf{x}_i .

In order to predict the energy prices, the training set was created from the price time series $p_{d,h}$, containing hourly energy prices, where the subscripts d and h denote the day and the hour respectively. As input database we used $\mathbf{x}_{d,h}$, $d = 1, \dots, L-1$, where each vector $\mathbf{x}_{d,h} \in \mathbb{R}^{47}$ is related to the hour $i = 24(d-1)+h$ and is composed as $\mathbf{x}_{d,h} = [p_{d,1}, p_{d,2}, \dots, p_{d,24}, p_{d+1,1}, p_{d+1,2}, \dots, p_{d+1,h-1}, 0, 0, \dots, 0]^T$. So, in practice, each $\mathbf{x}_{d,h}$ contains the 24 prices of the day d , the prices of the following day $d+1$ up to hour $h-1$, and the remaining elements are all zeros. The output target $z_{d,h}$ associated with $\mathbf{x}_{d,h}$ is the price $p_{d+1,h}$, which is the next price in the time series after the last value contained in $\mathbf{x}_{d,h}$, so $z_{d,h} = p_{d+1,h}$.

Overall, the training set is composed of $l = 24(L-1)$ training points $(\mathbf{x}_{d,h}, z_{d,h})$, $d = 1, \dots, L-1$, $h = 1, \dots, 24$. In this way the regressor is trained to predict the price at hour h of a current day, using all the known prices of the previous day plus those of the current day up to hour $h-1$; so that the time window of previous points used has a variable size of “24+h-1” and always begins at the hour 1 of the previous day. Such a choice is expected to perform better than a time window with a fixed number of previous points. In fact the price patterns show a strong day-based periodicity, and the use of input patterns that start at a fixed time of each repetition improves the learning of the day periodicity in the regression algorithm. The ϵ -SVR algorithm as formulated in [9] requires the choice of two real parameters, $C > 0$ and $\epsilon > 0$, and to select a kernel function $k(\mathbf{x}_i, \mathbf{x}_j)$ which defines the inner product for the input points \mathbf{x}_i . In this work we have used the linear kernel, which is the inner product in \mathbb{R}^n , $k(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_i^T \mathbf{x}_j$, and the

parameters $C > 0$ and $\epsilon > 0$ were obtained by means of 10-fold cross-validation. It is important to note that the data is preprocessed before the training, using a linear scaling of the price time series to the range $[-1, 1]$.

The calculation of the 25, 50 and 75 percentile for the upcoming day d is done as follows. First we train the ϵ -SVR with the known price data $p_{d,h}$ of the last $L = 40$ days, i.e., from day $d - 40$ to day $d - 1$. In this way the training set is composed by $24(L - 1) = 936$ training points $(\mathbf{x}_{d,h}, z_{d,h})$. After the training the regressor is used to predict the 24 values, for $h = 1, \dots, 24$, of the energy price $\hat{p}_{d,h}$ of the upcoming day d . This is accomplished by using the regressor 24 times, calculating $\hat{p}_{d,h}$ as the output of the regressor to the input composed as $\mathbf{x}_{d,h} = [p_{d-1,1}, \dots, p_{d-1,24}, \hat{p}_{d,1}, \dots, \hat{p}_{d,h-1}, 0, 0, \dots, 0]$.

4 Simulations

The algorithm illustrated in the previous section is now adopted to predict the hourly energy prices of the next day. The prediction is used to estimate the 25, 50 and 75 percentile and to estimate the price thresholds $p_{CL}^1, p_{CL}^2, p_{CL}^3$ required by the optimisation algorithm [1]. In this section we show that the proposed algorithm improves the original simple strategy used in [4] (i.e., estimating the same percentile from the previous day data). Figure 2 compares the true hourly price, with the one predicted using the illustrated SVM based methodology and the simple day-before strategy. The figure refers to energy price data taken from [8] and relative to the months of January and February 2012 in the Italian market. Finally, we use the daily prediction to obtain the thresholds, and we compare the accuracy of the percentile, i.e., we check whether the following day the energy price is indeed below the price thresholds the 25, 50 and 75% of the time, at the time instants of interest. For this purpose, we define the daily error




Fig. 2. Comparison between the true energy prices, the SVM-based prediction and the day-before prediction




Fig. 3. Comparison between the error obtained adopting a simple prediction strategy based on the day before data, and the more complex SVM based strategy. Overall, the number of wrong hours decreases from 183 to 162, with about a 10% improvement.

e_d for day d as the sum of the absolute values of the differences between the number of hours that the true price in day d is smaller than the three thresholds p_{CL}^1 , p_{CL}^2 , p_{CL}^3 , according to

$$e_d = |nh_{CL}^1 - 6| + |nh_{CL}^2 - 12| + |nh_{CL}^3 - 18|. \quad (2)$$

In this way e_d indicates the total number of wrong hours and should be as small as possible. Therefore, if the predicted thresholds p_{CL}^1 , p_{CL}^2 and p_{CL}^3 were correctly predicted, then the next day the controllable loads would stay connected for a number of hour (nh_{CL}^1 , nh_{CL}^2 and nh_{CL}^3) equal to respectively 6, 12 and 18 hours. Thus, e_d would be zero.

Figure 3 compares the error e_d obtained using the thresholds from the previous day and the thresholds from the SVM predictor. With the proposed method, in a range of 19 days, we obtain a 10% improvement. The result is particularly good as the regressor had been trained to learn the price pattern, and not to predict the thresholds directly. While the results can be in principle improved by training the predictor to learn the thresholds, the proposed approach has a general validity and can be used for other purposes than the specific ones investigated here, motivated by the adoption of the optimisation algorithm (II).

5 Conclusions

This paper describes the use of SVM theory to support the EMS in taking optimal power scheduling decisions. The implemented SVM predictor improves the results originally obtained in [4], where a simple day-before predictor was used, as it is reported in Section 4. It is important to remark that the implemented predictor not only estimates the price thresholds strictly required for the algorithm

(II), but the whole next day profile of hourly energy prices. Such a capability can be clearly used for general purposes in the context of the VPP, for instance to predict the best moments of the day for storing or releasing energy in the storage systems.

The proposed SVM predictor is a simple general purpose approach that improves the simple heuristic methods previously adopted. However, a more extensive comparison of current practises in the literature will be required to identify the most appropriate methodology as a decisional support for optimal power scheduling in a VPP. This problem is currently being investigated by the authors.

References

1. Lombardi, P., Powalko, M., Rudion, K.: Optimal operation of a virtual power. In: IEEE Power & Energy Society General Meeting (PES 2009) (2009)
2. El Bakari, K., Kling, W.L.: Virtual Power Plants: an answer to increasing distributed generation. In: Innovative Smart Grid Technologies Conference, Europe (2010)
3. European Commission, Community Research: New ERA for electricity in Europe. Distributed Generation: Key issues, challenges and proposed solutions, Brussels (2003)
4. Aloini, D., Crisostomi, E., Raugi, M., Rizzo, R.: Optimal Power Scheduling in a Virtual Power Plant. In: Innovative Smart Grid Technologies Conference, Europe (2011)
5. Mohsenian-Rad, A.-H., Leon-Garcia, A.: Optimal residential load control with price prediction in real-time electricity pricing environments. *IEEE Transactions on Smart Grid* 1(2) (2010)
6. Crisostomi, E., Tucci, M., Raugi, M.: Methods for Energy Price Prediction in the Smart Grid. In: Innovative Smart Grid Technologies Conference, Europe (2012)
7. Cartea, Á., Figueroa, M.G.: Pricing in Electricity Markets: A Mean Reverting Jump Diffusion Model with Seasonality. *Applied Mathematical Finance* 12(4), 313–335 (2005)
8. GME (Gestore Mercati Energetici): Daily and Monthly Electricity Prices and Volumes (2012), <http://www.mercatoelettrico.org/En>
9. Vapnik, V.: *Statistical Learning Theory*. Wiley, New York (1998)
10. Chang, C.C., Lin, C.J.: LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology* (2011)

Active Power Losses Constrained Optimization in Smart Grids by Genetic Algorithms

Gian Luca Storti¹, Francesca Possemato¹, Maurizio Paschero¹,
Silvio Alessandrini², Antonello Rizzi¹, and Fabio Massimo Frattale Mascioli¹

¹ University of Rome “La Sapienza”, Information Engineering,
Electronics and Telecommunications Department,
Polo per la Mobilità Sostenibile (POMOS) Laboratories

`maurizio.paschero@pomos.it`

² ACEA Distribuzione S.p.A.

`alessandrini.silvio@aceaspa.it`

Abstract. In this paper the problem of the minimization of active power losses in a real Smart Grid located in the area of Rome is faced by defining and solving a suited multi-objective optimization problem. It is considered a portion of the *ACEA Distribuzione S.p.A.* network which presents backflow of active power for 20% of the annual operative time. The network taken into consideration includes about 100 nodes, 25 km of MV lines, three feeders and three distributed energy sources (two biogas generators and one photovoltaic plant). The grid has been accurately modeled and simulated in the phasor domain by Matlab/Simulink, relying on the SimPowerSystems ToolBox, following a Multi-Level Hierarchical and Modular approach. It is faced the problem of finding the optimal network parameters that minimize the total active power losses in the network, without violating operative constraints on voltages and currents. To this aim it is adopted a genetic algorithm, defining a suited fitness function. Tests have been performed by feeding the simulation environment with real data concerning dissipated and generated active and reactive power values. First results are encouraging and show that the proposed optimization technique can be adopted as the core of a hierarchical Smart Grid control system.

1 Introduction

The wide diffusion of Distributed Generation (DG) represents a possible development of modern electrical distribution systems that can evolve towards Smart Grids (SG). A rigorous definition of the term “Smart Grid” is somewhat difficult. In fact, due to the fast evolution of the technology, it is quite hard to mark out a sharp boundary including all the aspects associated with this terminology. A widely adopted definition states that a SG is an electrical network able to perform an intelligent integration of all the users connected to it (*i.e.* producers and consumers), with the purpose of distributing the electrical power in a safe,

efficient and sustainable fashion. Tautly, it can be stated that a SG is a new generation electrical network where smartness, dynamicity, safety and reliability are achieved through the use of Information Communication Technologies (ICT) [47]. In fact, ICT can be considered an important support to the migration of traditional electrical infrastructure toward SG. Although the size of the actual electric networks has been improved in order to follow the always increasing power requests, this growth has been achieved without a global planning finalized to the optimization of the energy transportation. Moreover the backbones of the existing infrastructures have been built when the location of the main power users, such as industries, were much different from the actual configurations and when the DG was not even a theoretical concept. For these reasons distribution networks have been implemented in a hierarchic fashion. Electric power is distributed to the final user through an unidirectional transportation infrastructure. This configuration implies a considerable transportation consumption due to the long distance between producers and consumers. Finally, the available electric distribution infrastructures are inadequate to the future requirements. The main problems concerning actual networks are listed below:

- Losses due to long distance between producers and users
- Not optimal management of energetic flows
- Inefficient use of DG related to renewable energy generators
- Lag in the reaction time in case of blackout
- Incomplete and inaccurate knowledge on the instantaneous status of the infrastructure

As stated before, most of these problems can be solved by improving the actual infrastructures with the aid of ICT. More precisely a large number of sensors must be installed on the network in order to obtain a complete information on the instantaneous status of the infrastructure. This information can be used as the input of an optimization control algorithm capable to determine in real time the best network configuration in order to satisfy the instantaneous power request and to drive suitable actuators in order to achieve the optimal configuration. DGs can impact the bus voltage, line power flow, short-circuit current and power network reliability, so that it is very important in SG design and realization to be able to control DGs [35].

In the literature there is an increasing number of publications concerning the use of computational intelligence (CI) in SG [9,10]. Considering a SG as a complex, dynamic, nonlinear and stochastic system, CI can provide support for designing safer and more efficient control systems, in line with emerging technologies. From the point of view of the CI, the SG managing and control is a highly complex problem given the non-linearity and the dynamic of the system, as well as the heterogeneity of the elements that compose it (generators, transformers, transmission lines, time-variant loads, telecommunications system, market regulations). As well know, the main feature of control systems is the ability to run in real-time (unless the system is simulated). Neural and fuzzy approaches seems to be the main candidates, given the universality they offer to model any system. The distinction, considering a large-scale vision of a SG as a System

of Systems (SoS), is between Distributed Local and Wide-Area Monitoring and Control. In the first case neural techniques are used for learning and tracing the system dynamics in order to implement operations such as constraints and operating points settings. Neural algorithms based on Multilayer Perceptron (MLP) and Radial Basis Functions (RBFs) allow to control elements such as turbine generators, solar and wind installations or transmission lines. Techniques known as Artificial Immune System (AIS) allow adaptive control strategies without the need for off-line learning. The ability to control a power system depends on the quality of sensors and the reliability of communication infrastructures. Errors and failures in these systems may easily cause incorrect control schemes with serious consequences. To this aim Swarm Intelligence techniques can be used to recover data from faulty sensors.

Among the variety of techniques offered by CI the use of Genetic Algorithms (GA) seems to be a promising technique. In [6] an adaptive genetic algorithm is used to establish the best distributed generation siting and sizing on a distribution network, showing that the optimal siting and sizing of DG units can effectively reduce the network loss and improve the system voltage level. In [8] it is shown that GAs can deal well with the stochastic nature of the distribution grid and can be successfully used as an optimization method for solving the control problems. Beside theoretical studies it is important to have the opportunity to validate the designed optimization strategy on real data. Moving in this direction, a cooperation with *ACEA Distribuzione S.p.A.* [1] has been engaged with the aim to design a control strategy for the SG under development in the west area of Rome. The project concerning the upgrade of the actual network to up-to-date SG technology fulfils the requirements imposed by AEEG resolution 39/10 [2]. A complete simulator of the considered real network has been implemented; it is described in Sec. 2. In Sec. 3 the multi-objectives optimization problem is formulated and the use of a GA is proposed in order to solve it. In Sec. 4 it is shown how the proposed control strategy can be successfully used to modulate the power fed into the network by DGs in order to reduce active power losses taking into account suited constraints on voltages and currents levels, as well as the available working points of DGs. Finally, conclusions and works in progress are discussed in Sec. 5.

2 Network Simulation

The network under consideration is located in the west area of Rome. It is constituted of about 100 nodes and it is made up of:

- N.3 feeders at 20 kV
- N.2 transformers High Voltage/Middle Voltage (HV/MV)
- N.2 biogas generators
- N.1 photovoltaic generator
- 25 km of MV lines
- N.11 three phase breakers

The generators and the loads are driven by 2 inputs: Active Power (P) and Reactive Power (Q). The real behaviour of the SG can be simulated by means of P and Q yearly power profile. Moreover, in order to reproduce the possibility to supply each MV feeder from a different substation and therefore to change the topology of the network, a boundary switch (Breaker) is placed at the beginning and at the end of each line.

The proposed simulator has been implemented following a Multi-Level Hierarchical and Modular approach. The Multi-Level Hierarchical design improves, through the definition of suitable I/O interfaces, the readability of whole SG simulation model; the Modular approach allows to change, in a simple way, all the parameters of each component models. The structure of the SG simulator is implemented using the MatLab/Simulink SimPowerSystems ToolBox, which allows to rapidly and easily build models that simulate power systems. The SG simulation model is made up of 2 macro blocks: the Input Network and the Electrical Network. The interconnection of the blocks constituting the simulator are shown in Fig. 1. The first macro block, Input Network, is fed by the profiles of P and Q of all loads and generators coming from real measures. They have been saved in different *Excel* files, one for each feeder. These power profiles represent real data acquired with a time step of 1 hour. In the second macro block, Electrical Network, there are the HV, the MV and the LV networks, together with the State Breakers block, that, through several flags, sets the topology of the SG. In the HV network there are 2 transformers with 150 kV at the primary winding and 20 kV at the secondary winding; in the MV network the lines, the loads and the photovoltaic plant are modeled. In the LV network there are the two biogas generators. The lines are modeled using an equivalent model, given by MatLab/Simulink, based on lumped parameters modeling approach (Π model); the transformers are modeled using a block, also provided by MatLab/Simulink, called Three-Phase Transformer Inductance Matrix Type (Two Windings); the

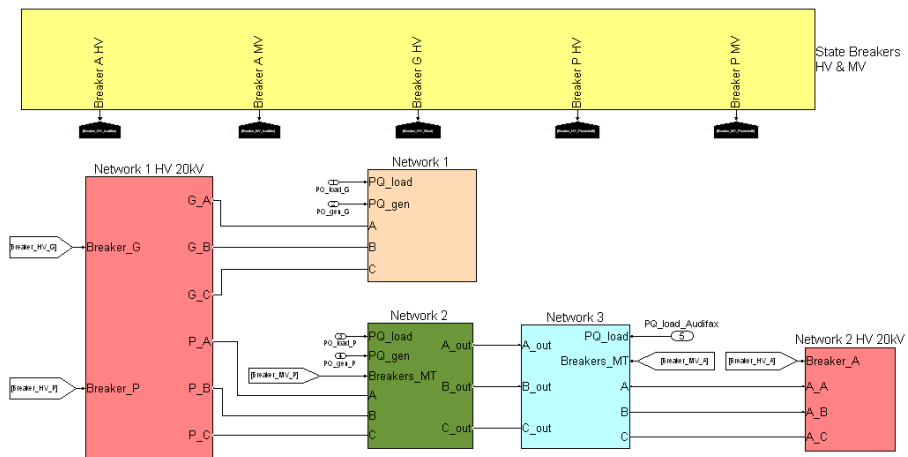


Fig. 1. Interconnection among simulator sub blocks

power driven loads and generators are modeled using the same custom block. This block is essentially a voltage controlled current source, with $I = (2P/V)^*$ where the value of P is read from the corresponding data file and the voltage V is measured at the three phase port of the block modeling the load (* represents the conjugate operator). The SimPowerSystem of MatLab/Simulink allows to use, for three phase network simulation, three solution methods: Continuous, Discrete and Phasor. Since the simulation sampling time is equal to one hour, it is possible to consider exhausted any transient response. For this reason the electrical network analysis has been carried out with the Phasor method.

3 Optimization Procedure

In this section it is described how the considered active power losses optimization problem can be formulated in terms of a multi-objective optimization problem and solved by adopting a suitable evolutionary computation approach. The faced problem consists in finding the optimal network parameters that minimize the value of the total active power losses in the network, considering the constraints imposed on voltages and currents due to safety or quality of service issues. Consider a linear space K , an admissible set E and a cost function $J : E \rightarrow \mathbb{R}$ that associates a real number to each element in E . The problem consists in minimizing the function J in $E \subset \mathbb{R}^\nu$, where ν is the dimension of the space K . The admissible set E is defined through the inequality constraints $g_i(\mathbf{k}) \leq 0 \quad i = 1, \dots, m$, where m is the number of constraints of the problem. Without loss of generality it has been considered possible to measure the voltage $\mathbf{V}(\mathbf{k})$ and the current $\mathbf{I}(\mathbf{k})$ at all locations in the network in order to compute the cost function. In this paper it is assumed that it is possible to control both the active and the reactive power of each biogas generator in a given range though a couple of parameters. More precisely, the ratio k between the nominal power and the active power and the phase ϕ . Moreover it is assumed that the active power of the photovoltaic generator can be modulated in a given range. Summarizing, the dimension ν of the parameters vector $\mathbf{k} = \{k_1, k_2, k_3, \phi_1, \phi_2\}$ has been set equal to 5.

Let's define:

$$A = \{\mathbf{k} \in \mathbb{R}^5 : 0.75 \leq k_1, k_2, k_3 \leq 1, -0.2 \leq \phi_1, \phi_2 \leq 0.45\} \tag{1}$$

$$B = \{\mathbf{k} \in \mathbb{R}^5 : V_j(\mathbf{k}) - 1.1V_{nom_j} \leq 0, \quad 0.9V_{nom_j} - V_j(\mathbf{k}) \leq 0, j = 1, \dots, N\} \tag{2}$$

in which N represents the number of nodes of the network and V_{nom_j} is the nominal value of the voltage of the j -th node,

$$C = \{\mathbf{k} \in \mathbb{R}^5 : |I_j(\mathbf{k})| - I_{max_j} \leq 0, j = 1, \dots, R\} \tag{3}$$

in which R represents the number of branches and I_{max_j} is the maximum current allowed in the j -th wire,

$$D = \{\mathbf{k} \in \mathbb{R}^5 : k_1 < \cos(\phi_1), k_2 < \cos(\phi_2)\} \tag{4}$$

The set D has been defined analysing the Capability Curve of biogas generators that establishes safe operational limits. A capability curve is defined as a curve which shows boundaries of the area on the kW-kVar diagram within which a machine may be operated continuously.

Now, it is possible to define the admissible set E as follows:

$$E = A \cap B \cap C \cap D \quad (5)$$

The cost function J has been defined as follows:

$$J(\mathbf{k}) = \frac{P_{loss}(\mathbf{k})}{P_{gen}(\mathbf{k})} = \frac{P_{gen}(\mathbf{k}) - P_{load}}{P_{gen}(\mathbf{k})} \quad (6)$$

where $P_{gen}(\mathbf{k})$ is the total power generated by all sources, P_{load} is the total power absorbed by the loads, and their difference $P_{loss}(\mathbf{k})$ represents the total losses in the network.

Given a particular determination $\bar{\mathbf{k}}$ of the vector \mathbf{k} , the value returned by (6) is equal to the normalized total active power losses in the network, and can be considered as a measure of how well $\bar{\mathbf{k}}$ solves the optimization problem. Since it is not practically possible to derive expression (6) in closed form as a function of \mathbf{k} , in this paper a genetic algorithm (derivative free approach) has been employed. A GA is a search method based on the principles of natural selection and evolution which selects individuals with high adaptation to environmental conditions as candidates to survive and being part of the following generation of individuals. Moreover, satisfying constraints (5) and minimizing expression (6) are two conflicting objectives, since active power loss is minimized when the voltage across the line is high. Consequently, the constrained optimization problem can be faced by defining a multi-objective optimization, by relying on the following fitness function:

$$F(\mathbf{k}) = \alpha J(\mathbf{k}) + (1 - \alpha) \Gamma(\mathbf{k}) \quad (7)$$

where α is a coefficient between 0 and 1 and it is used to adjust the relative weight of the power losses term $J(\mathbf{k})$ over the constraints term $\Gamma(\mathbf{k})$. The function $\Gamma(\mathbf{k})$ is defined as follows:

$$\Gamma(\mathbf{k}) = \beta \Gamma_{NLC}(\mathbf{k}) + (1 - \beta) [\gamma \Gamma_I(\mathbf{k}) + (1 - \gamma) \Gamma_V(\mathbf{k})] \quad (8)$$

in which β and γ are real numbers between 0 and 1. The function $\Gamma_V(\mathbf{k})$ is a measure of how much the constraints on voltages are violated, $\Gamma_I(\mathbf{k})$ evaluates the violation of the constraints on currents and $\Gamma_{NLC}(\mathbf{k})$ is a measure of the violation of the nonlinear constraints defined in (4). The parameter β is used to assign the relative weight of the nonlinear constraints violation with respect to the term measuring how much voltages and currents are far from the admissible range. In the same way, the γ parameter adjusts the relative weight of the violation of current constraints with respect to the term related to voltages violation.

4 Tests and Results

In order to test the effectiveness of the proposed optimization procedure, the first generation of the genetic algorithm has been generated in a random way, in the domain A of the network parameters defined in (1). This initialization does not necessarily guarantee the satisfaction of the constraints defined in (2), (3) and (4), considered in the definition of the chosen fitness function. In this way, it is possible to verify if the optimization algorithm is able to restore the network in a safe configuration satisfying all the constraints, possibly minimizing the total active power losses. For this test the behaviour of the control system has been simulated and validated in a single time sample (one hour). All the simulations have been realized using the Matlab Global Optimization Toolbox together with the developed network simulator.

Setting the nominal voltage for the low voltage lines (LV) to $V_{nom} = 400$ V and for the medium voltage lines (MV) to $V_{nom} = 20800$ V, the voltages constraints on the nodes of the network are expressed as follows:

$$360 \text{ V} < V_{LV} < 440 \text{ V}, \quad 18710 \text{ V} < V_{MV} < 22880 \text{ V} \tag{9}$$

Moreover, considering $I_{MAX} = 270$ A as the maximum allowed current in the network breakers and I as the current flowing in a given network branch, it is possible to define the currents constraints as follows:

$$|I| < 270 \text{ A} \tag{10}$$

Without loss of generality, voltages and currents have been measured only in some critical nodes and branches taken into consideration on the basis of the network topology. Four runs of the genetic algorithm have been done. The proposed solution represents the mean value of the results obtained in all the different simulations. The number of generations has been set to 55 and the number of individuals for each generation has been set to 10. A few comments can be made on the results showed in Fig. 2. First of all it should be noted that the GA produces 50 generations instead of 55 due to early convergence to a satisfying solution. In addition, it can be seen that the trend of the normalized voltage level remains almost constant into allowed ranges (Fig. 2 part (a)). Moreover the GA reduces the current level of the first and the second feeder (Fig. 2 parts (b) and (c)) approximately by 2% and 30% respectively. The decrease of the current level on the branches induces a reduction of the power loss as can be seen in Fig. 2 part (d). After 30 generations the trend of the power loss becomes almost stable and the overall decrease is approximately 2% of its original value. Looking at Fig. 3 it can be seen that the action performed by the GA is to increase the power contribution associated with all the DGs. In fact, the gains k_1 , k_2 and k_3 of all generators and the power factors $\cos(\phi_1)$ and $\cos(\phi_2)$ of the two biogas generators tend to unity, trying to achieve the maximum injection of active power in the network (see Fig. 3 part (a), (b), (c), (d) and (e)), reducing in this way the power request from the main grid. Finally in Fig. 3 (f) it can be seen the trend of the fitness function, that decreases with the increasing of the




Fig. 2. Evolution of the normalized voltage and current levels and power losses versus GA generations




Fig. 3. Evolution of GA parameters and fitness function versus GA generations

numbers of generations. Comparing Table 1 and 2 it is possible to figure out the percentage of active power absorbed (AAP), capacitive reactive power (CRP), active power generated (GAP) and inductive reactive power (IRP) generated by the loads, distributed generators, network and balance nodes in the initial and final configuration of the network. Some relevant comments can be made about

Table 1. Percentage of the distribution of power in the Loads, Distributed Generators, Network and at the Balance Nodes in the initial configuration

	Loads	Distributed Generators	Balance Nodes	Network
AAP	0,9955	0,0000	0,0000	0,0045
GAP	0,0000	0,1467	0,8533	0,0000
IRP	0,9914	0,0086	0,0000	0,0000
CRP	0,0000	0,0000	0,7266	0,2734

Table 2. Percentage of the distribution of power in the Loads, Distributed Generators, Network and at the Balance Nodes in the final configuration

	Loads	Distributed Generators	Balance Nodes	Network
AAP	0,9956	0,0000	0,0000	0,0044
GAP	0,0000	0,1717	0,8283	0,0000
IRP	1,0000	0,0000	0,0000	0,0000
CRP	0,0000	0,0161	0,7077	0,2762

the absorbed and generated active power in both conditions. In the non controlled configuration the active power is generated by 15% from DGs and 85% from balance nodes, while the network absorbs about 0.45%. In the optimized configuration, the active power is generated by 17% from DGs and 83% from balance nodes, while the network absorbs about 0.44%. Although the reduction of power dissipation of the network is small, it can be seen an increase in active power generated by the DGs, so it can be concluded that the evolutionary optimization procedure tries to minimize the total losses by decreasing the power requests from the main grid (balance nodes), while increasing the power fraction produced by DGs which, in the considered network, are closest to loads. Consequently, the distribution paths on the network from power sources to loads are shortened, decreasing the total active power dissipated on electrical lines. Moreover, in the initial configuration the DGs deliver inductive reactive power, while, at the end of the simulation, they deliver capacitive reactive power; this is a further aspect that helps to reduce the power losses on the overall network.

5 Conclusions

In this paper it has been proposed a control system able to reduce power losses in the *ACEA Distribuzione S.p.A.* SG in the west area of Rome by modulating the DGs active and reactive powers, while considering suitable constraints on voltages and currents imposed by safety and quality of service issues. Moreover the constraints imposed by safe operational limits established by the Capability Curve of biogas generators have been considered. The network has been accurately modelled and simulated relying on the MatLab/Simulink SimPowerSystems ToolBox, which allows to rapidly and easily build models to simulate power systems. The optimization problem has been faced as a multi-objective one, since

power losses minimization and constraints satisfaction are conflicting objectives. Since it is not practically possible to derive in closed form the expression of the power losses in terms of system parameters, a derivative free optimization procedure based on a genetic algorithm has been adopted. First results show that the proposed control strategy is able to reduce power losses and to achieve admissible voltage and current levels according to predefined constraints. Future works will concern on evaluating different derivative free optimization techniques, such as Particle Swarm Optimization. Moreover, a Thyristor Voltage Regulator (TVR) and Li-Po Energy Storage System (ESS) will be soon installed in the considered electrical network. An advanced control system able to deal with these two new components is currently under study.

References

1. <http://www.aceaspa.it/section.aspx/en/history?lang=en>
2. http://www.autorita.energia.it/allegati/docs/10/039-10arg_ti.pdf
3. Caihao, L., Xianzhong, D.: Distribution generation and its impaction on power system. *Automation of Electric Power System* 6, 53–56 (2001)
4. Dahu, L.: Electric power communication system of the future smart grid. *Telecommunication for Electric Power System* 32(222) (2011)
5. Jinlan, C., Tianqi, L.: Distributed generation and its grid interconnection issue. *Modern Electric Power* 3(24), 53–57 (2007)
6. Lili, X., Xiaoyan, Q., Xiwen, W., Xingyuan, L.: Research of distribution generation planning in smart grid construction. In: *Power and Energy Engineering Conference (APPEEC), Asia-Pacific* (2010)
7. Qiang, S.: The communication system of smart grid. *Equipment Manufacturing Technology* (3) (2011)
8. Ramaswamy, P.C., Deconinck, G.: Relevance of voltage control, grid reconfiguration and adaptive protection in smart grids and genetic algorithm as an optimization tool in achieving their control objectives. In: *International Conference on Networking, Sensing and Control, Delft, The Netherlands* (2011)
9. Venayagamoorthy, G.K.: Potentials and promises of computational intelligence for smart grids. In: *IEEE Power General Society General Meeting, Calgary, AB, Canada*, pp. 1–6 (2009)
10. Werbos, P.J.: Computational intelligence for the smart gridhistory, challenges and opportunities. *IEEE Computational Intelligence Magazine* 6(3), 14–21 (2011)

Solar Irradiation Forecasting for PV Systems by Fully Tuned Minimal RBF Neural Networks

Lucio Ciabattoni¹, Gianluca Ippoliti¹, Sauro Longhi¹,
Matteo Pirro¹, and Matteo Cavalletti²

- ¹ Università Politecnica delle Marche, Dipartimento di Ingegneria dell'Informazione,
Via Breccie Bianche 12, 60131 Ancona, Italy
{l.ciabattoni,g.ippoliti,s.longhi,m.pirro}@univpm.it
² Energy Resources S.p.A.,
Via Ignazio Silone, 60035 Jesi, Italy
m.cavalletti@energyresources.it

Abstract. An on-line prediction algorithm able to estimate, over a determined time horizon, the solar irradiation of a specific site is considered. The learning algorithm is based on Radial Basis Function (RBF) networks and combines the growing criterion and the pruning strategy of the minimal resource allocating network technique. An adaptive extended Kalman filter is used to update all the parameters of the Neural Network (NN). The on-line learning mechanism avoids the initial training of the NN with a large data set. The proposed solution has been experimentally tested on a 14 kWp PhotoVoltaic (PV) plant and results are compared to a classical RBF neural network.

Keywords: irradiation forecasting, minimal resource allocating networks, adaptive filtering, self learning algorithm, neural networks.

1 Introduction

Recently, energy saving and energy security have become major issues, especially in some countries where energy deficiency not only impacts economics, society and development of the country, but also results in the global warming. For these reasons, interest in renewable energy is growing around the world and electric system operators are addressing the challenge of how to integrate significant amounts of wind, solar and other forms of variable generation into electricity grids while ensuring system reliability. This calls for transmission additions and reinforcements, enhanced forecasting and planning techniques for variable generation, and access to flexible grid resources including customer participation in demand management programs, plug-in hybrid electric vehicles and large scale electricity storage to help reliably integrate variable resources to electricity systems. In particular, there are two tasks to integrate variable generation and Distributed Energy Resources (DER), both locally and globally: integrating them into the electricity network and into the energy market. One solution to decrease the problems caused by the variable output of some distributed generation is to add energy

storages into the systems (centralised or distributed energy storages). Forecast information on the expected DER power production plays a primary role for the optimization and management of those storages. Another solution is to use flexibility in electricity consumption [17]. In fact, today time of use tariffs, for domestic use of energy, penalize some periods of time with a higher price. Prosumers (customers and producers of energy at the same time), knowing their forecasted energy production profile can (re)arrange their processes to minimize costs, having great economic benefits. Since both solutions need power production forecasts it's a crucial task the finding of a valid and reliable forecasting scheme. This problem has been deeply investigated in literature, in particular for what regards the forecasting and power scheduling from wind plant [15]. Also the prediction of solar yields is becoming more and more important, especially for countries where legislation encourages the deployment of solar power plants [16]. The irradiation forecasted can be used as input for a PhotoVoltaic (PV) model to obtain the PV production, or used in building energy simulation program that model energy and water use in buildings. Depending on the application and the corresponding time scale different approaches for modeling and forecasting solar irradiation may be appropriate and there have been a lot of researchers engaged in the modeling of solar irradiance. It is possible to use satellite-based cloud motion vectors for cloud fields and, therefore, surface irradiance short-term predictions with a forecast horizon of up to approximately 4 h or irradiance models (clear-day solar radiation, half-sine, Seasonal Auto Regressive Integrated Moving Average (SARIMA), WD-SARIMA, Support Vector Regression, see [9] and reference therein). However, those methods have not been efficient in all cases (most of them are efficient only to forecast up to 5 – 10 minutes [143]) and contrarily they may yield to noised results. A possible solution is given by Neural Networks (NNs) which provide a nonlinear representation to implement mappings. In particular in this paper Radial Basis Function Networks (RBFNs) have been considered for this prediction and the system dynamics related to the irradiation have been taken into account through the RBFN input pattern that must be composed of a proper set of system input and output samples acquired in a finite set of past time instants [11]. These NNs have been widely used for nonlinear system identification [4,6,8,25] because they have the ability both to approximate complex nonlinear mappings directly from input-output data with a simple topological structure that avoid lengthy calculations [6] and to reveal how learning proceeds in an explicit manner. The considered on-line learning algorithm is based on the Minimal Resource Allocating Network (MRAN) technique [22], that adds hidden neurons to the network based on the innovation of each new RBFN input pattern which arrives sequentially. As stated in [22], to obtain a more parsimonious network topology a pruning strategy is introduced. This strategy detects and removes as learning progresses those hidden neurons which make little contribution to the network output. Pruning is necessary for the prediction of the irradiation changing dynamics because inactive hidden neurons could be present as the dynamics which caused their creation becomes nonexistent. If an observation has no novelty then the existing parameters of the network are adjusted by an Extended Kalman Filter (EKF) [22]. In this paper the performance of the filter is improved by an on-line adjustment of the noise statistics obtained by a suitably defined estimation algorithm; the proposed Adaptive Extended Kalman Filter (AEKF) is able to adaptively estimating the unknown statistical

parameters [13]. The main advantage of the proposed MRAN algorithm is that a large data set of irradiation measurements, weather forecast, temperature for a specific location is no longer required for the training of the NN, drastically reducing the setup time. Another important advantage is that, due to the adaptive algorithm, some singular seasonal weather situation can be rapidly identified and corrected. A comparison of the performance obtained by the MRAN AEKF RBF Neural Network with respect to the standard RBF Neural Network is presented, considering data taken from a PV plant located in Jesi, Italy. The paper is organized as follows. The on-line prediction algorithm is described in Section 2 and the performance of the considered NNs are discussed in Section 3. The paper ends with comments on the performance of the proposed solution.

2 Prediction Algorithm

The approach to implement a Minimal Resource Allocating Network (MRAN) is based on a sequential learning algorithm and an Extended Kalman Filter (EKF) [22,10]. In particular the sequential learning algorithm adds or removes neurons on-line to the network according to a given criterion [22] and an EKF is used to update the net parameters. In this paper the MRAN algorithm is improved by an Adaptive Extended Kalman Filter (AEKF) in order to take into account the time-varying noise statistics [13], as shown in the following.

2.1 Radial Basis Function Neural Network

A RBFN with input pattern $\mathbf{x} \in \mathbb{R}^m$ and a scalar output $\hat{y} \in \mathbb{R}$ implements a mapping $f: \mathbb{R}^m \rightarrow \mathbb{R}$ according to

$$\hat{y} = f(\mathbf{x}) = \lambda_0 + \sum_{i=1}^K \lambda_i \phi(\|\mathbf{x} - \mathbf{c}_i\|) \quad (1)$$

where $\phi(\cdot)$ is a given function from \mathbb{R}^+ to \mathbb{R} , $\|\cdot\|$ denotes the Euclidean norm, λ_i , $i = 0, 1, \dots, K$ are the weight parameters, $\mathbf{c}_i \in \mathbb{R}^m$, $i = 1, 2, \dots, K$, are the radial basis function centers (called also units or neurons) and K is the number of centers [6]. The terms:

$$o_i = \lambda_i \phi(\|\mathbf{x} - \mathbf{c}_i\|), i = 1, \dots, K \quad (2)$$

are called the hidden unit outputs.

In this paper the RBFN is used for the prediction of the output of a dynamical system and the system dynamics can be taken into account through the network input pattern \mathbf{x} , that must be composed of a proper set of system input and output samples acquired in a finite set of past time instants [11]. In particular in this paper, the system output is the solar irradiation (see Section 3) and the inputs are the weather forecast, the number of day of the year, the hour of the day (see Section 3).

Theoretical investigation and practical results show that the choice of the non-linearity $\phi(\cdot)$, a function of the distance d_i between the current input \mathbf{x} and the centre \mathbf{c}_i , does

not significantly influence the performance of the RBFN [6]. Therefore, the following gaussian function is considered:

$$\phi(d_i) = \exp(-d_i^2/\beta_i^2), \quad i = 1, 2, \dots, K \tag{3}$$

where $d_i = \|\mathbf{x} - \mathbf{c}_i\|$ and the real constant β_i is a scaling or “width” parameter [6].

2.2 Minimal Resource Allocating Network Algorithm

The learning process of MRAN involves allocation of new hidden units and a pruning strategy as well as adaptation of network parameters [22]. The network starts with no hidden units and as input-output data $(\mathbf{x}(\cdot), y(\cdot))$ are received, some of them are used to generate new hidden units based on a suitably defined growth criteria. In particular at each time instant n the following three conditions are evaluated to decide if the input $\mathbf{x}(n)$ should give rise to a new hidden unit:

$$\|e(n)\| = \|y(n) - f(\mathbf{x}(n))\| > E_1 \tag{4}$$

$$e_{rms}(n) = \sqrt{\sum_{j=n-(M-1)}^n \frac{e(j)^2}{M}} > E_2 \tag{5}$$

$$d(n) = \|\mathbf{x}(n) - \mathbf{c}_r(n)\| > E_3 \tag{6}$$

where $\mathbf{c}_r(n)$ is the centre of the hidden unit that is nearest to $\mathbf{x}(n)$ and M represents the number of past network outputs to calculate the output error $e_{rms}(n)$. The terms E_1 , E_2 and E_3 are thresholds to be suitably selected. As stated in [22], these three conditions evaluate the novelty in the data. If all the criteria of (4)–(6) are satisfied, a new hidden unit is added and the following parameters are associated with it:

$$\lambda_{K+1} = e(n) \tag{7}$$

$$\mathbf{c}_{K+1} = \mathbf{x}(n) \tag{8}$$

$$\beta_{K+1} = \alpha \|\mathbf{x}(n) - \mathbf{c}_r(n)\| \tag{9}$$

where α determines the overlap of the response of a hidden unit in the input space as specified in [22]. If the observation $(\mathbf{x}(n), y(n))$ does not satisfy the criteria of (4)–(6), an EKF is used to update the following parameters of the network:

$$\mathbf{w} = [\lambda_0, \lambda_1, \mathbf{c}_1^T, \beta_1, \dots, \lambda_N, \mathbf{c}_N^T, \beta_N]^T. \tag{10}$$

The update equation is given by:

$$\mathbf{w}(n) = \mathbf{w}(n-1) + \mathbf{k}(n)e(n) \tag{11}$$

where the gain vector $\mathbf{k}(n)$ is expressed by:

$$\mathbf{k}(n) = \mathbf{P}(n-1)\mathbf{a}(n) [\sigma_v^2(n) + \mathbf{a}^T(n)\mathbf{P}(n-1)\mathbf{a}(n)]^{-1} \tag{12}$$

with $\mathbf{a}(n)$ the gradient vector of the function $f(\mathbf{x}(n))$ (see Eq. (1)) with respect to the parameter vector $\mathbf{w}(n-1)$ [22], $\sigma_v^2(n)$ is the variance of the measurement noise and $\mathbf{P}(n-1)$ is the error covariance matrix. This matrix is updated by:

$$\mathbf{P}(n) = [\mathbf{I}_{z \times z} - \mathbf{k}(n)\mathbf{a}^T(n)] \mathbf{P}(n-1) + \sigma_\eta^2(n-1)\mathbf{I}_{z \times z} \quad (13)$$

where \mathbf{I} is the identity matrix and $\sigma_\eta^2(n-1)$ is introduced to avoid that the rapid convergence of the EKF algorithm prevents the model from adapting to future data [22]. The $z \times z$ matrix $\mathbf{P}(n)$ is positive definite symmetric and z is the number of parameters to be adjusted. When a new hidden neuron is allocated, the dimension of $\mathbf{P}(n)$ increases to:

$$\mathbf{P}(n) = \begin{pmatrix} \mathbf{P}(n-1) & 0 \\ 0 & p_0 \mathbf{I}_{z_1 \times z_1} \end{pmatrix}. \quad (14)$$

In (14), p_0 is an estimate of the uncertainty in the initial values assigned to the parameters and the dimension z_1 of the identity matrix \mathbf{I} is the number of new parameters introduced by adding the new hidden neuron. As stated in [22], to keep the RBF network in a minimal size a pruning strategy removes those hidden units that contribute little to the overall network output over a number of consecutive observations. To carry out this pruning strategy, for every observation $(\mathbf{x}(n), y(n))$ the hidden unit outputs are computed (see Eq. (2)) and normalized with respect to the highest output:

$$\bar{o}_i(n) = \frac{o_i(n)}{\max\{o_i(n)\}}, i = 1, \dots, K. \quad (15)$$

The hidden units for which the normalized output (15) is less than a threshold δ for ξ consecutive observations are removed and the dimensionality of all the related matrices are adjusted to suit the reduced network [22].

The weakness of the above algorithm is that all the parameters of the network, including all the centers of the hidden neurons, widths and weights, have to be update at every step; the size of the matrices to be update becomes large as the number of hidden neurons increases. Therefore, for the real-time implementation of the considered algorithm, it could be necessary to reduce the online computation effort and to this purpose a “winner neuron” strategy can be incorporate in the learning algorithm as proposed in [22]. The “winner neuron” is defined as the neuron in the network that is closest (in some norm sense) to the current input data. The criteria for adding and pruning the hidden neurons are all the same as in the above algorithm; the difference, in the “winner neuron” strategy, is that if the observation $(\mathbf{x}(\cdot), y(\cdot))$ does not meet the criteria to add a new hidden neuron (see Eqs. (4)–(6)), only the network parameters related to the selected “winner neuron” are updated by the EKF algorithm [22].

The EKF can be implemented once estimates of $\sigma_\eta^2(n)$ and $\sigma_v^2(n)$ are available. In general, a complete and reliable information about these estimates is not available; on the other hand it is well known how poor knowledge of noise statistics may seriously degrade the Kalman filter performance. This problem is here dealt with introducing an adaptive adjustment mechanism of $\sigma_\eta^2(n)$ and $\sigma_v^2(n)$ values in the EKF equations.

2.3 Adaptive Estimation of $\sigma_\eta^2(n)$ and $\sigma_v^2(n)$

A considerable amount of research has been performed on the adaptive Kalman filtering [13][23], but in practice it is often necessary to redesign the adaptive filtering scheme according to the particular characteristics of the faced problem. In the adaptive procedure here proposed, under proper assumptions given in [13], it is possible to define a simple and efficient estimation algorithm based on the condition of consistency, at each step, between the residual $e(n)$ and its predicted statistic $E\{e^2(n)\}$. Imposing such a condition, one-stage estimates $\hat{\sigma}_\eta^2(n-1)$ and $\hat{\sigma}_v^2(n)$, of $\sigma_\eta^2(n-1)$ and $\sigma_v^2(n)$, respectively, are obtained at each step. To increase their significance, the one-stage estimates $\hat{\sigma}_\eta^2(n-1)$ and $\hat{\sigma}_v^2(n)$ are average obtaining the relative smoothed version $\bar{\sigma}_\eta^2(n-1)$ and $\bar{\sigma}_v^2(n)$. After proper calculations [13], the following recursive form of estimates $\bar{\sigma}_\eta^2(n-1)$ and $\bar{\sigma}_v^2(n)$ is found:

$$\bar{\sigma}_\eta^2(n-1) = \bar{\sigma}_\eta^2(n-2) + \frac{1}{l_\eta + 1} [\hat{\sigma}_\eta^2(n-2) - \hat{\sigma}_\eta^2(n - (l_\eta + 1))] \tag{16}$$

$$\bar{\sigma}_v^2(n) = \bar{\sigma}_v^2(n-1) + \frac{1}{l_v + 1} (\hat{\sigma}_v^2(n) - \hat{\sigma}_v^2(n - l_v)) \tag{17}$$

where:

- $\hat{\sigma}_\eta^2(n-1) = \max\{(\mathbf{a}(n)\mathbf{a}^T(n))^{-1}[e(n)^2 - \mathbf{a}(n)\mathbf{P}(n-1)\mathbf{a}^T(n) - \bar{\sigma}_v^2(n)], 0\}$
- $\hat{\sigma}_v^2(n) = \max\{e^2(n) - [\mathbf{a}(n)\mathbf{P}(n-1)\mathbf{a}^T(n) + \mathbf{a}(n)\bar{\sigma}_\eta^2(n-1)\mathbf{I}\mathbf{a}^T(n)], 0\}$
- l_η and l_v are the number of one-stage estimates $\bar{\sigma}_\eta^2(n-1)$ and $\bar{\sigma}_v^2(n)$ respectively, yielding the smoothed estimates.

Parameters l_η and l_v of estimators (16) and (17) are chosen on the basis of two antagonist considerations: low values would produce noise estimators which are not statistically significant, large values would produce estimators which are scarcely sensitive to possible rapid fluctuations of the true $\sigma_\eta^2(n-1)$ and $\sigma_v^2(n)$ [13]. In other words, the one-stage estimates are made by averaging past samples in order to increase the statistical significance of estimators; if the samples are too far the filter has a low reactivity, while if the samples are too near estimators have a low statistical significance [13]. During filter initialization, the starting values $\hat{\sigma}_\eta^2(0)$ and $\hat{\sigma}_v^2(0)$ in Eqs. (16) and (17) respectively, must be chosen on the basis of the *a priori* available information. In case of lack of such information, a large value of $P(0,0)$ is useful to prevent divergence.

The MRAN prediction algorithm [22] enhanced by the AEKF, called MRANAIEKF algorithm, is summarized as follow:

1. For each observation $(\mathbf{x}(n), y(n))$ do: compute the overall network output: $\hat{y}(n) = f(\mathbf{x}(n)) = \lambda_0 + \sum_{i=1}^K \lambda_i \phi(\|\mathbf{x}(n) - \mathbf{c}_i\|)$ where K is the number of hidden units;
2. Calculate the parameters required by the growth criterion:
 - $\|e(n)\| = \|y(n) - f(\mathbf{x}(n))\|$
 - $e_{rms}(n) = \sqrt{\sum_{j=n-(M-1)}^n \frac{e(j)^2}{M}}$

- $d(n) = \|\mathbf{x}(n) - \mathbf{c}_r(n)\|$
3. Apply the criterion for adding a new hidden unit:
- if**
- $\|e(n)\| > E_1$ **and** $e_{rms}(n) > E_2$ **and** $d(n) > E_3$ allocate a new hidden unit $K + 1$ with:
- $\lambda_{K+1} = e(n)$
 - $\mathbf{c}_{K+1} = \mathbf{x}(n)$
 - $\beta_{K+1} = \alpha \|\mathbf{x}(n) - \mathbf{c}_r(n)\|$
- else**
- adapt the measurement noise coefficient:
 $\hat{\sigma}_v^2(n) = \hat{\sigma}_v^2(n-1) + \frac{1}{l_v+1} (\hat{\sigma}_v^2(n) - \hat{\sigma}_v^2(n-l_v))$
 - tune the network parameters:
 $\mathbf{w}(n) = \mathbf{w}(n-1) + \mathbf{k}(n)e(n)$
 - adapt the process noise coefficient:
 $\hat{\sigma}_\eta^2(n-1) = \hat{\sigma}_\eta^2(n-2) + \frac{1}{l_\eta+1} [\hat{\sigma}_\eta^2(n-2) - \hat{\sigma}_\eta^2(n-(l_\eta+1))]$
 - update the error covariance matrix:
 $\mathbf{P}(n) = [\mathbf{I}_{z \times z} - \mathbf{k}(n)\mathbf{a}^T(n)] \mathbf{P}(n-1) + \sigma_\eta^2(n-1)\mathbf{I}_{z \times z}$
- end**
4. Check the criterion to prune hidden units:
- compute the hidden unit outputs:
 $o_i(n) = \lambda_i \phi(\|\mathbf{x}(n) - \mathbf{c}_i\|), i = 1, \dots, K$
 - compute the normalized outputs:
 $\bar{o}_i(n) = \frac{o_i(n)}{\max\{o_i(n)\}}, i = 1, \dots, K$
 - **if** $\bar{o}_i(\cdot) < \delta$ for ξ consecutive observations **then**
 prune the i th hidden unit and reduce the dimensionality of the related matrices
- end**
5. $n = n + 1$ and **go** to step 1.

3 Neural Network Based Irradiation Forecasting

Tests are based on data acquired from January 2011 to December 2011 during PV plant standard working. The considered 14 kWp PV plant, equipped with polysilicon solar panels south oriented and tilt angle 27 deg., is located in Jesi (AN), Italy. It is composed by 8 strings of Renegies 220P/220 polysilicon panels [19] where each string is connected to a SMA Sunny Boy 1700IT solar inverter [21]. A lithium battery pack which is composed by the series of two sub-module with 80 ThunderSky modules 40 Ah, a Battery Management System (BMS) and a battery charger for each module [24]. A solar inverter (model SIAC soleil 10kW) is connected to this pack [20]. All communication is done through the TCP/IP protocol, using serial to TCP/IP converters. All devices are connected to a server, where it is located the software to manage the whole system.

For the above PV plant, two different situations are considered: in the first one the proposed algorithm is tested without a pre-trained net; this situation can occur if the solar irradiation has to be predicted on a plant without previous information on it. In the second situation, the MRANAELF learning algorithm starts with a pre-trained net




Fig. 1. Evolution of hidden neurons due to growing and pruning for the considered data set (Year 2011). The continuous red line is the pre-trained MRANAELKF network while the dashed blue line is the no pre-trained MRANAELKF network. The first data window is from day 7 (144th sample) to day 9 (216th sample); the second data window is from day 169 (4032th sample) to day 171 (4106th sample).

based on few historical information found on the WEB such as solar irradiation profile of clear sky days and cloudy days for the specified location, panel orientation and tilting. The information have been taken directly from the website of PVGIS [18]. This is a common operating condition, when no sensors and measures are available for the plant before the forecast starts. To measure the performance of the proposed algorithm, the Root Mean Square of the Error $e(\cdot)$ (RMSE) and its Standard Deviation (SD) have been calculated. Only hours with daylight (irradiance greater than zero) are considered for the calculation of the RMSE; night values with no irradiance are excluded from the evaluation. The set of experimental data is composed of 8000 pairs of input and output samples. Sampling time is 1 h and the data have been normalized, between 0 and 1, in order to have the same range. In particular the set of experimental data is given by the pairs $(x(n), y(n))$, $n = 1, 2, \dots$ where $x(n)$ is composed by the hourly weather forecast given by sunny, clean, partly cloud and overcast sky conditions taken on the WEB [12, 11], the number of day of the year, the hour of the day, the air temperature and $y(\cdot)$ is the solar irradiation measured by the SMA Sunny Sensorbox [21]. The considered algorithm requires careful selection of the threshold parameters E_1, E_2, E_3 and of parameter M , as defined in Eqs. (4)–(6), which control the growth characteristics of the network; i.e. if small thresholds are chosen more units are added to the NN. The parameters δ and ξ control the pruning strategy (Eq. (15)); it is important to take into account the system non stationarity to select these parameters. In other words, slowly dynamic variations imply a bigger δ and a smaller ξ . The parameters $\alpha, p_0, \sigma_v^2(0)$ and $\sigma_\eta^2(0)$ related to the AEKF algorithm used to update the network parameters of Eq. (10) are chosen by trial and error. In the considered experimental tests the numeric values of these parameters are selected, as: $E_1 = 0.01, E_2 = 0.02, E_3 = 0.4, M = 50, \delta = 0.0005, \xi = 2000, \alpha = 1.2, p_0 = 0.2, \sigma_v^2(0) = 0.03, \sigma_\eta^2(0) = 0.03$. Samples of the performed prediction tests are given in Fig. 1 through 4. In particular the Fig. 1




Fig. 2. First data window: day 7 to 9 of 2011 (7–9 January)




Fig. 3. Second data window: day 169 to 171 of 2011 (19–21 June)

shows the hidden neurons evolution history for the MRANAEOF algorithm as it learns sequentially from data. In this figure two data windows have been highlighted to compare the performance of the MRANAEOF RBF NN with respect to a classical RBF NN algorithm. As will be showed in Figs. 2 and 3, the first data window is relative to the starting learning period of the MRANAEOF (days 7 – 9) and the second data window is relative to the “steady state” operation (days 169 – 171). Results are compared with a previously trained classical RBF network with the same inputs and output. The training data set of the classical RBF network is composed by 5000 pairs of input-output and it is relative to the last 20 days of each month of 2010 as described in [7]. A sample of the performed tests is shown in Figs. 2 and 3, where continuous red line represents the irradiation predicted by the pre-trained MRANAEOF, the dashed blue line is the irradiation predicted by the no pre-trained MRANAEOF network, the dotted black line is the irradiation predicted by the classical RBF NN and the continuous green line is the measured irradiation. Results for the data set relative to the year 2011 have been




Fig. 4. Sample covariance of the residuals obtained by the prediction performed by the pre-trained MRANAIEKF network.

Table 1. Comparison of prediction results (values are in $[W/m^2]$) of the MRANAIEKF with (a) and without (b) pre-training and the classical fixed RBF network. The considered whole data set is relative to the year 2011. Data windows are highlighted in Fig. 1.

DATA	MRANAIEKF (a)		MRANAIEKF (b)		RBF	
	RMSE	SD	RMSE	SD	RMSE	SD
Whole data set	65.2	60.2	70.5	68.2	75.1	74.2
First data window	71.1	69.8	76.2	75.2	81.2	79.7
Second data window	50.4	48.3	55.9	53.4	59.3	57.1

summarized in Table 1. The network pre-trained with basic information on PV plant design and historical data available on the WEB [18] has shown good capability to generalize the prediction of the irradiation of a specific plant in the first two days of working and, with the growing on the net, the forecast becomes more accurate; the net starting without neurons (no pre-trained) needs up to 4 days to show a good performance. Both networks have shown better performance with respect to the classical RBF NN. The whiteness test on the sample covariance of the prediction errors has been used for the network validation. A sample is shown in Fig. 4.

4 Concluding Remarks

In this paper a minimal resource allocating network algorithm has been analyzed for on-line hourly site-specific irradiance forecast. This algorithm increases the number of RBFN hidden neurons depending on the input-output data and an adaptive extended Kalman filter is used to update all the parameters of the network. A pruning strategy is also considered to remove those hidden units which end up to give a contribution to the network output. Two ways of use the MRANAIEKF algorithm have been proposed:

with a pre-trained net based only on historical information found on the WEB and with a no pre-trained network. The results indicate that both situations give networks with better prediction accuracy with respect to a classical fixed RBF NN previously trained with a large data set. As future developments, the authors are actually considering the possibility to integrate a “winner neuron” strategy to minimize the computational effort of the overall algorithmic architecture. This will make the developed system even more appealing from a real-time implementation perspective.

References

1. <http://www.3bmeteo.com> (2012)
2. Antonini, P., Ippoliti, G., Longhi, S.: Learning control of mobile robots using a multiprocessor system. *Control Engineering Practice* 14(11), 1279–1295 (2006)
3. Ben Salah, C., Ben Mabrouk, A., Ouali, M.: Wavelet autoregressive forecasting of climatic parameters for photovoltaic systems. In: 2011 8th International Multi-Conference on Systems, Signals and Devices (SSD), pp. 1–6 (March 2011)
4. Cavalletti, M., Ippoliti, G., Longhi, S.: Lyapunov-based switching control using neural networks for a remotely operated vehicle. *International Journal of Control* 80(7), 1077–1091 (2007)
5. Cavalletti, M., Ippoliti, G., Longhi, S.: Intelligent control for a remotely operated vehicle. *International Journal of Systems Science* 40(11), 1099–1114 (2009)
6. Chen, S., Cowan, C., Grant, P.: Orthogonal least squares learning algorithm for radial basis function networks. *IEEE Transactions on Neural Networks* 2(2), 302–309 (1991)
7. Ciabattini, L., Ippoliti, G., Longhi, S., Cavalletti, M., Rocchetti, M.: Solar irradiation forecasting using RBF networks for PV systems with storage. In: 2012 IEEE International Conference on Industrial Technology (ICIT), pp. 699–704 (March 2012)
8. Corradini, M., Ippoliti, G., Longhi, S., Marchei, D., Orlando, G.: A quasi-sliding mode observer-based controller for PMSM drives. *Asian Journal of Control*, 1–8 (2012), <http://dx.doi.org/10.1002/asjc.555>
9. Dorvlo, A.S., Jervase, J.A., Al-Lawati, A.: Solar radiation estimation using artificial neural networks. *Applied Energy* 71(4), 307–319 (2002)
10. Giantomassi, A., Ippoliti, G., Longhi, S., Bertini, I., Pizzuti, S.: On-line steam production prediction for a municipal solid waste incinerator by fully tuned minimal RBF neural networks. *Journal of Process Control* 21(1), 164–172 (2011)
11. Hunt, K., Sbarbaro, D., Zbikowski, R., Gawthrop, P.: Neural networks for control systems – A survey. *Automatica* 28(6), 1083–1112 (1992)
12. <http://www.ilmeteo.it> (2012)
13. Jetto, L., Longhi, S., Venturini, G.: Development and experimental validation of an adaptive extended Kalman filter for the localization of mobile robots. *IEEE Transactions on Robotics and Automation* 15(2), 219–229 (1999)
14. Karner, O.: ARIMA representation for daily solar irradiance and surface air temperature time series. *Journal of Atmospheric and Solar-Terrestrial Physics* 71(8-9), 841–847 (2009)
15. Lange, M., Focken, U.: *Physical approach to short term wind power prediction*. Springer, New-York (2006)
16. Lorenz, E., Hurka, J., Heinemann, D., Beyer, H.: Irradiance forecasting for the power prediction of grid-connected photovoltaic systems. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 2(1), 2–10 (2009)
17. Palensky, P., Dietrich, D.: Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics* 7(3), 381–388 (2011)

18. <http://re.jrc.ec.europa.eu/pvgis> (2012)
19. <http://www.renergiesitalia.it> (2012)
20. <http://www.sielups.com> (2012)
21. <http://www.sma-italia.com> (2012)
22. Sundararajan, N., Saratchandran, P., Li, Y.: Fully tuned radial basis function neural networks for flight control. Kluwer Academic, London (2002)
23. Szabat, K., Orłowska-Kowalska, T.: Performance improvement of industrial drives with mechanical elasticity using nonlinear adaptive Kalman filter. *IEEE Transactions on Industrial Electronics* 55(3), 1075–1084 (2008)
24. <http://en.winston-battery.com> (2011)

Ontology-Based Device Configuration and Management for Smart Homes

Michele Nucci, Marco Grassi, and Francesco Piazza

D.I.I. Department of Information Engineering
Polytechnic University of Marche, Ancona, Italy
{m.nucci,m.grassi,f.piazza}@univpm.it
<http://www.semedia.dii.univpm.it>

Abstract. Energy saving is nowadays a mandatory requirement for buildings. In this paper, a novel holistic system for intelligent Smart Home environments is introduced, which is able to embrace energy production and consumption in a unique perspective and to handle them in conjunction with device and services management. Semantic Web technologies are employed to foster data interoperability and to supply inference power for intelligent tasks management, decisions making and energy saving. The proposed system uses an IP-based network as main communication channel and a semantic extension of the UPnP technology, to support zero-configuration, automatic discovering and intelligent control. The ontology framework, used to formally describe all relevant information of the Smart Home environment like devices, services and context, is also presented focusing on the device and the energy ontology. In addition, the issue of the lack of Semantic Web compliant device descriptions that currently are not provided by device vendors, posing a serious barrier toward the practical application of semantic technologies in Smart Home scenario, is also tackled.

1 Introduction

The multiplication of smart devices available for the home environment, able to accomplish to some extent autonomously operations of configuration and communication, is pushing forward the intelligent capabilities of houses. These can provide more comfortable living spaces to their inhabitants, resulting more adaptive and reactive to user needs, and more efficient behaviors, allowing the automatic accomplishment of several tasks like maintenance and energy management, as envisioned by Ambient Intelligent (AmI) paradigm [1]. To take full advantage of the potential of such smart devices, it's crucial to provide them the capability to self-describe and automatically communicate. In particular, providing machine processable and understandable descriptions of their capabilities, it is a key factor for the implementation of more efficient automatic control strategies. This not only to provide a better user comfort but also to improve energetic

efficiency, which in recent years has become a mandatory buildings requirement for both environment preservation and economic reasons [2].

In previous works [3] [4], we introduced a holistic vision for intelligent Smart Home (SH) task management, based on Semantic Web (SW) technologies to provide data interoperability among the numerous devices that compose and operate in a SH environment and to supply inference power for intelligent tasks management and energy saving. An ontology framework is used to model several aspects like: operative scenario, context, quality of service, user preferences and energetic production and consumption in a unique global Knowledge Base (gKB) used to support the implementation of efficient control logics.

In this paper, we move towards the practical implementation of our vision in the real world scenario. In particular, we discuss a semantic extension of the UPnP protocol used to foster automatic communications capabilities among devices. In addition, the issue of the lack of SW compliant device descriptions that currently are not provided by device vendors, posing a serious barrier toward the practical application of SW technologies in SH scenario, is also tackled.

The paper is organized as follows: Section 2 provides a brief state of the art about devices description in AmI systems. Section 3 summarizes our holistic vision for SH environments, describing also the developed ontology framework focusing on the Energy and Device ontologies. Section 4 presents the proposed semantic extension of UPnP protocol. Section 5 discusses a Device Description Repository (DDR) used to manage device descriptions and to promote their sharing and reuse, basing on a social model for collaborative description creation and exchange.

2 Device Description

AmI's smart devices have to be able to communicate and to automate their configuration and discovery tasks. Several proprietary protocols have been developed in recent years to enable the communication among smart home devices such as LonWorks [5], Z-Wave [6] and KNX [7]. Being proprietary, these protocols limit the communication capabilities with devices that are based on different communication buses (e.g. communication among devices from different manufacturers). The multiplication of the number of heterogeneous devices in smart homes, often produced by different manufacturers, has raised the need of using more general communication buses, pushing forward the implementation of IP-based solutions. A growing number of manufacturers has started to support opened IP-based protocols in their devices or, at least, to provide gateways to connect their proprietary networks with the IP-world.

Currently, UPnP [8] and other standards based on it like UPnP/AV [9] and DLNA [10], seem to be the most widely supported standards for devices description, configuration and discovering. These standards use the Extensible Markup Language (XML) for devices and services description and they classify devices using a limited set of predefined categories, each of those defines predefined services. However, the descriptive capabilities of this approach are not sufficient to

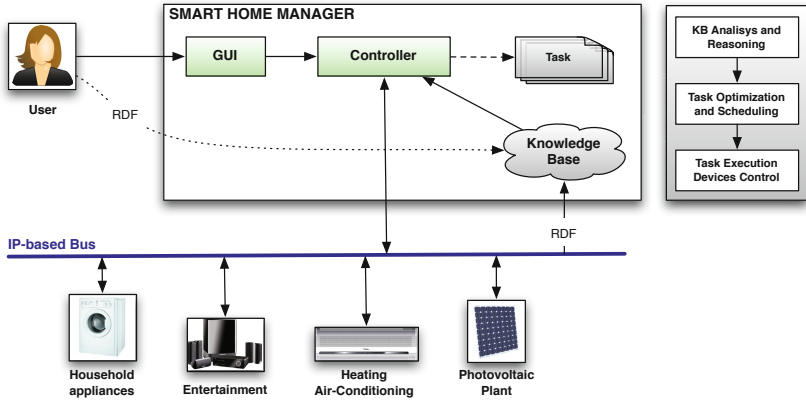


Fig. 1. A simplified sketch of system architecture

model the large number of devices currently available on the market as well as to encode the large number of information required to model the entire domain of an AmI (e.g. context, user preferences, etc.), as required to create a dynamic knowledge base and a reasoning system. Moreover, a description model based on XML lacks of formal semantics and provides only a strict syntactic representation of information based on preexisting schema. This limits the intelligent and adaptive capabilities of an AmI, considering that all the operations performed to search, discover or match different devices are based only on syntactic comparisons, without considering logical equivalences.

The importance that device description plays also in other scenarios like device design, system engineering and integration has led to several efforts in defining standards for device description: FIPA [11], CCPP [12], GSDML and FDCML [13], and CANopen [14]. However, the possibility to introduce a commonly accepted standard seems doomed to remain utopia for technical, time-related and political reasons. In order to overcome the lack of standardization, the limits of the current proposed solutions and to provide more expressive, flexible and effective descriptions, several proposals, like [15], [16] and [17], use Semantic Web technologies to develop generic architectures for devices and services description.

3 A Holistic Vision for Smart Home Environments

In this section we briefly describe the implementation of the proposed holistic vision for SH environment, Figure 1. More details can be founded in [3]. An IP-based network is used as main communication channel for the whole SH environment and a semantic extension of the UPnP protocol, developed in conformity with the specifications defined by the standard itself, is used for devices auto-configuration, discovering and controlling. The extended protocol, which is presented in details in section 4, enables intelligent management of devices and services, which have been semantically described using RDF. In order to design

a system to connect legacy devices (devices that do not support natively the IP connectivity and the UPnP technology) to the proposed SH environment, different classes of devices have been considered. Programmable devices can be easily extended to support the proposed implementation of UPnP protocol. For not programmable devices as well as for devices that do not support natively the IP connectivity dedicated low-cost gateways are used. These gateways allow to wrap one or more physical devices providing IP connectivity, offering transparent interfaces towards proprietary communication buses, and a shared implementation of the semantically extended UPnP protocol. The use of gateways ensures device independency in control logics allowing the implementation of high-level control algorithms. In this way, damaged physical devices can be substituted and obsolete devices can be replaced with new generation devices without requiring any intervention on control logic, since the semantic information required for their control are separated from the control application itself. Also, virtual devices can be implemented to emulate real devices capabilities (e.g. a virtual weather station can be created exploiting online weather forecast services and be later substituted by a sensor without requiring any further intervention).

A centralized Smart Home Manager (SHM) is used both for devices management (handling their communication, checking their status, sending controls, etc.), for automatic service administration (service composition, task scheduling, etc.) and for user interaction. SHM is responsible of retrieving and managing the semantic device descriptions, as discussed more in details in Sec. 5. Also, when a user interacts with any device placed in the SH environment, the SHM logs the user actions, trying to automatically profile the user preferences. In additions, users can directly interact with the SHM to enter their preferences, to ask the execution of a specific task, to schedule a complex set of actions, etc. All information obtained by devices or derived from user interaction are encoded in RDF and stored in the gKB.

In the proposed system, an ontology framework is used as modeling structure to encode all the relevant information regarding manifold aspects like operative scenario, context, quality of service, user preferences, energetic production and consumption in a unique gKB, as described in more details in a precedent work [4]. A wide set of different ontologies is introduced, each of those self-sufficient and specifically conceived to describe one aspect of the system but also strictly bounded with the others. The *Service Ontology* allows to describe and define a wide number of services, ranging from communication and security to domotic and entertainment, extending the limited number of classes defined by the UPnP standard and encoding several information like service type, services priority, required energy, etc. Several different facets of system context can be described using specific ontologies like the *User Ontologies*, for users and their preferences, and the *Energy Ontology*, for encoding information about energy actual and forecast status.

Currently the ontology framework is under development; however ontologies for device and energy description are already in an *alpha* status and can be used in conjunction to support intelligent energy management, as show in Figure 2.

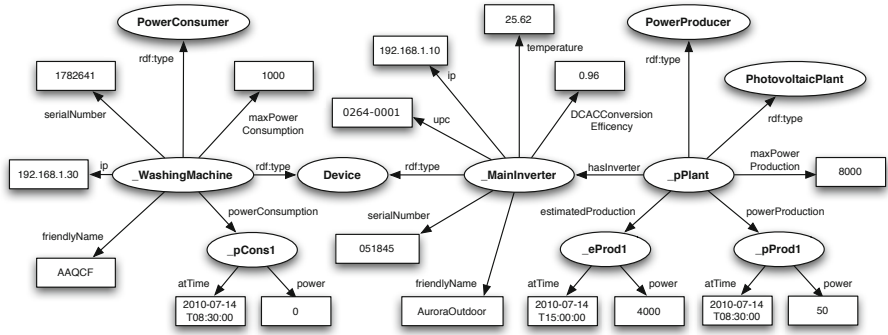


Fig. 2. An example of information encoding for intelligent task scheduling

Exploiting the inference capabilities provided by the used rich semantic descriptions, the gKB can be used to support the implementation of efficient control logics and to enhance system intelligence in task management. Energy saving can therefore be implemented not only at device level, simply using low consumption devices, but also at task level, through an intelligent task scheduling based on energy optimization strategy.

4 Semantic UPnP Extension

UPnP is a technology for pervasive peer-to-peer network connectivity of intelligent devices, designed to be flexible and to support zero-configuration, automatic discovering and controlling. UPnP devices can join a network, automatically advertise their presence and services and seamlessly discover the presence of other devices. UPnP devices advertise their presence multicasting a number of discovery messages. Such messages, which are based on part of the header field format of HTTP 1.1, contain a standardized set of headers, in particular, a *Location* header at which the device description and the description of the provided services can be retrieved. Control Points, which are special UPnP devices that are able to control other devices available on the network, can listen to the standard multicast address, for example, to get informed that new devices and services are available. When a new UPnP device is identified, the Control Point can process its advertising message to retrieve its XML description that can be used to control it.

To support RDF descriptions of UPnP devices, in the proposed extension of the UPnP technology the advertising message has been modified introducing additional headers as shown in Figure 3: *RDF-Location* and *Ontology*. The *RDF-Location* header contains the URI of the device’s RDF description. Such description is commonly stored internally to the device but can also be stored on an external repository. In the latter case, the description is required to be always accessible to avoid malfunctioning but can be more easily automatically updated,

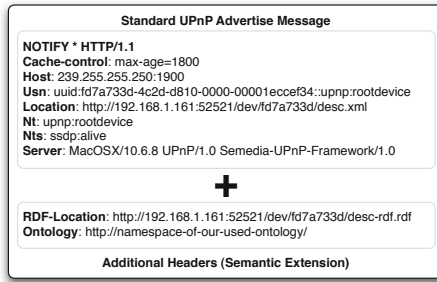


Fig. 3. An example of UPnP extended message

for example, to include additional information. The *Ontology* header contains a URI (or an array of URIs) pointing to the location at which the ontologies used to structure the RDF description are stored. Once retrieved, semantic descriptions are then stored in the gKB to be used in supporting intelligent systems control logics.

Adding new custom headers in the advertising message does not break the UPnP standard. Several tests have been done to identify possible problems due to the addition of new custom headers within the advertising message. These tests have been executed implementing software UPnP devices using Cling¹, CyberLinkC² and Platinum³ frameworks. In our tests, no problems have been identified. Traditional UPnP Control Points were able to correctly access to the standard XML description while intelligent Control Points were able to access the RDF description and the related ontologies.

5 Device Descriptions Repository

Currently, one of the main issues in fostering the application of SW technologies for device management is that manufacturers do not provide digital device descriptions using a semantic compliant format and often not even in electronic format. At best, smart device descriptions are supplied in XML, as in the case of UPnP and related standards, which can be converted in RDF using XSLT (EX-tensible Stylesheet Language Transformation). This requires the XML schema used in the description to be known and to define a mapping between the XML schema and the used ontology. Once that an XSLT is defined the conversion can be performed automatically for all the devices using the same XML schema. Commonly, most of the devices information is available only in human-readable format and provided in use and maintenance manuals or data-sheets. In addition, supplementary information has to be added with respect to the one provided in XML. This requires a preliminary activity to generate the required semantic

¹ <http://tealeal.org/projects/cling/>

² <http://www.cybergarage.org/twiki/bin/view/Main/CyberLinkForC>

³ <http://sourceforge.net/projects/platinum/>

device descriptions that is highly expensive and time consuming, in particular if it has to be performed manually and replied for each device and installation.

To ease the generation, the sharing and the retrieval of devices description, a centralized Device Description Repository (DDR) has been introduced. DDR provides tools for creating and adding new semantic device descriptions that are not yet available in the repository. This collaborative approach in data creation and sharing, which has been applied with success in several scenarios since the advent of Web 2.0, is expected to foster the availability of device descriptions. DDR has been conceived for being agnostic to the ontologies used for semantic descriptions, exploiting a general purpose RDF editor that works over a generic SPARQL endpoint. Moreover, DDR has been designed to be implemented as an open repository according to the Linked Data guidelines [18] to make SW compliant device descriptions accessible over the Web. In this way, available descriptions can be easily reused and consumed by other users and third party applications. DDR provides advanced and intuitive GUIs for devices descriptions searching and visualization based on the faceted browsing paradigm. This allows users to explore such information both using a search box, to perform keyword-based queries, and filtering the results using the facet menus (i.e. by adding or removing constraints on the facet properties). The repository can be used also for searching and comparing semantically enable devices, according for their type, performances and energetic consumption and can be also used to ease a system installation or configuration. The availability of semantic device descriptions can therefore become a commercial advantage and push vendors to publish themselves the descriptions in RDF format. DDR exposes restful APIs that can be used by software agents for device descriptions searching, authoring and consuming. In particular, the APIs allow to search for existing device descriptions and to create or edit new or existing ones according to the proposed ontologies, described in Section 3.

Figure 4 shows how the DDR can be used in the proposed system in supporting automatic devices management and control. When a semantically enabled device accesses the network, the SHM can retrieve its RDF description and store it in its local gKB. On the contrary, if the connected device supports only the standard UPnP protocol, the SHM can parse the information contained in the

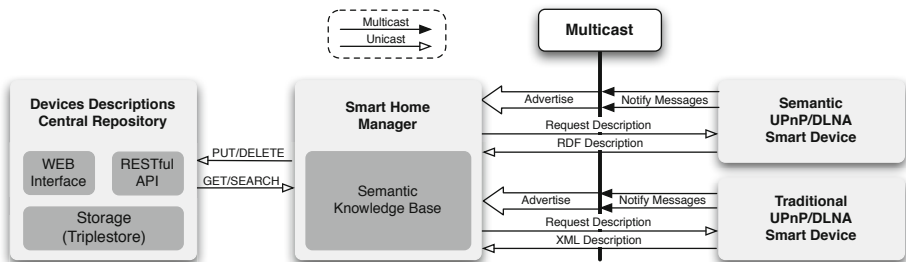


Fig. 4. Semantic compliant device description retrieval process

XML description and use them to query the DDR to look for an already existing semantic description for such device. Searching operation are based on device type, model and vendor and on some experimental heuristics. In case the search produces multiple results, user intervention is required to select the description to be used. If the searching operation did not produce any valid results, the SHM send the XML description retrieved from the local device to DDR for converting it to RDF using an XSLT transformation. In the implementation of the proposed system, some experimental XSLT transformations have been developed and tested to cover the main classes of devices defined by the UPnP Forum. More information is nevertheless required to provide a more intelligent device management, for example about, its energy consumptions. With such purpose, the SHM can require the user/installer to manually introduce the additional required information, using on purpose Web interface, shown in Figure 5.

In case of a device (or a device ensemble connected by proprietary bus) do not support the UPnP protocol, its configuration requires the human installer intervention to connect it to a semantic gateway and to fully edit its description. DDR can give aid also in this scenario. The DDR Web interface, in fact, can be used for editing the semantic device description. Once introduced in the repository, the description can be reused and its information employed to retrieve and to deploy the necessary middleware software that is required for the gateway

The screenshot displays the SAI - Semantic For Ambient Intelligence web interface. The top navigation bar includes 'Device Categories', 'Device Instances', 'Locations', and 'Manufacturers'. A search bar is present with a 'Show Search Form' button. The main content area shows search results for two devices: 'ACME UPnP Color Light' and 'iPhone 4s'. Each result includes a table of properties such as 'declared_power_consumption', 'description', 'type', and 'manufacturer'. The 'ACME UPnP Color Light' result shows a power consumption of 100 W and a manufacturer of 'ACME Inc.'. The 'iPhone 4s' result shows a manufacturer of 'Apple'. On the right side, there are three panels: 'On This Page' with a list of related items, 'Sort by' with a list of sortable fields, and 'View' with a list of viewable fields. At the bottom right, there is an 'Items per page' panel showing '10 > 25 > 50'.

Fig. 5. The Web interface for Device descriptions visualization

configuration. A prototype implementation of the DDR has been currently developed using a Sesame⁴ triplestore for data storage and Ruby on Rails⁵ for the backend logic.

6 Conclusions and Future Works

In this paper, the current efforts towards the implementation of a novel holistic system for a Smart Home environment has been discussed, in particular, focusing on supporting device auto-configuration and intelligent control. With such aims, a semantic extension of UPnP protocol has been developed to foster device automatic communication capabilities without breaking the official standard. In addition, a centralized Device Description Repository has been introduced to provide support for semantic device descriptions management and to foster their sharing and reuse.

References

1. Ducatel, K., Bogdanowicz, M., Scapolo, F., Leijten, J., Burgelman, J.-C.: Scenarios for ambient intelligence in 2010, Final Report, European Commission, Community Research (February 2001)
2. Li, X., Bowers, C.P. Schnier, T.: Classification of energy consumption in buildings with outlier detection. *IEEE Transactions on Industrial Electronics*, 3639–3644 (November 2010)
3. Grassi, M., Nucci, M., Piazza, F.: Towards a semantically-enabled holistic vision for energy optimisation in smart home environments. In: 2011 IEEE International Conference on Networking, Sensing and Control (ICNSC), pp. 299–304 (April 2011)
4. Grassi, M., Nucci, M., Piazza, F.: Towards an ontology framework for intelligent smart home management and energy saving. In: 2011 IEEE International Symposium on Industrial Electronics (ISIE), pp. 1753–1758 (June 2011)
5. Echelon: Lonmaker for windows release 3.1 users guide. Echelon Corporation, USA (2002)
6. Z-wave: The new standard in wireless remote control, <http://www.z-wave.com/modules/AboutZ-Wave/>
7. Knx system specification, <http://www.knx.org/>
8. Upnp forum, upnp device architecture version 1.1, Final Report, European Commission, Community Research (October 2008), <http://upnp.org/specs/arch/UPnP-arch-DeviceArchitecture-v1.1.pdf>
9. Chuah, J.W., Raghunathan, A., Jha, N.: Upnp forum, upnp av architecture:1 (September 2008), <http://www.upnp.org/specs/av/UPnP-av-AVArchitecture-v1.pdf>
10. Digital living network alliance, dlna overview and vision whitepaper (September 2007), http://www.dlna.org/news/DLNA_white_paper.pdf
11. Fipa specifications (September), <http://www.fipa.org/specs/fipa00091/>
12. Cc/pp information page, <http://www.w3.org/Mobile/CCPP/>

⁴ <http://www.openrdf.org/>

⁵ <http://rubyonrails.org/>

13. Kazuyuki, M., Mishina, K., Ren, F., Kuroiwa, S.: Industrial automation systems and integration - open systems application integration framework part 3: Reference description for iec 61158-based controlsystems, ISO 15745 (2003)
14. Can in automation - canopen - overview (2007),
<http://www.canopen.org/canopen/>
15. Togias, K., Goumopoulos, C., Kameas, A.: Ontology-based representation of upnp devices and services dynamic context-aware ubiquitous computing applications. In: Third International Conference on Communication Theory, Reliability and Quality of Service (2010)
16. Hegler, S., Wollschlaeger, M.: The semantic web in action: semantically enabled device descriptions, vol. 2, pp. 1013–1018 (June 2007)
17. Hydra: D2.5 initial requirements report, Hydra Project Deliverable, IST project 2005-034891, vol. 14(13) (2006)
18. Bizer, C., Heath, T., Berners-Lee, T.: Linked data - the story so far. Special Issue on Linked Data, International Journal on Semantic Web and Information Systems (2009)

A Comparison between Different Optimization Techniques for Energy Scheduling in Smart Home Environment

Francesco De Angelis, Matteo Boaro, Danilo Fuselli,
Stefano Squartini, and Francesco Piazza

Dipartimento di Ingegneria dell'Informazione
Università Politecnica delle Marche, Ancona, Italy

Abstract. Nowadays a correct use of energy is a crucial aspect, in fact cost and energy waste reduction are the main goals that must be achieved. To reach this objective an optimal energy management must be obtained through some techniques and optimization algorithms, in order to provide the best solution in terms of cost. In this work a comparison between different methods for energy scheduling is proposed and some analytical results are reported, in order to offer a clear overview for each technique, in terms of advantages and disadvantages. A residential scenario is considered for computer simulations, in which a system storage and renewable resources are available and exploitable to match the user load demand.

1 Introduction

The concept of smart grid faces many electrical power engineering requirements, so different solutions can be achieved for each specific application, from the generation to the customer level, where Computational Intelligence techniques can be very useful [1,2]. In this area the energy management has a main role in order to reduce costs and avoid its waste, also in a micro-grid for a residential or domestic scenario. A joint task and energy optimization framework has been already implemented [3], but several methods have been developed to accomplish efficiently only energy scheduling: linear programming techniques [4], Particle Swarm Optimization (PSO) [5], Fuzzy-Logic [6], Artificial Neural Networks [7], and also Adaptive Dynamic Programming (ADP) [8].

In this paper, the attention is focused in home environment connected to the main grid and also a photovoltaic (PV) system with a battery is considered to increase the saving. The load profile must be always satisfied managing renewable energy, battery and electrical grid in order to reduce costs. Therefore an optimal battery controller must be obtained, whose control policy is to minimize the energy cost imported from the grid managing the battery actions

(charge/discharge) and knowing the forecasted renewable resources, load profile and energy price. In this work a comparison among six different methods, the best promising chosen from literature, for battery management is proposed: an overview for each technique is provided and also a comparison is reported, in terms of advantages and disadvantages. Some of these methods have been already presented in other papers like [9] based on Adaptive Dynamic Programming (ADP) and on Particle Swarm Optimization (PSO) technique; some other methods, based on Linear Programming (LP) and PSO, are introduced for the first time.

The system description is reported in Section 2, the analytical issues of each optimization algorithm are discussed in Section 3 and the simulated scenario is shown in Section 4. Section 5 deals with the conducted computer simulations whereas Section 6 draws the work conclusions.

2 Home Energy System Description

The proposed home model is composed of a main electrical grid, external PV array, storage system and Power Management Unity (PMU), that ensures the meeting of load demand. As reported in Fig. 1, PMU unit (energy scheduler) manages the energy flows: battery can be charged from the grid and/or from PV, moreover if necessary it can be discharged to supply the load. If there is exceeded energy from PV not usable from the system, it is sold to the main grid. In addition the battery must satisfy the following constraints:

1. The charging and discharging rate can not be exceeded.
2. Battery level must be always included between the upper and lower bound.

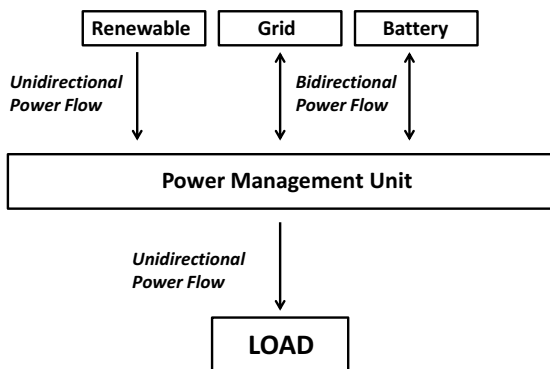


Fig. 1. Power Flows

3 Optimization Algorithms

3.1 Linear Programming Technique: LP Algorithm

The implemented algorithm is based on the “Linear Programming” (LP) paradigm. Its objective consists in maximizing or minimizing a given function, considering some constraints according to:

$$\begin{aligned} \max f(x) = c^T x \quad \text{or} \quad \min f(x) = c^T x \\ \text{subject to} \quad Ax \leq b \quad \text{or} \quad Ax \geq b \end{aligned}$$

where $x \geq 0$, $x \in \mathbb{R}^{n \times 1}$, $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^{m \times 1}$, $c \in \mathbb{R}^{1 \times n}$.

The cost function $U(t)$ used in this work with this optimization algorithm is the following:

$$U(t) = \sum_{t=1}^T \{ [L_u(t) - R_u(t) + Ch_u(t) - Dh_u(t)] \cdot C(t) \} \tag{1}$$

where $L_u(t)$ is the load demand at temporal slot t , $R_u(t)$ is the amount of renewables exploited at time t , Ch_u is the amount of energy used for charging the system storage, Dh_u is the discharged energy used for meeting (partially or totally) the load, $C(t)$ is the electricity cost at time t and T is the work horizon.

This kind of battery controller works in offline way, because it optimizes the given cost function in all the time horizon T .

The constraints of this problem, for $1 \leq t \leq T$, are reported as follows:

- Positive function: $L_u(t) - R_u(t) + Ch_u(t) - Dh_u(t) \geq 0$
- Do not exceed renewables: $R_u(t) \leq R(t)$
- Load is fixed $L_u(t) = L(t)$
- Charge and discharge limits: $Ch_u(t) \leq Ch_{rate}$, $Dh_u(t) \leq Dh_{rate}$
- Battery level: $SL(t) = SL(t - 1) + Ch_u(t) - Dh_u(t)$
- Battery level limits: $SL_{MIN} \leq SL(t) \leq SL_{MAX}$

where $R(t)$ and $L(t)$ are respectively the total renewable available and load demand, while SL is the State of Charge (SoC) of the system storage.

3.2 Particle Swarm Optimization Algorithm

PSO is a technique inspired to certain social behaviors, and it is used to explore a search parameter space to find values allowing to minimize an objective function [10]. The PSO algorithm works by maintaining simultaneously various candidate solutions (particles in the swarm) in the search space. In PSO, the coordinates of each particle represent a possible solution associated with two vectors, the position x and velocity v vectors in N -dimensional search space. A swarm consists of a number i of particles “or possible solutions” that flies through the feasible

solution space to find the optimal one. Each particle updates its position x_i on the basis of its own best exploration p_i , its best swarm overall experience p_g , and its previous velocity vector $v_i(k-1)$ according to (2) and (3).

$$x_i(k) = x_i(k-1) + v_i(k) \quad (2)$$

$$v_i(k) = v_i(k-1) + \rho_1 \cdot rand_1 \cdot [p_i - x_i(k-1)] + \rho_2 \cdot rand_2 \cdot [p_g - x_i(k-1)] \quad (3)$$

where ρ_1 and ρ_2 are two positive correction factors, k is the iteration step while $rand_1$ and $rand_2$ are two random numbers [0.0, 1.0]. The PSO algorithm can be described in general as follows:

1. For each particle, randomly initialize the position and velocity vectors with the same size as the problem dimension.
2. Measure the fitness (utility function value) of each particle and store the particle with the best fitness value (minimum utility function value).
3. Update velocity and position vectors according to (2) and (3) for each particle.
4. Repeat steps 2 and 3 until a termination criterion is satisfied.

As already done in [9] we introduce in (4) an utility function that must be minimized for each temporal slot t .

$$U(t) = \sqrt{\{[L(t) - R(t) + u(t)] \cdot C(t)\}^2 + \{SL_{cap} - [SL(t) + u(t)]\}^2} \quad (4)$$

where $u(t)$ is the optimized value of battery charge ($u(t) > 0$) or discharge ($u(t) < 0$) that must be found by the algorithm for each time t , SL_{cap} is the battery capacity and $SL(t)$ is the actual battery level. Minimizing $U(t)$ means charging the battery when renewable is high and/or when cost is low, and discharging the battery when renewable is lower than the load and/or the cost is high. Obviously $u(t)$ must satisfy the two battery constraints discussed in Section 2. If one of these constraints is not satisfied, the obtained solution $u(t)$ is not valid and must be discarded. So the function is multiplied with a penalty factor which is set to a higher value.

It is important to note that this battery controller works in online way, because the cost function is evaluated step by step without knowing the energy horizon profiles.

3.3 Extended Particle Swarm Optimization Algorithm

Similar to the scheme proposed in Section 3.2, an extended version of PSO has been realized. The operation is not online anymore, but offline in order to give an optimal solution on an extended period, for which all scenario profiles are considered in the work horizon, as well as the forecasted data about renewable energy. Differently from (4), the utility function adopted in this case does not

include battery terms, and also a sum over the entire period is considered, in order to provide an optimization for the entire work horizon T .

$$U(t) = \sum_{t=1}^T \sqrt{\{[L(t) - R(t) + u(t)] \cdot C(t)\}^2} \quad (5)$$

Obviously $u(t)$ must satisfy the battery constraints discussed in Section 2.

3.4 Adaptive Dynamic Programming

Combining approximate dynamic programming and reinforcement learning, Werbos proposed a new optimization technique [11], whose goal is to design an optimal control policy, which can be able to minimize a given cost function called “utility function” (especially in nonlinear and noisy environments), adapting two neural networks: the Action Network and the Critic Network. The Action Network, taking the current state as input, has to drive the system to a desired one, providing a control to the latter. The Critic Network, knowing the state and the control provided by the Action Network, has to check its performances and return to the Action Network a feedback signal to reach the optimal state over time. To check Action performances, the Critic Network approximates the following Bellman equation associated with optimal control theory:

$$J(t) = \sum_{i=0}^{\infty} \gamma^i U(t+i) \quad (6)$$

where γ is the discount factor (0, 1] and $U(t)$ is the utility function.

As already implemented in [9], that was inspired by [8], an Action-Dependent Heuristic Dynamic Programming (ADHDP) model free approach is adopted for the design of an optimal controller, whose goal is to manage the battery, knowing forecasted data (Load, Price, Renewable Energy), in order to save money during an overall time-horizon. As reported in Fig. 2 the input to the Action network is the system state $x(t)$, and the output $u(t)$ is the amount of energy used to charge or discharge the battery; the input of the Critic Network consists of the current system state and the current control provided by the Action Network.

The used Critic network is composed by 15 linear neurons in input, 40 sigmoidal hidden neurons and 1 linear in output, while Action network by 4 linear neurons in input, 40 sigmoidal hidden neurons and 1 linear in output. In this study the proposed utility function $U(t)$ is reported in (7).

$$U(t) = \sqrt{\{[L(t) - R(t) + u(t)] \cdot C(t)\}^2} \quad (7)$$

where $u(t)$ is the optimized value of battery charge ($u(t) > 0$) or discharge ($u(t) < 0$) that must be found for each time t . Obviously $u(t)$ must satisfy the battery constraints discussed in Section 2 and in this case it is forced after Action Network output to respect these limits. When the utility function is minimized the control policy is optimal and the cost is the lowest. The squaring of the equation is necessary to avoid that $U(t)$ is negative.




Fig. 2. ADHDP Scheme

The online training is based on the “Backpropagation algorithm”: the iterative training used for both neural networks is repeated for a fixed number of epochs and explained step by step below:

1. The Action and Critic weights are initialized before the training: with random values $[-1,1]$ or pre-trained with extended PSO.
2. Train Critic Network refreshing its weights using computing Critic error (E_c), then refresh Action Network computing Action error (E_a).
3. Evaluate the system performance computing the total cost to minimize in the time horizon. If the cost decreases, the control policy is improving and the new action weights are the best; if not, revert to old action weights and add a small random perturbation. Then restart the training from Step 2.

Also this type of battery controller is an offline one because it optimizes the utility function during all the time horizon T , since it employs forecasted data.

3.5 Self-Learning Procedure Based on ADHDP Scheme

Like in [12] this optimization procedure is based on a simplified ADHDP scheme because only few actions can be done by the controller, in fact the battery is limited to a ternary choice (charge, discharge or idle). In this way we consider only a critic network in the scheme. If a network is trained correctly, whenever power demand occurs the critic network verifies which is the action that involves the smallest output value, so the most convenient action is chosen. The training procedure is the following:

1. Data are collected: the action is taken randomly, the state is characterized by the cost rate, the load profile, the battery level and the renewable energy;
2. Compute $U(t)$ and $Q(t)$ in order to obtain the target, since the training is based on the mapping: $\{x(t - 1); u(t - 1)\} \rightarrow \{U(t) + \gamma Q(t)\}$, where $x(t - 1)$ and $u(t - 1)$ are the previous state and control, $U(t)$ is the actual utility function, $\gamma \in (0, 1]$ is a discount factor and $Q(t)$ is the actual critic network output;

3. The critic network is trained with the “Levenberg-Marquardt backpropagation” algorithm;
4. Eventually the neural network can be re-trained whenever there are consistent changes in the scenario.

The utility function that we want to minimize is:

$$U(t) = [L(t) - R(t) + u(t)] \cdot C(t) \tag{8}$$

where $u(t) = u'(t)q(t)$ is the battery charge ($u(t) > 0$) or discharge ($u(t) < 0$), $u'(t)$ is the battery action $(1, -1, 0)$ and $q(t)$ is the charging/discharging battery quantity. Also in this case $u(t)$ must satisfy the battery constraints discussed in Section 2 and in this case it is forced to respect these limits.

Also this battery control strategy is offline and considers the overall working horizon T .

4 Simulated Scenario

All the simulations reported in Section 5 refer to the same scenario: a system storage is supposed to be available, as well as renewable resources deriving from solar energy. In Texas, in Austin city, we consider an area of 30 m^2 covered by some photovoltaic (PV) panels, whose efficiency is 30 %, and irradiation data is taken from [13]. According to [14], the available renewable energy is computed with $P = GHI \cdot \eta_{pv} \cdot A_{pv}$, where GHI is the Global Horizontal Irradiance in Wh/m^2 received on a horizontal surface, η_{pv} is the efficiency of the PV and A_{pv} is the total area of the PV panel in m^2 . The difference between the simulations is the considered time horizon: 48-h, 96-h and 168-h horizons are simulated in order to test performances of each technique for the short and long-term period. In the Tab. 1 system storage parameters used for all the simulations are reported, and since the resolution time used is one hour, kWh and kW agree so we can consider the same unit of measurement both for energy and power parameters.

Table 1. Storage system parameters (in kW)

SL_0	SL_{MIN}	SL_{MAX}	Ch_{rate}	Dh_{rate}
5	0	10	1	1

In Tab. 1 SL_0 , SL_{MIN} and SL_{MAX} are respectively the initial, minimum and maximum State of Charge (SoC), while Ch_{rate} and Dh_{rate} are the maximum charge and discharge rate of the considered storage system. The efficiency η has been considered equal to 100%.

5 Computer Simulations

In this study a battery management problem is considered in order to minimize the imported energy from the main grid and increase the money saving. The cost (expressed in dollars) related to each optimization technique for three different time horizon (48-h, 96-h and 168-h) is reported in Tab. 2, while the money saving, compared to an online baseline algorithm applied in the same scenario, is given in percentage in Tab. 3. The baseline approach follows the next simple rules for each time step:

- if the load is greater than the available renewable energy, the difference is supplied discharging the battery (according with Dh_{rate} in Tab. 1). If the battery support is not enough, the needed energy to supply totally the load is imported from the main grid;
- if the available renewable energy is greater than the load demand, the surplus is used to charge the battery (according with Ch_{rate} in Tab. 1). If the battery is already full or the surplus is greater than the charging rate, the amount of energy in excess, not usable in other ways, is sold to the main grid.

Table 2. Cost comparison for energy scheduling (in \$)

T	LP	ADP	$Ext\ PSO$	$Self-L\ ADP$	PSO	$Baseline$
48 h	6.46	6.48	6.49	6.63	6.97	7.12
96 h	12.80	12.84	12.86	13.08	13.74	14.09
168 h	25.08	25.16	25.22	25.75	26.04	26.79

Looking at the results reported in Tables 2-3 it is evident that the LP offline algorithm provides the best solution, due to the linear behavior of the energy scheduling problem. Furthermore it has no convergence problems and the computational cost is very low, but as mentioned this method can not work in real-time and it needs forecasted data relative to the considered horizon. Whenever a linear approximation of a nonlinear model is not valid, this linear approach cannot be used and a different method should be chosen.

Table 3. Money saving (in percentage) compared to baseline algorithm

T	LP	ADP	$Ext\ PSO$	$Self-L\ ADP$	PSO
48 h	9.27%	8.99%	8.85%	6.89%	2.10%
96 h	9.21%	8.87%	8.73%	7.17%	2.55%
168 h	6.38%	6.08%	5.86%	3.88%	2.80%

Differently, PSO is an online algorithm able to work without forecasted data, and it optimizes step by step a given utility function, with very low computational complexity. For this reason it is not possible to offer an optimal solution

over a large time horizon, so the cost reduction is limited. Different is the case of Extended PSO, which gives a good solution over a considered work horizon, even if its performances degrade gradually while the horizon increases (due to the fact that the number of unknowns increases). This offline method needs forecasted data relative to all the temporal steps considered, and it has a computational cost higher than the previous mentioned techniques.

As mentioned, the Extended PSO is used to pretrain the neural networks used in ADP method. The ADP, adapting the Action and Critic weights, can improve the performances of the Extended PSO and find a better solution with an higher saving. This saving is remarkable especially in longer periods, where the ADP overcomes the performances of the Extended PSO, which finds an optimal solution over longer period because of the variables number increase.

The initial computational cost of the ADP algorithm is not very small, but it has the advantage to adapt itself quite quickly when the time horizon and the scenario change; in fact the optimization process can continue from the best weights of the neural networks stored on the previous training step, and it does not need to restart from the beginning like the other proposed methods.

Finally, the self-learning procedure based on ADHDP scheme offers a trade off between the goodness of the solution and the computational cost: slightly sacrificing the cost reduction, a much shorter time spent for the neural network training can be obtained.

6 Conclusions

A comparison between different optimization techniques for the energy scheduling in a smart home environment has been proposed, and an evaluation in monetary terms has been given in order to highlight the performances also referred to a baseline approach. Although all the shown methods are valid, the LP algorithm offers the best solution, since the energy scheduling problem is linear. Also the other techniques provide good solutions, and obviously they could solve optimization problems in more complex and nonlinear systems. PSO works fine in online configuration, but its extended offline version gives a better cost reduction. Both the two ADP procedures can be advantageous concerning cost reduction and computational cost for long-term periods. In conclusion, the choice among these different offline and online methods depends on the linearity of the problem, the computational cost and the availability of forecasted profiles.

References

1. Harley, R.G., Liang, J.: Computational Intelligence in Smart Grids. In: IEEE SSCI 2011 Symposium Series on Computational Intelligence CIASG (2011)
2. Venayagamoorthy, G.K.: Potentials and Promises of Computational Intelligence for Smart Grids. In: 2009 IEEE Power and Energy Society General Meeting, pp. 1–6 (2009)

3. De Angelis, F., Boaro, M., Fuselli, D., Squartini, S., Piazza, F., Wei, Q., Wang, D.: Optimal Task and Energy Scheduling in Dynamic Residential Scenarios. In: Wang, J., Yen, G.G., Polycarpou, M.M. (eds.) ISSN 2012, Part I. LNCS, vol. 7367, pp. 650–658. Springer, Heidelberg (2012)
4. Morais, H., Kádár, P., Faria, P., Vale, Z.A., Khodr, H.M.: Optimal scheduling of a renewable micro-grid in an isolated load area using mixed-integer linear programming. *Renewable Energy - An International Journal* 35(1), 151–156 (2009)
5. Gudi, N., Wang, L., Devabhaktuni, V., Depuru, S.S.S.R.: A Demand-Side Management Simulation Platform Incorporating Optimal Management of Distributed Renewable Resources. In: IEEE/PES Power Systems Conference and Exposition (PSCE), pp. 1–7 (2011)
6. Liang, R.H., Liao, J.H.: A Fuzzy-Optimization Approach for Generation Scheduling with Wind and Solar Energy Systems. *IEEE Transactions on Power Systems. Journals and Magazines* 22(4), 1665–1674 (2007)
7. Vale, Z.A., Faria, P., Morais, H., Khodr, H.M., Silva, M., Kadar, P.: Scheduling Distributed Energy Resources in an Isolated Grid –An Artificial Neural Network Approach. In: IEEE Power and Energy Society General Meeting, pp. 1–7 (2010)
8. Welch, R.L., Venayagamoorthy, G.K.: Energy Dispatch Controllers for a Photovoltaic System. *Engineering Applications of Artificial Intelligence*, 249–261 (2008)
9. Fuselli, D., De Angelis, F., Boaro, M., Liu, D., Wei, Q., Squartini, S., Piazza, F.: Optimal Battery Management with ADHDP in Smart Home Environments. In: Wang, J., Yen, G.G., Polycarpou, M.M. (eds.) ISSN 2012, Part II. LNCS, vol. 7368, pp. 355–364. Springer, Heidelberg (2012)
10. Del Valle, Y., Venayagamoorthy, G.K., Mohagheghi, S., Hernandez, J.C., Harley, R.G.: Particle Swarm Optimization: Basic Concepts, Variants and Applications in Power Systems. *IEEE Transactions on Evolutionary Computation. Journals and Magazines* 12(2), 171–195 (2008)
11. Werbos, P.J.: Approximate Dynamic Programming for Real-Time Control and Neural Modeling. In: *Handbook of Intelligent Control* (1992)
12. Huang, T., Liu, D.: Residential Energy System Control and Management using Adaptive Dynamic Programming. In: International Joint Conference on Neural Networks (IJCNN), pp. 119–124 (2011)
13. National Renewable Energy Laboratory (NREL) of U.S. Department of Energy, <http://www.nrel.gov/rredc/>, Office of Energy Efficiency and Renewable Energy, operated by the Alliance for Sustainable Energy, LLC
14. Markvart, T.: *Solar Electricity*, 2nd edn. Wiley, New York (2000)

Part V

**Special Session on
“Computational Intelligence in
Emotional or Affective Systems”**

Towards Emotion Recognition in Human Computer Interaction

Günther Palm^{1,2,*} and Michael Glodek²

¹ Domenico Labate, MECMAT - Mediterranean University of Reggio Calabria, via Graziella Feo di Vito, 89122 Reggio Calabria, Italy

² Institute of Neural Information Processing, University of Ulm, Germany
firstname.lastname@uni-ulm.de

Abstract. The recognition of human emotions by technical systems is regarded as a problem of pattern recognition. Here methods of machine learning are employed which require substantial amounts of 'emotionally labeled' data, because model based approaches are not available. Problems of emotion recognition are discussed from this point of view, focusing on problems of data gathering and also touching upon modeling of emotions and machine learning aspects.

Keywords: Affective computing, emotion recognition, human computer interaction, companion systems.

1 Introduction

In the field of human-computer interaction (HCI) there is a growing interest in using not only explicit function-oriented signals for the cooperation between a human user and a technical system, but also implicit “affective” or “emotional” signals.

In spite of a long tradition in psychology (e.g. [1,2,3,4,5,6,7,8]) there still is no undisputed definition of emotions, neither of particular emotions nor of the general concept of emotion. From a biological perspective emotions can be described as modulations of animal behavior in response to certain extreme circumstances. For example, in extreme danger, an animal may focus its attention to the relevant sensory modalities and to particular spatio-temporal constellations in these modalities; in addition it may mobilize all its physical energy and strength in particular groups of muscles to enable quick reactions like fleeing or fighting, and it may shift its evaluative weights (which will be explained in a bit more detail below) towards taking more risks of injuries and pain in the course of actions. This scenario describes a modulation of the three main ingredients of goal-directed behavior: sensation, action and evaluation. It is plausible that in the course of evolution and also of individual development animals have found

* Corresponding author.

a small number of such combined modulatory settings that are useful for the modulation of behavior in certain situations. These settings may be what we call emotions.

In social animals it may also be useful to signal these settings to other animals in order to prevent, for example, that animals of the same species hurt each other unnecessarily. So emotional states may also be conveyed by social signals, which in humans often are unconsciously produced in addition to or as modification of our explicit communication.

The biological approach sketched here may sound provocative to some psychologists, but it boils down to defining emotions as particular modulators of behavior which may also have some physiological consequences, for example concerning heart-rate, blood-flow, body-heat and skin resistance.

Some attempts have been made to define “basic” emotion categories (typically around six of them, see for example [9], but the varying degree of similarity between these categories suggests an embedding of them into a 2- or 3-dimensional space; most commonly in a circular fashion ([3] or the “emotion wheel” of [8]). The dimensions or axes of this space have also been identified ([10,6] as “value” (from negative or “bad” to positive or “good”), “activation” (from low or “calm” to high or “agitated”) and “dominance” or perhaps also “competence” or “influence” (from low or “passive”, “enduring” to high or “active”, “in control”). Another aspect of the classification of emotions is their variable strength or expressiveness.

When we try to incorporate emotionality in HCI, we also have to consider emotions from the computer science point-of-view. Here there are three tasks that can be distinguished although they are often used in combination:

1. detection and classification of human emotions,
2. modeling of emotions in an artificial agent,
3. displaying emotional signals towards the human user.

1. *Emotion detection* can be viewed as a particular branch of pattern recognition and will be in the focus of this article.
2. *Emotion modeling* is an interesting enterprise which may follow the biological approach described above and thereby relate to neurobiological and physiological research on this topic (e.g. [11,12,13,14,15,16]). The neuroscientific accounts of emotions are concentrating on fMRI experiments showing emotional reactions in particular brain areas and also on the idea of explaining the cognitive faculty of empathy as a use of the own “emotion generating system” in those areas to simulate or “mirror” the emotions of others ([17,18,19,20,21,22,23,24]), similar to the mirror neuron system representing intentional actions. Theoretical approaches often rely on variations of reinforcement learning ([25,26,27,12,28,29,30,31,32,33,15,34]). The basic idea of reinforcement learning (RL) is that an agent learns to predict an evaluative signal (called reward) based on its sensory input and the course of action it is about to take. It then takes the course of action that maximizes the predicted reward. Biologically realistic versions of RL assume a small number

of different motives or objectives like eating, drinking, having sex, exploring territory, avoiding pain, which the agent may want to optimize with different weights or different priorities depending on its current situation (these are the evaluative weights mentioned above).

There is also an interesting modeling approach by D. Dörner and his scholars ([35,36,37,38]) that is closely related to these ideas. Another approach to the modeling of emotions is motivated from cognitive psychology ([39,40,41]) treating the relation between emotions and motivations or appraisals on a higher cognitive level with less emphasis on the underlying physiological processes and on the dynamical temporal aspects of emotions.

3. *Emotion display* has become quite popular in HCI in applications concerned with entertainment and computer games. It makes use of the fact that humans readily project emotions into artificial agents they interact with. In a feedback situation of HCI it can be useful to combine emotion recognition with emotion display, for example in order to improve and stabilize the emotion recognition performance on both sides in a dyadic HCI.

In fact, one can perhaps divide the whole area of affective computing into those domains where the expression of emotions by the computer is more important, which will most likely occur in recreational activities like entertainment, edutainment, or game playing, and those domains where the recognition of human emotions by the computer is more important, which will most likely occur in the context of work, when the computer helps to achieve a given task. In our collaborative research center (see <http://www.sfb-trr-62.de/>) we focus on the second type of interaction and we define *Companion technology* as the technology that helps to improve this kind of HCI.

2 Recognition of Emotions or User Dispositions

In HCI the recognition of emotions has turned out to be a very hard problem of pattern recognition and machine learning due to the following reasons:

- Emotions are rare. They do not occur often in human interaction and perhaps even less in HCI.
- Most emotions can occur in different degrees and they occur only in weak degrees most of the time.
- This often leads to a high degree of uncertainty about the recognized emotion and, in addition, it may be necessary to distinguish the degree of recognition uncertainty from the degree of expressiveness of the present emotion.
- Emotions are expressed multi-modally. So it is often useful and even necessary to combine several sensory modalities (audition, vision, and biophysical measurements, if available) in their detection and classification.
- Emotion recognition is context dependent. Humans often cannot correctly identify emotions unless they are provided within broader context of the given situation. For artificial systems this implies that one should make use

Table 1. Multimodal datasets of HCI with annotated emotions

Name	Modalities ¹	Size ²	Annotation ³	HCI type ⁴	Remarks ⁵	Reference
AVEC	A,V	95/13/7.5h	4 D 2-3 raters	D	T,E	[78]
EmoRec	A,V,P	110/110/73h	3 D not req.	S,T	E	[75]
Humaine	A,V	50/?/3.3h	4 D n/a	D	E	[79]
Last Minute	A,V,P	126/126/63h	not yet annotated	S,C	T,C	[80]
Nimitek	A,V	10/10/15h	6 C n/a	T	T,E	[81]
PIT	A,V	37/74/9h	> 9 C 2-3 raters	C,O	C	[82]
SAL	A,V	20/4/10h	6 C, 4 D n/a	D	E	[83]
Smartkom	A,V	224/224/16.8h	9 C n/a	C,O	T	[84]

¹ (A) audio, (P) physiological, (V) video

² #/#/# number of recordings/number of subjects/hours of recording

³ # number of labels, (C) categories (D) dimensions;
number of raters, (n/a) number of raters not available,
(not req.) rating not required

⁴ (T) trainer/teacher, (M) monitor, (O) organizer, (C) consultant, (S) servant,
(D) discourse

⁵ (T) transcript available, (E) emotions are evoked, (C) multiple cameras

of high-level symbolic information, if available, and at least use temporal integration of the given multimodal signals over a longer time (e.g. covering a whole interaction sequence, not just one utterance or turn in a dialogue).

These observations have been made in our own attempts to create a “*Companion technology*” ([42]) for HCI that allows the computer to react to human emotional signals. There is not much scientific literature yet on these topics to substantiate these preliminary observations (e.g. [43,44,45,46,47,48,49,50]), since the field of “affective computing” is still in its infancy. Early experiments on unimodal emotion recognition ([51,7,52,53,54,55]) had to rely on artificial acted emotions. And even on these data (e.g. [56,57,58,59]) the typical performance for 4 or 5 emotion categories was around 70 to 80 percent, usually about as bad or just a little worse than human performance on the same data.

There are some projects providing large multimodal databases of almost natural human conversations (see [60,61,62,63,64,65,66,67,68,69,70]) which can be used to substantiate the 5 points above, but not much work on natural emotion recognition has appeared so far (e.g. [48,71,49,72,73,74,75,76,77]).

A serious problem of emotion recognition from the engineering point of view is that, strictly speaking, there is no ground truth that can be used for labeling. Even if one tries to work with induced emotions the exact emotional category that is induced may strongly depend on the character of the subject, in addition to ethical problems involved in the induction of strong negative emotions. The only available ‘gold-standard’ is human rating of the observed emotional behavior and this is far from being unequivocal.

Concerning emotions in HCI the available data are still very sparse. Currently we are aware just of a few reasonably extensive multimodal HCI databases (see Table II). In fact, any project of data-collection for affective HCI is faced with a number of very particular challenges:

- Technical problems with gathering and displaying large amounts of synchronized multi-modal data in a reasonable HCI scenario (often requiring WOZ methodology).
- Determining the relevant label categories for the data, which may even be application- or scenario-dependent (see below).
- Developing tools that make it possible for annotators to add (un)certainity values to their labels (like [85,86,87]).

2.1 Identifying User Dispositions in Companion Technology

In spite of these problems we expect much further progress in affective HCI for the development of *Companion* systems, based on a few conceptual and technical developments in the next years.

First, we said that emotions are rare in HCI, and indeed some of the genuine human emotions may not even be relevant for typical *Companion* systems. Hence one can try to restrict the emotional categories that have to be recognized and distinguished to the practically relevant ones. To this end it would of course be useful not to consider every application separately, but to identify a small set of typical applications or *Companion tasks*. In our understanding of *Companion* systems we are not focusing on entertainment systems, but rather on applications where the *Companion* system helps the user to achieve some goal. A typical problem of this kind occurs if the user has to rely on a conventional technical system which is complicated or partially unknown to the user. In such a case the *Companion* system can mediate between the user and the conventional system. Another type of application occurs in training or teaching where the system guides the user in extending his knowledge or his practical or physical capabilities.

To arrive at a short list of generic *Companion* tasks we can consider the typical examples of assistance that are proposed or already offered by modern technical systems in the household (cleaning, cooking, preparing or organizing meals, invitations), in the car (navigation, driving assistance), at work (keeping track of dates, meetings and duties, finding information, e.g. on products, transportation, addresses, making connections to other people), simplifying or personalizing use of machinery (cameras, coffee-machines, cell-phones, audio-video-displays, diagnosis and repair), in education (learning, edutainment), in buying, ordering things, in the hospital and rehabilitation (monitoring, control of medication, motivation of physical exercises).

We have attempted to condense this down to the following five types of *Companions*: **Trainer/teacher, monitor, organizer, consultant, servant.**

When we now consider the emotional dispositions of the user that will be relevant for the reactions of the *Companion* in these cases, it turns out that in

practically all cases it is the disposition or attitude of the user towards the task at hand and towards the *Companion* system that matters. The central affective goal of the system is to maintain a positive attitude of the user, for example by adjusting the amount, redundancy and complexity of the information it is providing, or by giving ensuring or just evaluating feedback to the user (and, of course, by functioning properly). Thus the main practical objective in emotion recognition for such systems will be to distinguish different kinds of **negative** emotions or attitudes from each other, and also from neutral or “no problem”. Here we propose to distinguish the following seven: **Bored, disengaged, frustrated, helpless, over-strained, angry, impatient.**

Again this short list of specific application oriented user attitudes or dispositions has to be regarded as preliminary and it certainly needs further explanation, for example in terms of emotion dimensions like activation and dominance, in particular when it is used for labeling. Such a short list, however, would help a lot to simplify the task of emotion recognition for *Companion* systems and to unify the labeling of corresponding affective HCI databases for *Companion* systems.

Another reason for anticipating a fast progress in emotion recognition lies in the development of new ideas and techniques in pattern recognition and machine learning. There are a number of new ideas concerning information and classifier fusion ([88,89,90,91,92,93,94,95,96,97]) which have lead to better theoretical understanding and practical results in the fusion of information from different sources and have already been applied to the recognition of user dispositions or emotions (e.g. [96,98,99,100,101,102,103]) There are also new results on semi-supervised learning ([104,105,106,107,108,109,110,111,112,113,114,115,116,117,118]) and on learning from uncertain teacher signals ([119,120,121]) that can be used to simplify the labeling of large datasets with uncertain or partially missing labels.

We believe that, based on these ideas and developments, the field of affective computing for *Companion* systems and also for HCI in general is going to make considerable progress in the next years.

Acknowledgment. This work was stimulated by extensive discussions and financially supported within the Transregional Collaborative Research Centre SFB/TRR 62 ”Companion-Technology for Cognitive Technical Systems” funded by the German Research Foundation (DFG).

References

1. Cannon, W.: The James-Lange theory of emotions: A critical examination and an alternative theory. *The American Journal of Psychology* 39(1/4), 106–124 (1927)
2. Ekman, P., Friesen, W.: The repertoire of nonverbal behavior: Categories, origins, usage, and coding. *Semiotica* 1(1), 49–98 (1969)

3. Russell, J.: A circumplex model of affect. *Journal of Personality and Social Psychology* 39(6), 1161 (1980)
4. Ekman, P.: Facial expression and emotion. *American Psychologist* 48, 384–392 (1993)
5. Picard, R.W.: *Affective Computing*. The MIT Press (1997)
6. Russell, J., Barrett, L.: Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant. *Journal of Personality and Social Psychology* 76(5), 805 (1999)
7. Cowie, R., Douglas-Cowie, E., Tsapatsoulis, N., Votsis, G., Kollias, S., Fellenz, W., Taylor, J.: Emotion recognition in human-computer interaction. *IEEE Signal Processing Magazine* 18(1), 32–80 (2001)
8. Scherer, K.: What are emotions? And how can they be measured? *Social Science Information* 44(4), 695–729 (2005)
9. Ekman, P.: An argument for basic emotions. *Cognition & Emotion* 6(3–4), 169–200 (1992)
10. Mehrabian, A.: Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in temperament. *Current Psychology* 14(4), 261–292 (1996)
11. Damasio, A.: Descartes' error and the future of human life. *Scientific American* 271(4), 144–145 (1994)
12. Dayan, P., Balleine, B.: Reward, motivation, and reinforcement learning. *Neuron* 36(2), 285–298 (2002)
13. Taylor, J., Fragopanagos, N.: The interaction of attention and emotion. *Journal of Neural Networks* 18(4), 353–369 (2005)
14. Fogassi, L., Ferrari, P.: Mirror neurons and the evolution of embodied language. *Current Directions in Psychological Science* 16(3), 136–141 (2007)
15. Dayan, P., Huys, Q.: Serotonin in affective control. *Annual Review of Neuroscience* 32, 95–126 (2009)
16. Niv, Y.: Reinforcement learning in the brain. *Journal of Mathematical Psychology* 53(3), 139–154 (2009)
17. Rizzolatti, G., Arbib, M.: Language within our grasp. *Journal of Trends in Neurosciences* 21(5), 188–194 (1998)
18. Rizzolatti, G., Fogassi, L., Gallese, V.: Mirrors in the mind. *Scientific American* 295(5), 54–61 (2006)
19. Rizzolatti, G., Sinigaglia, C.: *Mirrors in the brain: How our minds share actions, emotions*. Oxford University Press (2008)
20. Arbib, M.: *Beyond the Mirror: Biology and Culture in the Evolution of Brain and Language*. Oxford University Press (2005)
21. Arbib, M.: *Action to language via the mirror neuron system*. Cambridge University Press (2006)
22. Bonaiuto, J., Rosta, E., Arbib, M., et al.: Extending the mirror neuron system model, I. Audible actions and invisible grasps. *Journal of Biological Cybernetics* 96(1), 9 (2007)
23. Bonaiuto, J., Arbib, M.: Extending the mirror neuron system model, II: What did i just do? A new role for mirror neurons. *Journal of Biological Cybernetics* 102(4), 341–359 (2010)
24. Gallese, V., Goldman, A.: Mirror neurons and the simulation theory of mind-reading. *Journal of Trends in Cognitive Sciences* 2(12), 493–501 (1998)

25. Albus, J.: Outline for a theory of intelligence. *IEEE Transactions on Systems, Man and Cybernetics* 21(3), 473–509 (1991)
26. Bertsekas, D., Tsitsiklis, J.: *Neuro-dynamic programming*. *Journal of Optimization and Neural Computation* 3 (1996)
27. Sutton, R., Barto, A.: *Reinforcement learning: An introduction*, vol. 1. Cambridge University Press (1998)
28. Wörgötter, F.: Actor-critic models of animal control — A critique of reinforcement learning. In: *Proceedings of International Symposium on Engineering of Intelligent Systems*, EIS (2004)
29. Wörgötter, F., Porr, B.: Temporal sequence learning, prediction, and control: A review of different models and their relation to biological mechanisms. *Journal of Neural Computation* 17(2), 245–319 (2005)
30. Oudeyer, P., Kaplan, F., Hafner, V.: Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation* 11(2), 265–286 (2007)
31. Izhikevich, E.: Solving the distal reward problem through linkage of STDP and dopamine signaling. *Journal of Cerebral Cortex* 17(10), 2443–2452 (2007)
32. Botvinick, M., Niv, Y., Barto, A.: Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *Journal of Cognition* 113(3), 262–280 (2009)
33. Bhatnagar, S., Sutton, R., Ghavamzadeh, M., Lee, M.: Natural actor-critic algorithms. *Automatica* 45(11), 2471–2482 (2009)
34. Huys, Q., Dayan, P.: A Bayesian formulation of behavioral control. *Journal of Cognition* 113(3), 314–328 (2009)
35. Bartl, C., Dörner, D.: PSI: A theory of the integration of cognition, emotion and motivation. In: *Proceedings of the European Conference on Cognitive Modelling (ECCM)*, DTIC Document, pp. 66–73 (1998)
36. Dörner, D.: *Bauplan für eine Seele (Blueprint of a Soul)*. RoRoRo (1999)
37. Hille, K.: Synthesizing emotional behavior in a simple animated character. *Journal of Artificial Life* 7(3), 303–313 (2001)
38. Bach, J.: *Principles of synthetic intelligence — PSI: an architecture of motivated cognition*. Oxford University Press (2008)
39. Ortony, A., Clore, G., Collins, A.: *The cognitive structure of emotions*. Cambridge University Press (1988)
40. Tomasello, M., Carpenter, M., Call, J., Behne, T., Moll, H., et al.: Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences* 28(5), 675–690 (2005)
41. Marsella, S., Gratch, J.: EMA: A process model of appraisal dynamics. *Journal of Cognitive Systems Research* 10(1), 70–90 (2009)
42. Wendemuth, A., Biundo, S.: A Companion Technology for Cognitive Technical Systems. In: Esposito, A., Esposito, A.M., Vinciarelli, A., Hoffmann, R., Müller, V. (eds.) *Cognitive Behavioural Systems 2011*. LNCS, vol. 7403, pp. 89–103. Springer, Heidelberg (2012)
43. Scherer, K., Ceschi, G.: Criteria for emotion recognition from verbal and non-verbal expression: studying baggage loss in the airport. *Personality and Social Psychology Bulletin* 26(3), 327–339 (2000)

44. Batliner, A., Fischer, K., Huber, R., Spilker, J., Nöth, E.: How to find trouble in communication. *Journal of Speech Communication* 40(1), 117–143 (2003)
45. Batliner, A., Zeißler, V., Frank, C., Adelhardt, J., Shi, R., Nöth, E.: We are not amused-but how do you know? User states in a multi-modal dialogue system. In: *Proceedings of the European Conference on Speech Communication and Technology, Eurospeech* (2003)
46. Batliner, A., Hacker, C., Steidl, S., Nöth, E., Haas, J.: From Emotion to Interaction: Lessons from Real Human-Machine-Dialogues. In: André, E., Dybkjær, L., Minker, W., Heisterkamp, P. (eds.) *ADS 2004. LNCS (LNAI)*, vol. 3068, pp. 1–12. Springer, Heidelberg (2004)
47. Esposito, A.: The amount of information on emotional states conveyed by the verbal and nonverbal channels: Some perceptual data. *Journal of Progress in Nonlinear Speech Processing*, 249–268 (2007)
48. Schuller, B., Seppi, D., Batliner, A., Maier, A., Steidl, S.: Towards more reality in the recognition of emotional speech. In: *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, vol. 4, pp. 941–944. IEEE (2007)
49. Wendt, C., Popp, M., Karg, M., Kuhnlenz, K.: Physiology and HRI: Recognition of over-and underchallenge. In: *Proceedings of the Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 448–452. IEEE (2008)
50. Batliner, A., Steidl, S., Schuller, B., Seppi, D., Vogt, T., Wagner, J., Devillers, L., Vidrascu, L., Aharonson, V., Kessous, L., et al.: Whodunnit-searching for the most important feature types signalling emotion-related user states in speech. *Journal of Computer Speech & Language* 25(1), 4–28 (2011)
51. Tian, Y., Kanade, T., Cohn, J.: Recognizing action units for facial expression analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23(2), 97–115 (2001)
52. Lee, C., Yildirim, S., Bulut, M., Kazemzadeh, A., Busso, C., Deng, Z., Lee, S., Narayanan, S.: Emotion recognition based on phoneme classes. In: *Proceedings of the Annual Conference of the International Speech Communication Association (ISCA), Interspeech* (2004)
53. Wagner, J., Vogt, T., André, E.: A Systematic Comparison of Different HMM Designs for Emotion Recognition from Acted and Spontaneous Speech. In: Paiva, A.C.R., Prada, R., Picard, R.W. (eds.) *ACII 2007. LNCS*, vol. 4738, pp. 114–125. Springer, Heidelberg (2007)
54. Vlasenko, B., Schuller, B., Wendemuth, A., Rigoll, G.: Frame vs. Turn-Level: Emotion Recognition from Speech Considering Static and Dynamic Processing. In: Paiva, A.C.R., Prada, R., Picard, R.W. (eds.) *ACII 2007. LNCS*, vol. 4738, pp. 139–147. Springer, Heidelberg (2007)
55. Maganti, H.K., Scherer, S., Palm, G.: A Novel Feature for Emotion Recognition in Voice Based Applications. In: Paiva, A.C.R., Prada, R., Picard, R.W. (eds.) *ACII 2007. LNCS*, vol. 4738, pp. 710–711. Springer, Heidelberg (2007)
56. Kanade, T., Cohn, J., Tian, Y.: Comprehensive database for facial expression analysis. In: *Proceedings of the International Conference on Automatic Face and Gesture Recognition (FG)*, pp. 46–53. IEEE (2000)
57. Wendt, B., Scheich, H.: The “Magdeburger Prosodie-Korpus”. In: *Proceedings of the International Conference on Speech Prosody*, pp. 699–701 (2002)

58. Burkhardt, F., Paeschke, A., Rolfes, M., Sendlmeier, W., Weiss, B.: A database of German emotional speech. In: Proceedings of the European Conference on Speech Communication and Technology, Eurospeech (2005)
59. Bänziger, T., Scherer, K.R.: Using Actor Portrayals to Systematically Study Multimodal Emotion Expression: The GEMEP Corpus. In: Paiva, A.C.R., Prada, R., Picard, R.W. (eds.) ACII 2007. LNCS, vol. 4738, pp. 476–487. Springer, Heidelberg (2007)
60. Douglas-Cowie, E., Campbell, N., Cowie, R., Roach, P.: Emotional speech: Towards a new generation of databases. *Journal of Speech Communication* 40(1), 33–60 (2003)
61. Douglas-Cowie, E., Cowie, R., Sneddon, I., Cox, C., Lowry, O., McRorie, M., Martin, J.-C., Devillers, L., Abrilian, S., Batliner, A., Amir, N., Karpouzis, K.: The HUMAINE Database: Addressing the Collection and Annotation of Naturalistic and Induced Emotional Data. In: Paiva, A.C.R., Prada, R., Picard, R.W. (eds.) ACII 2007. LNCS, vol. 4738, pp. 488–500. Springer, Heidelberg (2007)
62. Campbell, N.: Databases of expressive speech. *Journal of Chinese Language and Computing* 14(4), 295–304 (2004)
63. Campbell, N.: Listening between the lines: A study of paralinguistic information carried by tone-of-voice. In: Proceedings of the International Symposium on Tonal Aspects of Languages: With Emphasis on Tone Languages (2004)
64. Campbell, N., Sadanobu, T., Imura, M., Iwahashi, N., Noriko, S., Douxchamps, D.: A multimedia database of meetings and informal interactions for tracking participant involvement and discourse flow. In: Proceedings of the International Conference on Language Resources and Evaluation, LREC (2006)
65. Cowie, R., Douglas-Cowie, E., Cox, C.: Beyond emotion archetypes: Databases for emotion modelling using neural networks. *Journal of Neural Networks* 18(4), 371–388 (2005)
66. Martin, O., Kotsia, I., Macq, B., Pitas, I.: The eNTERFACE'05 audio-visual emotion database. In: Proceedings of the International Conference on Data Engineering Workshops (ICDE), p. 8. IEEE (2006)
67. Grimm, M., Kroschel, K., Narayanan, S.: The Vera am Mittag German audio-visual emotional speech database. In: Proceedings of the International Conference on Multimedia and Expo (ICME), pp. 865–868. IEEE (2008)
68. Strauß, P.M., Hoffmann, H., Minker, W., Neumann, H., Palm, G., Scherer, S., Traue, H., Weidenbacher, U.: The PIT corpus of german multi-party dialogues. In: Proceedings of the International Conference on Language Resources and Evaluation, LREC (2008)
69. McKeown, G., Valstar, M., Cowie, R., Pantic, M.: The SEMAINE corpus of emotionally coloured character interactions. In: Proceedings of the International Conference on Multimedia and Expo (ICME), pp. 1079–1084. IEEE (2010)
70. Oertel, C., Scherer, S., Wagner, P., Campbell, N.: On the use of multimodal cues for the prediction of involvement in spontaneous conversation. In: Proceedings of the Annual Conference of the International Speech Communication Association (ISCA), Interspeech, pp. 1541–1544. ISCA (2011)
71. Caridakis, G., Castellano, G., Kessous, L., Raouzaoui, A., Malatesta, L., Asteriadis, S., Karpouzis, K.: Multimodal Emotion Recognition from Expressive Faces, Body Gestures and Speech. In: Boukis, C., Pnevmatikakis, L., Polymenakos, L. (eds.) Artificial Intelligence and Innovations 2007: from Theory to Applications. IFIP, vol. 247, pp. 375–388. Springer, Boston (2007)

72. Batliner, A., Steidl, S., Hacker, C., Nöth, E.: Private emotions versus social interaction: A data-driven approach towards analysing emotion in speech. *Journal of User Modeling and User-Adapted Interaction* 18(1), 175–206 (2008)
73. Zeng, Z., Pantic, M., Roisman, G., Huang, T.: A survey of affect recognition methods: Audio, visual, and spontaneous expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31(1), 39–58 (2009)
74. Glodek, M., Tschechne, S., Layher, G., Schels, M., Brosch, T., Scherer, S., Kächele, M., Schmidt, M., Neumann, H., Palm, G., Schwenker, F.: Multiple Classifier Systems for the Classification of Audio-Visual Emotional States. In: D’Mello, S., Graesser, A., Schuller, B., Martin, J.-C. (eds.) *ACII 2011, Part II. LNCS*, vol. 6975, pp. 359–368. Springer, Heidelberg (2011)
75. Walter, S., Scherer, S., Schels, M., Glodek, M., Hrabal, D., Schmidt, M., Böck, R., Limbrecht, K., Traue, H., Schwenker, F.: Multimodal Emotion Classification in Naturalistic User Behavior. In: Jacko, J.A. (ed.) *HCI International 2011, Part III. LNCS*, vol. 6763, pp. 603–611. Springer, Heidelberg (2011)
76. Schels, M., Glodek, M., Meudt, S., Schmidt, M., Hrabal, D., Böck, R., Walter, S., Schwenker, F.: Multi-modal classifier-fusion for the classification of emotional states in WOZ scenarios. In: *Proceeding of the International Conference on Affective and Pleasurable Design, APD* (2012)
77. Scherer, S., Glodek, M., Schwenker, F., Campbell, N., Palm, G.: Spotting laughter in natural multiparty conversations: A comparison of automatic online and offline approaches using audiovisual data. *ACM Transactions on Interactive Intelligent System* 2(1), 4:1–4:31 (2012)
78. Schuller, B., Valstar, M., Eyben, F., McKeown, G., Cowie, R., Pantic, M.: AVEC 2011—The First International Audio/Visual Emotion Challenge. In: D’Mello, S., Graesser, A., Schuller, B., Martin, J.-C. (eds.) *ACII 2011, Part II. LNCS*, vol. 6975, pp. 415–424. Springer, Heidelberg (2011)
79. Douglas-Cowie, E., Cox, C., Martin, J., Devillers, L., Cowie, R., Sneddon, I., McRorie, M., Pelachaud, C., Peters, C., Lowry, O., et al.: The HUMAINE database. In: *Emotion-Oriented Systems: The Humaine Handbook*, pp. 243–284. Springer (2011)
80. Frommer, J., Michaelis, B., Rösner, D., Wendemuth, A., Friesen, R., Haase, M., Kunze, M., Andrich, R., Lange, J., Panning, A., Siegert, I.: Towards emotion and affect detection in the multimodal LAST MINUTE corpus. In: *Proceedings of the International Conference on Language Resources and Evaluation, LREC* (2012)
81. Gnjatovic, M., Rösner, D.: On the role of the Nimitex corpus in developing an emotion adaptive spoken dialogue system. In: *Proceedings of the International Conference on Language Resources and Evaluation Conference, LREC* (2008)
82. Strauß, P.M., Hoffmann, H., Minker, W., Neumann, H., Palm, G., Scherer, S., Traue, H., Weidenbacher, U.: The PIT corpus of german multi-party dialogues. In: *Proceedings of the International Conference on Language Resources and Evaluation, LREC* (2008)
83. Douglas-Cowie, E., Cowie, R., Cox, C., Amir, N., Heylen, D.: The sensitive artificial listener: An induction technique for generating emotionally coloured conversation. In: *Proceedings of the Workshop Corpora for Research on Emotion and Affect at the International Conference on Language Resources and Evaluation, LREC* (2008)
84. Türk, U.: The technical processing in Smartkom data collection: A case study. In: *Proceedings of the European Conference on Speech Communication and Technology, Eurospeech* (2001)

85. Scherer, S., Siegert, I., Bigalke, L., Meudt, S.: Developing an expressive speech labeling tool incorporating the temporal characteristics of emotion. In: Proceedings of the International Conference on Language Resources and Evaluation (LREC). European Language Resources Association, ELRA (2010)
86. Böck, R., Siegert, I., Haase, M., Lange, J., Wendemuth, A.: ikannotate – A Tool for Labelling, Transcription, and Annotation of Emotionally Coloured Speech. In: D’Mello, S., Graesser, A., Schuller, B., Martin, J.-C. (eds.) ACII 2011, Part I. LNCS, vol. 6974, pp. 25–34. Springer, Heidelberg (2011)
87. Siegert, I., Böck, R., Philippou-Hübner, D., Vlasenko, B., Wendemuth, A.: Appropriate emotional labeling of non-acted speech using basic emotions, Geneva emotion wheel and self assessment manikins. In: Proceedings of the International Conference on Multimedia and Expo (ICME). IEEE (2011)
88. Dietrich, C., Schwenker, F., Palm, G.: Classification of time series utilizing temporal and decision fusion. *Journal of Multiple Classifier Systems*, 378–387 (2001)
89. Dietrich, C., Palm, G., Schwenker, F.: Decision templates for the classification of bioacoustic time series. *Journal of Information Fusion* 4(2), 101–109 (2003)
90. Dietrich, C., Palm, G., Riede, K., Schwenker, F.: Classification of bioacoustic time series based on the combination of global and local decisions. *Journal of Pattern Recognition* 37(12), 2293–2305 (2004)
91. Brown, G., Wyatt, J., Harris, R., Yao, X.: Diversity creation methods: A survey and categorisation. *Journal of Information Fusion* 6(1), 5–20 (2005)
92. Kahsay, L., Schwenker, F., Palm, G.: Comparison of multiclass SVM decomposition schemes for visual object recognition. *Journal of Pattern Recognition*, 334–341 (2005)
93. Mayer, G., Utz, H., Palm, G.: Information integration in a multi-stage object classifier. *Journal of Autonome Mobile Systeme*, 211–217 (2005)
94. Plumpton, C.O., Kuncheva, L.I., Oosterhof, N.N., Johnston, S.J.: Naive random subspace ensemble with linear classifiers for real-time classification of fMRI data. *Journal of Pattern Recognition* 45(6), 2101–2108 (2012)
95. Schwenker, F., Kestler, H., Palm, G.: Three learning phases for radial-basis-function networks. *Journal of Neural Networks* 14(4-5), 439–458 (2001)
96. Scherer, S., Schwenker, F., Palm, G.: Classifier fusion for emotion recognition from speech. In: Proceedings of the International Conference on Intelligent Environments (IE), pp. 152–155 (2007)
97. Schels, M., Thiel, C., Schwenker, F., Palm, G.: Classifier Fusion Applied to Facial Expression Recognition: An Experimental Comparison. In: Ritter, H., Sagerer, G., Dillmann, R., Buss, M. (eds.) *Human Centered Robot Systems. Cognitive Systems Monographs*, vol. 6, pp. 121–129. Springer, Heidelberg (2009)
98. Scherer, S., Schwenker, F., Palm, G.: Classifier fusion for emotion recognition from speech. In: Proceeding of the International Conference Intelligent Environments (IE), pp. 95–117. Springer (2009)
99. Scherer, S., Trentin, E., Schwenker, F., Palm, G.: Approaching emotion in human computer interaction. In: Proceedings of the International Workshop on Spoken Dialogue Systems (IWSDS), pp. 156–168 (2009)
100. Scherer, S., Schwenker, F., Campbell, N., Palm, G.: Multimodal Laughter Detection in Natural Discourses. In: Ritter, H., Sagerer, G., Dillmann, R., Buss, M. (eds.) *Human Centered Robot Systems. Cognitive Systems Monographs*, vol. 6, pp. 111–120. Springer, Heidelberg (2009)
101. Schuller, B., Vlasenko, B., Eyben, F., Wöllmer, M., Stuhlsatz, A., Wendemuth, A., Rigoll, G.: Cross-corpus acoustic emotion recognition: Variances and strategies. *IEEE Transactions on Affective Computing* 1(2), 119–131 (2010)

102. Wöllmer, M., Metallinou, A., Eyben, F., Schuller, B., Narayanan, S.: Context-sensitive multimodal emotion recognition from speech and facial expression using bidirectional lstm modeling. In: Proceedings of the Annual Conference of the International Speech Communication Association (ISCA), Interspeech, pp. 2362–2365 (2010)
103. Esparza, J., Scherer, S., Schwenker, F.: Studying Self- and Active-Training Methods for Multi-feature Set Emotion Recognition. In: Schwenker, F., Trentin, E. (eds.) PSL 2011. LNCS, vol. 7081, pp. 19–31. Springer, Heidelberg (2012)
104. Blum, A., Mitchell, T.: Combining labeled and unlabeled data with co-training. In: Proceedings of the Annual Conference on Computational Learning Theory (COLT), pp. 92–100. Morgan Kaufmann (1998)
105. Bennet, K., Demiriz, A., Maclin, R.: Exploiting unlabeled data in ensemble methods. In: Proceedings of the International Conference on Knowledge Discovery and Data Mining (KDD), pp. 289–296 (2002)
106. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. *Journal of Advances in Neural Information Processing Systems* 17, 529–536 (2005)
107. Lawrence, N., Jordan, M.: Semi-supervised learning via Gaussian processes. *Journal of Advances in Neural Information Processing Systems* 17, 753–760 (2005)
108. Schwenker, F., Dietrich, C., Thiel, C., Palm, G.: Learning of decision fusion mappings for pattern recognition. *Proceedings of the International Journal on Artificial Intelligence and Machine Learning (AIML)* 6, 17–21 (2006)
109. Thiel, C., Scherer, S., Schwenker, F.: Fuzzy-Input Fuzzy-Output One-Against-All Support Vector Machines. In: Apolloni, B., Howlett, R.J., Jain, L. (eds.) KES 2007, Part III. LNCS (LNAI), vol. 4694, pp. 156–165. Springer, Heidelberg (2007)
110. Wang, W., Zhou, Z.-H.: Analyzing Co-training Style Algorithms. In: Kok, J.N., Koronacki, J., Lopez de Mantaras, R., Matwin, S., Mladenič, D., Skowron, A. (eds.) ECML 2007. LNCS (LNAI), vol. 4701, pp. 454–465. Springer, Heidelberg (2007)
111. Abdel Hady, M.F., Schwenker, F.: Decision Templates Based RBF Network for Tree-Structured Multiple Classifier Fusion. In: Benediktsson, J.A., Kittler, J., Roli, F. (eds.) MCS 2009. LNCS, vol. 5519, pp. 92–101. Springer, Heidelberg (2009)
112. Abdel Hady, M.F., Schels, M., Schwenker, F., Palm, G.: Semi-supervised Facial Expressions Annotation Using Co-Training with Fast Probabilistic Tri-Class SVMs. In: Diamantaras, K., Duch, W., Iliadis, L.S. (eds.) ICANN 2010, Part II. LNCS, vol. 6353, pp. 70–75. Springer, Heidelberg (2010)
113. Abdel Hady, M.F., Schwenker, F., Palm, G.: Semi-supervised learning for tree-structured ensembles of RBF networks with co-training. *Neural Networks* 23(4), 497–509 (2010)
114. Abdel Hady, M.F., Schwenker, F., Palm, G.: When classifier selection meets information theory: A unifying view. In: Proceedings of the International Conference on Soft Computing and Pattern Recognition (SoCPar), pp. 314–319. IEEE (2010)
115. Abdel Hady, M., Schwenker, F.: Combining committee-based semi-supervised learning and active learning. *Journal of Computer Science and Technology (JCST): Special Issue on Advances in Machine Learning and Applications* 25(4), 681–698 (2010)
116. Abdel Hady, M., Schwenker, F., Palm, G.: Semi-supervised learning for tree-structured ensembles of RBF networks with co-training. *Journal of Neural Networks* 23(4), 497–509 (2010)

117. Zhu, X.: Semi-supervised learning literature survey. Technical Report 1530, Department of Computer Sciences, University of Wisconsin at Madison (2008)
118. Settles, B.: Active learning literature survey. Technical Report 1648, Department of Computer Sciences, University of Wisconsin-Madison (2009)
119. El Gayar, N., Schwenker, F., Palm, G.: A study of the robustness of KNN classifiers trained using soft labels. *Journal of Artificial Neural Networks in Pattern Recognition*, 67–80 (2006)
120. Thiel, C., Sonntag, B., Schwenker, F.: Experiments with Supervised Fuzzy LVQ. In: Prevost, L., Marinai, S., Schwenker, F. (eds.) *ANNPR 2008. LNCS (LNAI)*, vol. 5064, pp. 125–132. Springer, Heidelberg (2008)
121. Thiel, C., Giacco, F., Schwenker, F., Palm, G.: Comparison of neural classification algorithms applied to land cover mapping. In: *Proceeding of the International Conference on New Directions in Neural Networks: Italian Workshop on Neural Networks: (WIRN)*, pp. 254–263. IOS Press (2009)

Towards Causal Modeling of Human Behavior

Matteo Campo¹, Anna Polychroniou¹, Hugues Salamin¹,
Maurizio Filippone¹, and Alessandro Vinciarelli^{1,2}

¹ University of Glasgow - Sir A. Williams Bldg. - G12 8QQ Glasgow (UK)

² Idiap Research Institute - CP 592 - 1920 Martigny (Switzerland)

firstname.lastname@glasgow.ac.uk

Abstract. This article proposes experiments on decision making based on the “Winter Survival Task”, one of the scenarios most commonly applied in behavioral and psychological studies. The goal of the Task is to identify, out of a predefined list of 12 items, those that are most likely to increase the chances of survival after the crash of a plane in a polar area. In our experiments, 60 pairs of unacquainted individuals (120 subjects in total) negotiate a common choice of the items to be retained after that each subject has performed the task individually. The results of the negotiations are analyzed in causal terms and show that the choices made by the subjects individually act as a causal factor with respect to the outcome of the negotiation.

1 Introduction

In the last years, automatic analysis of human behavior has attracted a large deal of attention in the computing community (see [1,2] for extensive surveys). The efforts focused on two main directions, namely (i) the synthesis of human behavior - in particular when it comes to social and affective phenomena that make embodied conversational agents believable, and (ii) the automatic understanding of human communication dynamics, with particular attention to the prediction of behavioral outcomes and the inference of socially relevant information from nonverbal communication.

Current approaches tend to adopt a purely computational perspective, i.e. they do not try to understand the phenomena they synthesize or analyze, but simply to maximize performance metrics like the recognition rate (percentage of times an approach makes the correct prediction) or the Mean Opinion Score (average appreciation score assigned by users). Such a perspective is certainly effective, but the interdisciplinary collaboration with human sciences, inevitable when dealing with human behavior, shows that no technology can be effective in the field without understanding human-human and human-machine interactions in terms of causes and effects [3,4,5].

The statistical literature has studied extensively approaches aimed at learning cause-effect relationships from data (see [6] for an extensive survey). However, these approaches were largely neglected in the computing community, in part




Fig. 1. The plot shows the percentage of calls where there was a disagreement on each item

because they require the formulation of untested causal assumptions about the phenomena under analysis, in part because they adopt notations and terminology different from those commonly used in the machine intelligence community [7]. This paper shows that overcoming these two barriers can be beneficial to the computational analysis of social interactions from at least two points of view. The first is that it makes social phenomena more predictable and, hence, easier to analyze automatically; the second is that it allows a better understanding of the data being modeled.

2 The Experiment

The experiment adopts the “*Winter Survival Task*”, a scenario where participants must identify, out of a list of 12 predefined items, those that maximize the chances of survival after an emergency landing in Northern Canada (in the middle of the winter). The main advantage of the scenario, often used in psychological and behavioral experiments, is that the average subject is unlikely to have experienced a plane crash or to know survival techniques suitable for a winter beyond the Polar Circle. Hence, the outcomes of the experiment are likely to depend on social and psychological phenomena during the interactions and not on skills and knowledge the participants have before and independently of the experiment.

In this work, the participants earn three British Pounds each time they select a correct item (there is a gold standard for the task), but lose the same amount of money when they select a wrong one. In this way, whenever the subjects disagree about a certain item, they are motivated to persuade their interlocutor.

2.1 Experimental Protocol

The task was performed by 60 pairs of fully unacquainted subjects that have never met before the experiment (120 subjects in total). For each pair, the protocol included the following steps:

- The two subjects are accompanied to different rooms without meeting or crossing one another.
- Once in their room, the subjects receive a mobile phone (same model for all participants) and the documentation accompanying the experiment (scenario, questionnaires, etc.), including the list of the 12 items at the core of the Winter Survival Task.
- Before starting the call, the subjects fill a form where they must write a decision (“*Yes*” or “*No*”) for each item of the list. This makes it possible to know, for each item, what is the decision made by each subject before any interaction with their counterpart.
- One of the two subjects, selected randomly, calls the other and starts the discussion item per item.
- During the call, the subjects discuss item by item and negotiate a common solution (“*Yes*” or “*No*”) that is the final outcome of the task.

At the end of the call, it is possible to know what are the items on which the participants disagree and, most importantly, what are the subjects that persuade their interlocutors in case of disagreement, i.e. the subjects that convince others to adopt their initial decision in case of disagreement. Figure 1 shows, for each item, the percentage of calls where discussion was needed to reach a common decision. While some items (e.g., the clothing) were discussed only a few times, others were frequently debated between participants.

3 Causal Analysis

The main question behind the experiment is what are the causal factors that increase the chances of persuading others. In other words, whether the decision about a given item is random (which is what the scenario seems to suggest), it depends on the characteristics of the subjects, or on their choice prior to the discussion. The next sections show how the problem was modeled and the resulting findings.

3.1 Modeling

The problem can be modeled with a set \mathbf{X} of observable binary variables:

Role \mathbf{R} : $R \in \{0, 1\}$
 Gender \mathbf{G} : $G \in \{0, 1\}$
 Initial Choice \mathbf{Y} : $Y \in \{0, 1\}$
 Result \mathbf{W} : $W \in \{0, 1\}$

The variable R stands for the role a participant had in the conversation in terms of “*caller*”, the subjects who makes the phone call, or “*receiver*”, the subject who receives the call (associated to 1 and 0, respectively). The variable G corresponds to the gender of the participant, with 0 and 1 for male and female, respectively. The variable Y accounts for the initial decision of the participants, with 0 and 1 corresponding to “*No*” and “*Yes*”, respectively. Finally, the variable W accounts for whether the subject persuades the other person ($W = 1$) or not ($W = 0$).

The aim of this work is to make statements about the causal relationship between Y and W . In particular, we would like to estimate the following quantities:

Post-intervention distribution: The post-intervention distribution estimates the probability of a subject persuading or not the other person given that the “treatment” $Y = 1$ is applied (i.e., given that the initial decision of the subject is “*Yes*”):

$$p(W|\text{do}(Y = 1))$$

where $\text{do}(Y = 1)$ is the “do” operator [8] and expresses the probability of an effect given an “action” on the model (the action will be, for our purposes, taking an initial choice Y). It is important to point out that this quantity has a different and stronger meaning than the probability of an effect given an “observation”, which is the usual conditional probability.

Counterfactuals: The probabilities of a change of effect, given a change of treatment. These can be expressed in terms of the following distributions:

$$\text{PN} = \frac{p(W = 1) - p(W = 1|\text{do}(Y = 0))}{p(W = 1, Y = 1)}$$

$$\text{PS} = \frac{p(W = 1|\text{do}(Y = 1)) - p(W = 1)}{p(W = 0, Y = 0)}$$

$$\text{PNS} = p(W = 1|\text{do}(Y = 1)) - p(W = 1|\text{do}(Y = 0))$$

that account, respectively, for the probability of Y being a necessary cause of W , the probability of Y being a sufficient cause of W , and the probability of Y being a necessary and sufficient cause of W .

Within the framework of causal inference, causal dependencies between variables are encoded by means of Direct Acyclic Graphs (DAG) [8]. Therefore, DAGs represent assumptions on the causal relationships between variables that can then be inferred after data are observed. Once the DAG is built, the first step to undertake in order estimate causal effects and counterfactuals is to verify whether these can actually be computed: this problem is called “identifiability” [8]. At the core of identifiability in DAG is the d-separation criterion, by which it is possible to derive conditional independence relationships between variables. In our experiments, d-separation was checked using TETRAD, a publicly available software package dealing with causal models [9].




Fig. 2. The picture shows the Directed Acyclic Graph corresponding to the causal assumptions behind the experiments of this work

3.2 The Causal Model

This section shows the causal model used for the experiments of this work. The corresponding DAG is depicted in Figure 2 and the underlying causal assumptions are as follows:

- An unobserved variable U_1 influences the role of the participants;
- Role and gender influence the initial decision of the subjects and the result of the discussion;
- An unobserved variable U_2 influences the initial choice;
- The initial choice influences the result of the discussion;
- An unobserved variable U_3 influences the result of the discussion.

Unobserved variables are assumed to be deterministically related to their children and mutually independent. According to the graph (and its underlying assumptions), the joint probability distribution of the four observed variables is as follows:

$$p(\mathbf{X}) = p(R) p(G) p(Y|G, R) p(W|G, Y, R) . \tag{1}$$

Table 1. The table reports the Maximum Likelihood estimate of the joint probability distribution of \mathbf{X} as obtained from the data observed in the 60 conversations used in the experiments

<table style="width: 100%; border-collapse: collapse;"> <tr> <td></td> <td style="text-align: center; border-bottom: 1px solid black;">Y = 0</td> <td style="text-align: center; border-bottom: 1px solid black;">Y = 1</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">G = 0</td> <td style="text-align: center;">0.036</td> <td style="text-align: center;">0.021</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">G = 1</td> <td style="text-align: center;">0.026</td> <td style="text-align: center;">0.016</td> </tr> <tr> <td></td> <td colspan="2" style="text-align: center;">(R = 0, W = 0)</td> </tr> </table>		Y = 0	Y = 1	G = 0	0.036	0.021	G = 1	0.026	0.016		(R = 0, W = 0)		<table style="width: 100%; border-collapse: collapse;"> <tr> <td></td> <td style="text-align: center; border-bottom: 1px solid black;">Y = 0</td> <td style="text-align: center; border-bottom: 1px solid black;">Y = 1</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">G = 0</td> <td style="text-align: center;">0.019</td> <td style="text-align: center;">0.028</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">G = 1</td> <td style="text-align: center;">0.038</td> <td style="text-align: center;">0.028</td> </tr> <tr> <td></td> <td colspan="2" style="text-align: center;">(R = 1, W = 0)</td> </tr> </table>		Y = 0	Y = 1	G = 0	0.019	0.028	G = 1	0.038	0.028		(R = 1, W = 0)	
	Y = 0	Y = 1																							
G = 0	0.036	0.021																							
G = 1	0.026	0.016																							
	(R = 0, W = 0)																								
	Y = 0	Y = 1																							
G = 0	0.019	0.028																							
G = 1	0.038	0.028																							
	(R = 1, W = 0)																								
<table style="width: 100%; border-collapse: collapse;"> <tr> <td></td> <td style="text-align: center; border-bottom: 1px solid black;">Y = 0</td> <td style="text-align: center; border-bottom: 1px solid black;">Y = 1</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">G = 0</td> <td style="text-align: center;">0.071</td> <td style="text-align: center;">0.123</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">G = 1</td> <td style="text-align: center;">0.075</td> <td style="text-align: center;">0.134</td> </tr> <tr> <td></td> <td colspan="2" style="text-align: center;">(R = 0, W = 1)</td> </tr> </table>		Y = 0	Y = 1	G = 0	0.071	0.123	G = 1	0.075	0.134		(R = 0, W = 1)		<table style="width: 100%; border-collapse: collapse;"> <tr> <td></td> <td style="text-align: center; border-bottom: 1px solid black;">Y = 0</td> <td style="text-align: center; border-bottom: 1px solid black;">Y = 1</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">G = 0</td> <td style="text-align: center;">0.062</td> <td style="text-align: center;">0.141</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">G = 1</td> <td style="text-align: center;">0.064</td> <td style="text-align: center;">0.120</td> </tr> <tr> <td></td> <td colspan="2" style="text-align: center;">(R = 1, W = 1)</td> </tr> </table>		Y = 0	Y = 1	G = 0	0.062	0.141	G = 1	0.064	0.120		(R = 1, W = 1)	
	Y = 0	Y = 1																							
G = 0	0.071	0.123																							
G = 1	0.075	0.134																							
	(R = 0, W = 1)																								
	Y = 0	Y = 1																							
G = 0	0.062	0.141																							
G = 1	0.064	0.120																							
	(R = 1, W = 1)																								

3.3 Post-Intervention Distribution

The model is *Markovian* because the associated graph is undirected and acyclic and the unobserved variables U_1, U_2, U_3 are mutually independent. This guarantees the causal effect $p(W|\text{do}(Y))$ to be identifiable. Performing an “intervention” on variable Y , that is, using the do-operator $\text{do}(Y = 1)$, the *post-intervention* distribution is:

$$p(W, R, G|\text{do}(Y = 1)) = p(R) p(G) p(W|Y = 1, R, G) .$$

Therefore, the following holds:

$$p(W = 1|\text{do}(Y = 1)) = \sum_{g,r \in \{0,1\}} p(r) p(g) p(W = 1|Y = 1, r, g) = 0.849$$

$$p(W = 1|\text{do}(Y = 0)) = \sum_{g,r \in \{0,1\}} p(r) p(g) p(W = 1|Y = 0, r, g) = 0.700 .$$

The difference between the two probabilities is:

$$p(W = 1|\text{do}(Y = 1)) - p(W = 1|\text{do}(Y = 0)) = 0.149$$

and it expresses the difference between the effects of two different treatments. Note that the quantity above has also a counterfactual interpretation, as it is the probability of Y being a necessary and sufficient cause of W (PNS). A binomial test proportion with null hypothesis for $p(W = 1|\text{do}(Y = 1))$ and $p(W = 1|\text{do}(Y = 0))$ to be binomially distributed with the same success probability shows that the observed difference is statistically significant with p -value lower than $2 \cdot 10^{-4}$. This suggests that the value of Y acts as a causal factor and starting from an initial positive decision significantly increases the chances of persuading the counterpart.

3.4 Counterfactuals

The observed data show that the following relationship holds for the interactions used in the experiments:

$$p(W = 1|\text{do}(Y = 1)) \geq p(W = 1) \geq p(W = 1|\text{do}(Y = 0))$$

The relation above corresponds to the *monotonicity* of W relative to Y and it guarantees the identifiability of the three counterfactuals PN, PS, PNS [8]. According to the data, the values of the counterfactuals are:

$$\text{PN} = 0.171$$

$$\text{PS} = 0.503$$

The values are particularly interesting as they show the probability of obtaining a different result were the initial choices different. In particular, PN gives the

probability that changing the initial choice from $Y = 1$ to $Y = 0$ would have changed the result from $W = 1$ to $W = 0$. Conversely, PS gives the probability that changing the initial choice from $Y = 0$ to $Y = 1$ would have changed the result from $W = 0$ to $W = 1$. The latter represents the probability of Y being a sufficient cause of W , and suggests the interesting conclusion that in half of the cases where the initial choice $Y = 0$ led to $W = 0$, a different initial choice $Y = 1$ would have turned the result to $W = 1$.

4 Conclusions

This paper has presented a causal analysis of the decision-making behavior of individuals involved in the “Winter Survival Task”. The experiments have involved 120 subjects and show that, when it comes to binary decisions about the acceptance of an item in the task, the initial choice of a subject acts as a causal factor for the final outcome of the discussion. In particular, subjects that start with an initial positive decision (“*the item should be retained*”) have a probability of persuading others three times higher than the subjects starting with a negative decision (“*the item should not be retained*”).

While being relatively simple, the experiments involve a large number of subjects that allow one to reliably estimate the causal effects. The main difference with respect to the application of traditional, associative statistics is that the estimated probabilities do not simply tell how frequently two or more variables take certain values, but what is the probability that one or more variables cause the value of one or more other variables. In this respect, the application promises to be fruitful not only from a technological point of view, making the prediction of interaction outcomes easier, but also from a scientific point of view, providing explanations about the observed results.

In the case of these experiments, all observations of interest could be modeled with binary variables, but in real-world scenarios, variables of interest are more likely to be continuous. This does not represent a major problem because all equations used in this work are independent of the actual expression used to estimate the probabilities. In other words, the tables used in these experiments can be replaced by distributions as complex as necessary without changing model assumptions, identifiability considerations, and formulas estimating the post-intervention probabilities or counterfactuals.

Future work will focus on those cases where causal modeling can actually make a major difference with respect to associative statistics, i.e. the analysis of non-experimental data. Those are data where conditions cannot be manipulated (as it typically happens in naturalistic settings for human-human and human-machine interactions), and hence, major effects might be missed by associative statistics simply because they are less frequent.

Acknowledgements. The Authors wish to acknowledge the support from the College of Science and Engineering of the University of Glasgow, the FP7 funded European Network of Excellence SSPNet, the Swiss National Center for Competence in Research on Interactive Multimodal Information Management (IM2) and the project “Human Emotional Interaction” funded by Finnish Ministry for the Technological Innovation (TEKES).

References

1. Vinciarelli, A., Pantic, M., Bourlard, H.: Social Signal Processing: Survey of an emerging domain. *Image and Vision Computing Journal* 27(12), 1743–1759 (2009)
2. Vinciarelli, A., Pantic, M., Heylen, D., Pelachaud, C., Poggi, I., D’Érrico, F., Schroeder, M.: Bridging the gap between social animal and unsocial machine: A survey of social signal processing. *IEEE Transactions on Affective Computing* 3(1), 69–87 (2012)
3. Brunet, P., Cowie, R.: Towards a conceptual framework of research on social signal processing. *Journal on Multimodal User Interfaces* (to appear, 2012)
4. Mehu, M., Scherer, K.: A psycho-ethological approach to social signal processing. *Cognitive Processing* 13(2), 397–414 (2012)
5. Poggi, I., D’Érrico, F.: Social signals: a framework in terms of goals and beliefs. *Cognitive Processing* 13(2), 427–445 (2012)
6. Pearl, J.: Causal inference in statistics: An overview. *Statistics Survey* 3(1), 96–146 (2009)
7. Pearl, J.: Statistics and causal inference: A review. *Test* 12(2), 281–345 (2003)
8. Pearl, J.: *Causality: Models, Reasoning and Inference*. Cambridge University Press (2000)
9. Glymour, C., Scheines, R., Spirtes, P., Kelly, K.: *Discovering causal structure: Artificial intelligence, philosophy of science, and statistical modeling*. Academic Press (1987)

How Social Signal Processing (SSP) Can Help Assessment of Bonding Phenomena in Developmental Psychology?

Emilie Delaherche¹, Sofiane Boucenna¹,
Mohamed Chetouani¹, and David Cohen^{1,2}

¹ Institut des Systèmes Intelligents et de Robotique,
CNRS UMR 7222,

Université Pierre et Marie Curie, Paris, France

² Department of Child and Adolescent Psychiatry,
AP-HP, Groupe Hospitalier Pitie-Salpetriere,
Université Pierre et Marie Curie, Paris, France

{emilie.delaherche@isir.upmc.fr, sofiane.boucenna}@isir.upmc.fr,
mohamed.chetouani@upmc.fr,
david.cohen@psl.aphp.fr

Abstract. In the field of biology, the study of bonding has been renewed by the discovery of non genetic transmission of behavioural traits through early mother-infant interaction and the role of stress hormones and oxytocin. However, the study of early interaction is complex and Social Signal Processing (SSP) can help in addressing some issues. Based on works from our group, we will show data from diverse sources (e.g. experiments, home movies) showing how SSP was used to address synchrony between partners (e.g. infant, child, care giver, agent) and characteristics that participates to interpersonal exchanges (e.g. motherese, emotional prosody or faces).

1 Introduction

Recent advance has shown that human learning and cultural evolution are supported by paradoxical biological adaptation. We are born immature; yet, immaturity has value: "Delaying maturation of cerebral cortex allows initial learning to influence the neural architecture in ways that support later, more complex learning" [28]. Early learning appears to be computational [22], to be based on perceptual-action mapping (meaning that it is supported by brain circuits linking perception and action). Learning is also social and supported by skills present in infancy: imitation, shared attention and empathic understanding [28]. The paradigmatic situation of very early empathic skills is described as bonding that occurs in all species and starts at birth. Assessment of early interaction between an infant and his/her caregiver remains challenging as it requires considering both partners and mutual influences they provide during interaction.

In this paper, our aim is to show that SSP can help in addressing some issues. Based on works from our group, we will show data from diverse sources (e.g. experiments, home movies) showing how SSP was used to address synchrony between partners (e.g. infant, child, care giver, agent) and characteristics that participates to interpersonal exchanges (e.g. motherese, emotional prosody or faces). To do so, we will first describe the bonding phenomena and consider how recent developments in the field of biology have renewed its core importance during development. Second, we will review computational modelling of interpersonal synchrony. Third, we will show that automatic detection of non-verbal cues of interpersonal synchrony may help addressing the complex issue of the emotional implication in interpersonal exchange. Finally, we will propose some prospects in new methods and tools for the study of children development.

2 Description of the Bonding Phenomena

The first evidences of the bonding phenomena came from two sources in the fifties: (1) ethologists studying early development and interaction showed in birds and monkeys, the printing phenomenon; (2) child psychoanalysts studying the impact of severe deprivation and mother-infant early separation described infant depression and its reversibility by providing infants warm and individualised care. Later, Bowlby proposed the attachment theory and provided a theoretical background for many studies trying to understand consequences of early adversities in infants. Early infant-caregiver interaction was then specifically studied to understand bonding, attachment and early development [26].

Despite the pioneering efforts of Denenberg, who first showed the non-genomic transmission of behavioral traits in animals when studying early separation [12], we only recently understood the biological implications of early stress. Using rodent models, Meaney and Champagne showed that early stress, maternal care and stress during the gestation affected the development of future generations of rats through the hypothalamus-pituitary axis (HPA) and epigenetic modifications. These modifications could be transferred from generation to generation and were independent of an animal's genetic inheritance. The following briefly lists some important points learned from these experiments: 1) Early experience has a long-term effect on behavior and the biological system, especially when the mother and offspring are separated or when the quality of maternal care varies dramatically [23]; 2) Certain early experiences can affect future generations, providing a non-genomic mechanism for the transmission of behavioral traits [16]. It appears that maternal care affects development through a behavioral program and the future adult's pathological responses to stress. The quality of maternal care influenced the stress response HPA axes of offspring [23] and greatly influenced the epigenesis in the following generations (through DNA epigenomic marking) [44]. Furthermore, naturally occurring variations in maternal behavior are associated with differences in estrogen-inducible central oxytocin receptors, which are involved in pro-social behaviors [8]. Oxytocin appears to enhance both maternal as well as affiliative behaviors and is considered as the bonding hormone. These developments have pushed developmental psychologists to study

early interaction not only as the addition of two behaviors but rather as a single phenomenon with a dialogue between two partners engaged in behavioral and emotional exchange. Developmental psychologists give now importance to rhythm, synchrony and emotion, regarded as key expression of proper early interaction [15]. However, the study of synchrony and emotion in children interacting with a partner or a caregiver is complex and SSP can help in addressing some issues.

3 Computational Modeling of Interpersonal Synchrony

Synchrony refers to individuals' temporal coordination during social interactions [11,20]. The analysis of this phenomenon is complex, requiring the perception and integration of multimodal communicative signals [11]. The evaluation of synchrony has received multidisciplinary attention because of its role in early development [14], language learning [18] and social connection [19]. Initially, instances of synchrony were directly perceived in the data by trained observers. Several methods have been proposed to evaluate interactional synchrony, ranging from behavior micro-analysis [6] to global perception of synchrony [3]. Synchrony has now captured the interest of researchers in such fields as social signal processing, robotics and machine learning [21,33].

3.1 Fully Automatic Measures of Movement Synchrony

To exploit synchrony cues in human-machine interaction, automatic techniques can be used to capture relevant social signals and assess movement synchrony in human-human interactions. This studies aim at measuring the degree of similarity between the dynamics of the non-verbal behaviors of dyadic partners. The goals of these studies are generally divisible into two categories: (a) compare the degree of synchrony under different conditions (e.g., with or without visual feedback) [39,42] and (b) study the correlation between the degree of synchrony and an outcome variable (e.g., friendship, relationship quality) [1,34].

The first step in computing synchrony is to extract the relevant features of the dyad's motion with motion-tracking devices [2], image-processing techniques (tracking algorithms, image differencing) [9,42] or physiological sensors [42]. After extracting the motion features, a measure of similarity is applied. Correlation is the most commonly used method to assess interactional synchrony [1,34]. A time-lagged cross-correlation is applied between the movement time series of the interactional partners using short windows of interaction. Another method to assess the similarity of motion of two partners is recurrence analysis [35]. Recurrence analysis assesses the points in time that two systems show similar patterns of change or movement, called "recurrence points". Spectral methods constitute an interesting alternative to temporal methods when dealing with rhythmic tasks. Spectral methods measure the evolution of the relative phase between the two partners as an indication of a stable time-lag between them [31,37]. Spectral

methods also measure the overlap between the movement frequencies of the partners, called cross-spectral coherence [9,36,37] or power spectrum overlap [31].

A critical question when attempting to detect dependence relationships between features is the level of significance of the synchrony metrics. A well-spread method consists of applying surrogate statistical testing [2,9,36,40]. Video images of dyadic partners are isolated and re-combined in a random order to synthesize surrogate data (pseudo-interactions). Synchrony scores are assessed using the original and surrogate datasets. The synchrony scores on the surrogate dataset constitute a baseline for judging for the dyad's coordination. Fully automatic measures of movement synchrony are subject to several criticisms in the context of studying naturalistic interaction data. First, the measures provided by these methods are mostly global and do not shed light on what happened locally during the interaction; they do not provide a local model of the communication dynamics. Second, the importance of speech and multimodality is often concealed in these methods. Third, these methods are suitable for analyzing a database but do not provide direct insights on how to equip a machine with such coordination skills.

3.2 Modeling Communication Dynamics

Given these criticisms, many in the field adopted the alternative practice of modeling the timing and occurrence of higher-level behavioral events such as smiles, head gestures, gazes and speaker changes. These behavioral events can be either extracted from a human-annotated database or predicted from low-level signals automatically extracted from data. These methods arise from a great interest in identifying the dynamical patterns of interaction and characterizing recurrent interpersonal behaviors.

Machine learning methods offer an interesting framework for the exploration of interactive behaviors. A key challenge is proposing models with the content and temporal structure of dyadic interactions. Various sequential learning models, such as Hidden Markov Models (HMMs) or Conditional Random Fields (CRFs), are usually used to characterize the temporal structure of social interactions. Messinger et al. employ related techniques for the understanding of communicative development, which is characterized by mutual influences during interaction: infants and parents influence and respond to one another during communication [29]. In section 3.3, we will further develop these prospects in children with autism.

Among interpersonal behaviors, the prediction of turn-taking and back-channels has been largely studied in the perspective of building fluent dialog systems. The central idea is to develop "predictive models of communication dynamics that integrate previous and current actions from all interlocutors to anticipate the most likely next actions of one or all interlocutors" [32]. The purpose of the turn-taking prediction is to accurately predict the timing between speaker transitions and the upcoming type of utterance (speaker holding the floor, speaker changes) as it occurs in human-human interactions [43]. Back-channel behavior assures the speaker that the listener is paying attention and is

in the same state in the conversation [41]. Several teams have investigated how the speaker behavior triggered listeners' back-channels [30].

3.3 Interaction and Communication of Children with Autism Spectrum Disorder and Typically Developing Children

Here, we present two examples of the use of computational analysis to describe interaction and communication of children with autism. In the first experiment, we asked children with autism and TD controls to build a clown in three different situations: (1) the child imitates the speech therapist; (2) the child follows vocal instruction; (3) the child gives vocal instructions. Using automatic extraction of speech turn taking, gestural turn tacking and synchronized motion coupled with machine learning, we found that features characterizing the rhythm of the therapist and the duration of his gestural pauses were particularly adequate to predict the child clinical group. We also found that the performance in the tasks also depended on the age of the child, especially when the child gives instruction to the therapist. The volume of speech interventions, the duration of the therapist interventions and the duration of the therapist gestural pauses were found to be predictive of the age of the child in this task [10].

In the second example, we aimed to assess whether taking into account interaction synchrony would help to better differentiate autism (AD) from typical development (TD) in family home movies (HM) of infants aged less than 18 months. An integrative approach was proposed to explicitly consider the interaction synchrony of behaviors. We estimated transitions between behaviors of the infant and the parent by analyzing behaviors co-occurring in a 3s window. Assuming a Markovian process, we used a maximum likelihood estimation to estimate the probability of each interactive pattern, resulting in bi-gram models characterizing the temporal structure. We also considered the two directions of interaction (Parent→Infant and Infant→Parent). Compared to TD children, infant with AD exhibit a growing deviant development of interactive patterns. Parents of AD did not differ very much from parents of TD when responding to their child. However, when initiating interaction, parents use more touching and intense stimulation as early as the first semester [38].

4 Automatic Detection Non-verbal Cues of Interpersonal Synchrony

As said previously, early interaction is not only based on behavioural cues but also on emotional cues. These appeared to be crucial although assessment in infant and CG is complex. Developmental psychologists have shown that motherese (the way CG talk with their infant) has specific characteristics and plays a key role in early interaction and language learning. Here, we summarized two aspects of automatic detection of emotion based either on audio or video extraction in the context of human-robot interaction and home movies.

4.1 Facial Expressions Assessment through Human-Robot Interaction

We are interested in understanding how babies learn to recognize facial expressions without having a teaching signal allowing to associate a facial expression to a given abstract label (i.e the name of the facial expression 'sadness', 'happiness'...). Our starting point was a mathematical model showing that if the baby uses a sensory motor architecture for the recognition of the facial expression then the parents must imitate the baby facial expression to allow the on-line learning. A first series of robotics experiments showing that a simple neural network model can control the robot head and learn on-line to recognize the facial expressions (the human partner imitates the robot prototypical facial expressions) is presented. We emphasize the importance of the emotions as a mechanism to ensure the dynamical coupling between individuals allowing to learn more complex tasks.

Using the cognitive system algebra [17], we showed that a simple sensory-motor architecture based on a classical conditioning paradigm could learn online to recognize facial expressions if and only if we suppose that the robot produces first facial expressions according to his internal emotional state and that next the parents imitate the facial expression of their robot allowing in return the robot to associate these expressions with his internal state.

4.2 Experimental Set-Up

Using a minimal robotic set-up (Figure 1), we adopt the following experimental protocol: In a first phase of interaction, the robot produces a random facial expression (sadness, happy, anger, surprised) plus the neutral face during 2s, then returns to a neutral face to avoid human misinterpretations of the robot facial expression during 2 sec. The human subject is asked to mimic the robot head. After this first phase lasting between 2 to 3 min according to the subject "patience". The generator of random emotional states is stopped. If the N.N has learned correctly, the robot must be able to mimic the facial expression of the human partener. The computational architecture (Figure 2) allows to recognize the visual features of the people interacting with the robot head and to learn if these features are correlated with its own facial expression.

4.3 Neural Network Model

We use a visual system independent from face framing. The visual system is based on a sequential exploration of the image focus points (Figure 3). The focus points are the result of a DOG filter convolved with the gradient of the input image. This process allows the system to focus more on the corners and end of lines in the image for example eyebrows, corners of the lips, but also distractors (hair, background). Its main advantages over the SIFT (Scale Invariant Feature Transform) [24] method are its computational speed and a fewer extracted focus points (the intensity of the point is directly its level of interest).




Fig. 1. Examples of robot facial expressions: a) sadness, b) surprise, c) happiness. d) Example of a typical human / robot interaction game (here the human imitating the robot).




Fig. 2. The global architecture to recognize facial expression and imitate. A visual processing allows to extract sequentially the local views.




Fig. 3. Visual processing: This visual system is based on a sequential exploration of the image focus points

One after the other, the most active focus points of the same image are used to compute local views: either a log polar¹ transform centered on the focus point is performed to obtain an image more robust to small rotations and distance variations and his radius is 20 pixels, and gabor filters are performed (robust to rotations and distance variations). The features extracted for the convolution between the gabor filter and the focus point are the mean and the standard deviation.

This collection of local views is learned by the recruitment of new neurons (visual features). Of course, there is no constraint on the selection of the local views. This means that numerous distractors can be present (local views in the background, or inexpressive parts of the head). Therefore, distractors can be learned. Nevertheless, the architecture will tend to learn and reinforce only the expressive features of the face (Figure 2). In our face to face situation, the distractors are present for all the facial expressions so their correlation with an emotional state tends toward zero.

A simple conditioning mechanism (the Least Mean Square rule [45]) is able to associate the visual features with the internal state. A sensory-motor architecture learn online to recognize facial expressions if and only if we suppose that the robot produces first facial expressions according to his internal emotional state and that the parents imitate the facial expression of their robot allowing in return the robot to associate these expressions with his internal state.

Arbitrary, a limited amount of time is fixed for the visual exploration of one image. The system succeeds to analyse 10 local views on each image. It is a quite small number of points but since the system usually succeeds to take 3 to 4 relevant points on the face (mouth, eyebrow). Yet, it is enough in most cases and it allows to maintain real time interaction (3 to 5 images/second) in order to test our model.

After learning, the robot head can imitate the human's facial expression and the focus points are associated to each facial expression i.e these focus points vote for the recognition of a given facial expression. Each facial expression is mainly characterized by a specific set of focal points corresponding to local areas on the face which are relevant for the recognition of that expression. For example, some local view around the mouth (lip) characterize the "happyness" facial expressions, some others around the eyebrows characterize the anger facial expression. Figure 4 shows that the model can generalize to people who were not present during the learning phase. A possible explanation for the bad result concerning sadness is that the people have difficulties to display sadness without a context. Each partner imitating the robot displays the sadness in a different way.

Our robot learns thanks to the interaction with a human partner. Our model has allowed us to show that in order to learn online to recognize the facial expressions, the learner must produce facial expressions first and be mimicked by his/her caregiver [4]. The system proposed had no real interaction capability dur-

¹ The local polar transform increases the robustness of the extracted local views to small rotations and scale variations.




Fig. 4. Generalisation to new faces: After 20 persons interacted with the robot head (learning phase), the robot had to imitate new persons never seen

ing the learning phase since this phase was completely predefined. In conclusion, this work suggests the baby/parents system is an autopoietic social system [27] in which the emotional signal and the empathy are important elements of the network to maintain the interaction and to allow the learning of more and more complex skills as the social referencing² [5].

4.4 Motherese: An Emotional Based Process for Sustaining Mother-Infant Interaction

Given the role of motherese in early interaction of typically developing children [13], we aimed to explore whether or not this emotional prosody was implicated in a different way in early interaction of infant who will later develop autism. To do so, we developed an automatic algorithm based on prosodic features to classify motherese versus other speech in Home Movies (HM) [25]. We then assessed the course of infants' responses to parents' vocalisation in the same HM data base described earlier. We found: that parents of infants developing autism displayed more intense solicitations rich in motherese; that motherese increased infant responses towards people and infant receptive behaviours; that fathers of infants developing autism assumed a greater part in vocalisations addressed to infants, and appeared to increase infant intersubjective responses and active behaviours. We conclude that parents of infants who will later develop autism change their interactive pattern of behaviour by both increasing motherese and father's commitment as they improve infant's social responses. Taken together

² The ability to recognize, understand, respond to and alter behavior in response to the emotional expressions of a social partner.

these results stress that parents are aware of the pervasive development of their child, and that they try to adapt long before diagnosis are given [7].

5 Conclusion

We conclude that SSP can help to address some of the issues related to the study of early interaction. SSP can be used for several purposes such as modelling, assessing synchrony between partners and characterizing specific cues that participates to interpersonal exchanges. SSP may also be of interest for developing specific tools with human-like abilities to stimulate social behaviors in a controlled context.

Acknowledgments. This work was supported by the UPMC "Emergence 2009" program, the European Union Seventh Framework Programme under grant agreement n°288241 and the Fondation de France.

References

1. Altmann, U.: Studying movement synchrony using time series and regression models. In: Esposito, I.A., Hoffmann, R., Hübner, S., Wrann, B. (eds.) Program and Abstracts of the COST 2102 Final Conference Held in Conjunction with the 4th COST 2102 International Training School on Cognitive Behavioural Systems, p. 23 (2011)
2. Ashenfelter, K.T., Boker, S.M., Waddell, J.R., Vitanov, N.: Spatiotemporal symmetry and multifractal structure of head movements during dyadic conversation. *J. Exp. Psychol. Hum. Percept. Perform.* 35(4), 1072–1091 (2009)
3. Bernieri, F.J., Reznick, J.S., Rosenthal, R.: Synchrony, pseudo synchrony, and dis-synchrony: Measuring the entrainment process in mother-infant interactions. *Journal of Personality and Social Psychology* 54(2), 243–253 (1988)
4. Boucenna, S., Gaussier, P., Andry, P.: What should be taught first: the emotional expression or the face? In: *epirob* (2008)
5. Boucenna, S., Gaussier, P., Hafemeister, L., Bard, K.: Towards a new social referencing paradigm. In: *epirob 2009*, pp. 201–202 (2009)
6. Cappella, J.N.: Behavioral and judged coordination in adult informal social interactions: vocal and kinesic indicators. *Pers. Soc. Psychol.* 72, 119–131 (1997)
7. Cassel, R.S., Saint-Georges, C., Mahdhaoui, A., Chetouani, M., Laznik, M.-C., Muratori, F., Adrien, J.-L., Cohen, D.: Course of maternal prosodic incitation (motherese) during early development in autism: an exploratory home movie study. *Interaction Studies* (in Press)
8. Champagne, F., Diorio, J., Sharma, S., Meaney, M.J.: Naturally occurring variations in maternal behavior in the rat are associated with differences in estrogen-inducible central oxytocin receptors. *Proceedings of the National Academy of Sciences of the United States of America* 98, 12736–12741 (2001)
9. Delaherche, E., Chetouani, M.: Multimodal coordination: exploring relevant features and measures. In: *Second International Workshop on Social Signal Processing. ACM Multimedia* (2010)

10. Delaherche, E., Chetouani, M., Bigouret, F., Xavier, J., Plaza, M., Cohen, D.: Assessment of communicative and coordination skills of children with pervasive developmental disorders and typically developing children (submitted, 2012)
11. Delaherche, E., Chetouani, M., Mahdhaoui, M., Saint-Georges, C., Viaux, S., Cohen, D.: Interpersonal synchrony: A survey of evaluation methods across disciplines. *IEEE Transactions on Affective Computing* (to appear, 2012)
12. Denenberg, V.H., Whimby, A.E.: Behavior of adult rats is modified by the experiences their mothers had as infants. *Science* 142, 1192–1193 (1963)
13. Falk, D.: Prelinguistic evolution in early hominins: whence motherese? *Behavioral and Brain Sciences* 27(4), 491–503 (2004)
14. Feldman, R.: Infant-mother and infant-father synchrony: the coregulation of positive arousal. *Infant Mental Health Journal* 24(1), 1–23 (2003)
15. Feldman, R.: Parent-infant synchrony and the construction of shared timing; physiological precursors, developmental outcomes, and risk conditions. *The Journal of Child Psychology and Psychiatry and Allied Disciplines* 48(3-4), 329–354 (2007)
16. Francis, D., Diorio, J., Liu, D., Meaney, M.J.: Nongenomic transmission across generations of maternal behavior and stress responses in the rat. *Science* 286, 1155–1158 (1999)
17. Gaussier, P.: Toward a cognitive system algebra: A perception/action perspective. In: *European Workshop on Learning Robots (EWRL)*, pp. 88–100 (2001)
18. Goldstein, M.H., King, A.P., West, M.J.: Social interaction shapes babbling: Testing parallels between birdsong and speech. *Proceedings of the National Academy of Sciences of the United States of America* 100(13), 8030–8035 (2003)
19. Harriot, A.W., Waugh, R.M.: Dyadic synchrony: Its structure and function in children's development. *Developmental Review* 22(4), 555–592 (2002)
20. Cappella, J.: *Coding Mutual Adaptation in Dyadic Nonverbal Interaction*, pp. 383–392. Lawrence Erlbaum (2005)
21. Kozima, H., Michalowski, M., Nakagawa, C.: *Keepon*. *International Journal of Social Robotics* 1, 3–18 (2009)
22. Kuhl, P.K.: Early language acquisition: cracking the speech code. *Nat. Rev. Neurosci.* 5(11), 831–843 (2004)
23. Liu, D., Diorio, J., Tannenbaum, B., Caldji, C., Francis, D., Freedman, A., Sharma, S., Pearson, D., Plotsky, P.M., Meaney, M.J.: Maternal care, hippocampal glucocorticoid receptors, and hypothalamic-pituitary-adrenal responses to stress. *Science* 277, 1659–1662 (1997)
24. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 2, 91–110 (2004)
25. Mahdhaoui, A., Chetouani, M., Cassel, R.S., Saint-Georges, C., Parlato, E., Laznik, M.C., Apicella, F., Muratori, F., Maestro, S., Cohen, D.: Computerized home video detection for motherese may help to study impaired interaction between infants who become autistic and their parents. *International Journal of Methods in Psychiatric Research* (2011)
26. Marcelli, D., Cohen, D.: *Enfance et psychopathologie*. Masson (2012)
27. Mataruna, H.R., Varela, F.J.: *Autopoiesis and Cognition: the realization of the living*. Reidel, Dordrecht (1980)
28. Meltzoff, A.N., Kuhl, P.K., Movellan, J., Sejnowski, T.J.: Foundations for a new science of learning. *Science* 325(5938), 284–288 (2009)
29. Messinger, D.M., Ruvolo, P., Ekas, N.V., Fogel, A.: Applying machine learning to infant interaction: The development is in the details. *Neural Networks* 23(8-9), 1004 (2010); *Social Cognition: From Babies to Robots*

30. Morency, L.-P., de Kok, I., Gratch, J.: Predicting Listener Backchannels: A Probabilistic Multimodal Approach. In: Prendinger, H., Lester, J.C., Ishizuka, M. (eds.) IVA 2008. LNCS (LNAI), vol. 5208, pp. 176–190. Springer, Heidelberg (2008)
31. Oullier, O., de Guzman, G.C., Jantzen, K.J., Scott Kelso, J.A., Lagarde, J.: Social coordination dynamics: Measuring human bonding. *Social Neuroscience* 3(2), 178–192 (2008)
32. Ozkan, D., Sagae, K., Morency, L.-P.: Latent mixture of discriminative experts for multimodal prediction modeling. *Computational Linguistics*, 860–868 (August 2010)
33. Prepin, K., Pelachaud, C.: Shared understanding and synchrony emergence: Synchrony as an indice of the exchange of meaning between dialog partners. In: ICAART 2011 International Conference on Agent and Artificial Intelligence, vol. 2, pp. 25–30 (January 2011)
34. Ramseyer, F., Tschacher, W.: Nonverbal synchrony in psychotherapy: Coordinated body movement reflects relationship quality and outcome. *Journal of Consulting and Clinical Psychology* 79(3), 284–295 (2011)
35. Richardson, D., Dale, R., Shockley, K.: Synchrony and swing in conversation: Coordination, temporal dynamics, and communication. Oxford University Press (2008)
36. Richardson, D.C., Dale, R.: Looking to understand: The coupling between speakers' and listeners' eye movements and its relationship to discourse comprehension. *Cognitive Science* 29(6), 1045–1060 (2005)
37. Richardson, M.J., Marsh, K.L., Isenhower, R.W., Goodman, J.R.L., Schmidt, R.C.: Rocking together: Dynamics of intentional and unintentional interpersonal coordination. *Human Movement Science* 26(6), 867–891 (2007)
38. Saint-Georges, C., Mahdhaoui, A., Chetouani, M., Cassel, R.S., Laznik, M.-C., Apicella, F., Muratori, P., Maestro, S., Muratori, F., Cohen, D.: Do parents recognize autistic deviant behavior long before diagnosis? Taking into account interaction using computational methods. *PLoS ONE* 6(7), e22393 (2011)
39. Shockley, K., Santana, M.-V., Fowler, C.A.: Mutual interpersonal postural constraints are involved in cooperative conversation. *Journal of Experimental Psychology: Human Perception and Performance* (29) (2003)
40. Sun, X., Truong, K., Nijholt, A., Pantic, M.: Automatic visual mimicry expression analysis in interpersonal interaction. In: Proceedings of IEEE Int'l Conf. Computer Vision and Pattern Recognition (CVPR-W 2011), Workshop on CVPR for Human Behaviour Analysis, Colorado Springs, USA, pp. 40–46 (2011)
41. Thorisson, K.R.: Natural Turn-Taking Needs No Manual: Computational Theory And Model, From Perception to Action, pp. 173–207. Kluwer Academic Publishers, Dordrecht (2002)
42. Varni, G., Volpe, G., Camurri, A.: A system for real-time multimodal analysis of nonverbal affective social interaction in user-centric media. *IEEE Transactions on Multimedia* 12(6), 576–590 (2010)
43. Ward, N.G., Fuentes, O., Vega, A.: Dialog prediction for a general model of turn-taking. In: INTERSPEECH, pp. 2662–2665 (2010)
44. Weaver, I.C., Cervoni, N., Champagne, F.A., D'Alessio, A.C., Sharma, S., Seckl, J.R., Dymov, S., Szyf, M., Meaney, M.J.: Epigenetic programming by maternal behavior. *Nat. Neuroscience* 7, 847–854 (2004)
45. Widrow, B., Hoff, M.E.: Adaptive switching circuits. In: IRE WESCON, New York, pp. 96–104 (1960); Convention Record

Emotion and Complex Tasks: Writing Abilities in Young Graders

Michaël Fartoukh¹, Lucile Chanquoy^{1,*}, and Annie Piolat²

¹ BCL, University of Nice – Sophia-Antipolis, France
{lucile.chanquoy,michael.fartoukh}@unice.fr

² PSYCLÉ, University of Aix-Marseille, France
annie.piolat@univ-amu.fr

Abstract. Writing processes depend on the development and the capacity of working memory. Their execution is highly costly in cognitive resources. During writing, emotions are potentially present. According to Ellis and Ashbrook's (1988) model, emotions are expected to cause interferences in working memory by creating extra cognitive load. Our main hypothesis was that emotions should be compared to a secondary task, overloading working memory capacities. Two experiments using emotional induction procedures were carried out on two different writing tasks (text production and dictation) with young graders. Results showed that emotional content interfered as cognitive overload within the limited working memory resources and had an impact on orthographic abilities. In terms of computational intelligence, as emotions seem to have an impact on the availability of cognitive resources, this could lead to important theoretical and practical implications for the elaboration of interactive scenarios or modeling learning and processing procedures.

Keywords: Emotion, cognition, working memory, writing processes.

1 Introduction

It is nowadays well known that emotions can modify and influence adults' cognitive processing and performance during complex problem solving task. For example, several models focus on the relationships between emotion and memory [5], information processing [11], processing capacity [14] or creative thinking [19]. However, there is little understanding on how emotions can influence cognition.

2 Influence of Emotions on Cognition in Adults

There is empirical evidence that cognitive performance may be improved or reduced according to emotional contexts. A positive emotional state can sometimes lead to a

* To contact the authors, please write to : Lucile Chanquoy, Université de Nice - Sophia Antipolis, UMR 7320, Bases, Corpus, Langage, Campus Saint Jean d'Angély, Institut des SHS de Nice, 24, Avenue des Diables Bleus, 06357 NICE cedex 4, France.

better performance in decision making [20], problem solving [21], thinking flexibility [19], or during intuitive coherence judgments [1]. Conversely, in case of positive emotional states a reduction in cognitive performance leading to a decrease of cognitive processing [8], a reduction in processing capacity [25], and sometimes a certain decline of motivation [4] has also been noticed.

Similar effects have been reported for negative emotional states, with a benefit in the development of systematic, analytic and elaborate cognitive processing [11], [30]. However, negative emotional states are more generally associated to an increase of time processing and a decrease in cognitive performance, namely during problem solving and/or creative thinking activities [5], [14], [19], [25]. This reduction in cognitive performance in case of negative emotional states has been explained, *inter alia*, by a decrease in information processing capacity. Indeed, according to Ellis and Ashbrook's Resource Allocation Model [14], negative emotional states could be responsible for extra-task processing or task-irrelevant processing (intrusive thinking, ruminations). Nevertheless, some researchers consider that both positive and negative emotional states are resource consuming and are responsible for a decrease in cognitive performance [15]. This effect could be mediated by an overload of working memory capacities.

3 Influence of Emotions on Working Memory

Emotion effects are noticeable on working memory capacities [13-14], [24], [28]. Positive and negative emotional states seem to have the same effect on storing and retrieving capacities in working memory. For example, Martin and Kerns [24] have shown a decrease in storage capacities in case of positive emotional state. Moreover, according to Oaksford, Morris, Grainger and Williams [25], a positive emotional state can impair the ability of planning (*i.e.*, in a Tower of London task situation) in loading the central executive resources. This had led some researchers to consider emotions as an interfering task costly in attentional resources [15], [17], [25]. In this case, the effect of emotional states seems to be essentially negative and could be due to intrusive thinking and/or ruminations [14], [16].

Indeed, intrusive thinking is highly costly in working memory resources [29]. However, according to Phillips, Smith and Gilhooly [26], the effects of emotional states depend on working memory span or capacity. Participants with small capacities are more impacted by an emotional component and thus, children who are in the process of cognitive development could be more affected by emotions.

4 Influence of Emotions on Cognition in Children

Research on the interaction between cognition and emotion in children started more than 30 years ago. However, nowadays, there are still few papers on this issue. One of the first studies carried out with children was led by Bartlett and Sandtröck [2] and it

was about the recall of a story after a positive or negative emotional induction procedure in five years-old children. Participants had to listen to a story and had to recall it without cues first and with cues afterwards. Results showed that the positive condition recall was better than the negative condition one. In another study [23], the learning capacities and the problem solving duration were also improved in case of positive emotional state whereas it was the opposite in case of negative emotional states. More recently, Rader and Hugues [27] have shown better results in a visual problem task in case of positive emotional states and Burkitt and Barnett [7] have shown that children drawings were richer in case of positive emotional state than in negative emotional states.

In a more academic context, Bryan and Bryan [6] found that when children were under positive emotional induction they performed better when facing mathematical problem solving. Cuisinier, Sanguin-Bruckert, Bruckert and Clavel [12] have observed an increase in the number of orthographic errors with an emotional content in 5th graders. In their experiment, after an emotional induction procedure through the reading of three emotional content texts (neutral, positive or negative), children were asked to write a dictation from the text. It thus seems that emotional states could equally modify the writing processes. This was considered within the framework of Hayes' model of writing [18].

5 Emotion and Writing

According to Hayes's writing model [18], many aspects of the human cognition interact with the different writing processes, such as long-term memory, working memory, but also motivation or affect. More precisely, working memory, due to its limited capacities, plays a very important role overall the writing activity. However, through the motivation and affect components, Hayes affirms that emotions are fundamental during writing.

In parallel, learning how to write a text is a long and complex activity, more than ten years are necessary to become an expert [9]. The development of working memory would permit efficient information retrieval and storing by a more efficient cognitive resource allocation and/or sharing during writing [3]. Slowly, some processes become automated and the control process – that controls the online distribution of processes – becomes more proficient; this leads to freeing up attention resources for other activities, mainly those that are not yet automated or that cannot be automated. Concerning spelling issues, as all the rules and procedures are not totally efficient or “fixed”, the associated cognitive cost is high and some weakly automated orthographic skills are easily concerned by a cognitive overload. For instance, Chenoweth and Hayes [10] have observed an increase in spelling errors with a resource-consuming secondary task.

Depending on the goal, the potential reader, and/or the emotional state, a text produced by a writer could be affected in terms of length, orthographic accuracy or more generally, its content. A positive or negative emotional state, which could require

additional working memory capacities, could lead to an increase in the proportion of spelling errors during a sentence production task or a dictation. This was partially shown in the study carried out by Cuisinier and her collaborators [12].

Thus, just as emotional states seem to be able to modify complex cognitive activities, more precisely working memory resources, writing processes being a complex task, could lead to cognitive overload. Hence, we hypothesize that the introduction of emotional content before or during the writing activity could be resource consuming. To test the effect of emotion on writing, two experiments were conducted using two different writing tasks: a text production and a dictation tasks. In agreement with the literature, we predicted that emotions could have an impact on text length and on orthographic accuracy, given the lack of cognitive resources available.

6 Experiments

The experiments were led in 4 classrooms (1 for experiment 1 and 3 for experiment 2) of French 4th graders, in their school. All the parents have given their written approval for their child to participate in these experiments.

6.1 Experiment 1

In the first experiment, the children ($n = 20$, 11 boys and 9 girls, mean age: 9 years and 10 months, ranking from 9 to 10.8 years-old) had to write a text under three emotional conditions: positive (children had to write about the happiest day of their life), negative (the saddest day of their life), and neutral (their classroom activities of the morning before). To estimate each grader's initial spelling level and working memory level, the spelling test from the "ECS-Cycle 3" scale [22] (a global assessment for the last years of French elementary school elaborated by Khomsi in 1998) and the Working Memory Index (WMI) of the Wechsler Intelligence Scale for Children [31] (4th edition, French version) were used (see Table 1).

Table 1. Mean scores for orthographic and WM tests (standard deviations in italics)

Spelling results	WM index
19.70 (<i>7.91</i>)	88.75 (<i>12.23</i>)

Participants were asked to produce a text in one of three conditions during three successive weeks. The average time for each writing session was approximately 45 minutes. In parallel, the working memory test was individually administrated in a quiet room of the school.

After evaluating the number of words in each text, the spelling correction of texts was made for both grammatical and lexical errors. Thus, both types of errors were estimated and a percentage of errors for each text was computed by dividing the total number of errors by the total number of words and multiplied by 100 (see Table 2).

Table 2. Mean number of words and percentages of spelling errors according to emotional instructions and error categories

		Emotional instruction		
		Neutral	Positive	Negative
Mean number of words		96.70 (8.84)	93.00 (7.95)	66.20 (5.74)
Spelling errors	Grammatical	11.23 (1.28)	13.46 (1.30)	14.14 (1.22)
	Lexical	5.14 (1.22)	4.91 (0.96)	6.03 (1.05)

First of all, a qualitative analysis of the texts has shown that the children have respected their writing instructions. The “positive emotion” instruction has led to autobiographical texts with positive emotions, states or events. The “negative emotion” instruction has also led to autobiographical texts with negative emotions. Finally, for the “neutral emotion” instruction, all of the graders have more or less detailed those events concerning their morning in the classroom.

To observe the effect of the emotional instruction on the writing processes, an analysis of variance with repeated measures (as all graders wrote 3 texts) was carried out on the number of words in the different texts. The emotional instruction had a significant effect on the number of words ($F(2, 38) = 9.99, p < .001$). Planned comparisons revealed that the difference between the number of words in texts written under neutral (96.7) and positive (93) conditions was not significant ($F(1, 19) < 1$) whereas the difference between negative (66.2) and neutral conditions and between negative and positive conditions were both significant (respectively: $F(1, 19) = 19.01, p < .001$ and $F(1, 19) = 12.06, p < .01$).

Concerning spelling errors, a 3 (emotional instructions) x 2 (types of error: grammatical and lexical) analysis of variance with repeated measures was carried out on the percentage of spelling errors. There was an effect of the types of errors ($F(1, 19) = 57.94, p < .0001$). Children make more grammatical (12.95%) than lexical errors (5.36%). There was no effect of emotional instruction on spelling errors (Neutral = 8.18%; Positive = 9.18%; Negative = 10.08%; $F(2, 38) = 1.95, p = .15$) and the interaction between emotional instruction and type of errors was not significant ($F(2, 38) = 1.38, p = .27$).

Data from the working memory test were used to make simple regression analyses. These analyses were carried out to verify if spelling errors were linked to children’s working memory capacities. Results showed a significant effect of working memory span on previous orthographic test, and on the proportion of grammatical errors, only in neutral condition. However, regardless of positive or negative instructions, no significant effect was reported and working memory capacity did not explain the proportion of spelling errors. With positive or negative condition, it seems that there was interference within the spelling processing and, in these cases, children’s capacity became less predictable (see Table 3).

Table 3. Simple regression analyses for the spelling test (first line of the table) and for the grammatical spelling errors (SE) in written texts

Explicative variable	Explained variables	R^2	B	SE		t	p
WM capacities	Spelling test score	0.09	0.17	0.08	0.30	1.97	.05
WM capacities	SE in neutral condition	0.12	0.14	0.06	0.34	2.26	.02
	SE in positive condition	-	-	-	-	-	> .05
	SE in negative condition	-	-	-	-	-	> .05

6.2 Experiment 2

In this experiment, participants were 4th graders from 3 different classrooms ($n = 77$, 39 boys and 38 girls, mean age: 10 years and 4 months, ranking from 9 years to 11.3 years-old). For the emotional induction procedure, children were distributed within three independent groups, in which they had to listen to a story with positive emotion, negative emotion or neutral emotion. To estimate each grader's initial spelling level and working memory level, the spelling test from the "ECS-Cycle 3" scale [22] and the "Letter Number Sequencing" subtest from the Working Memory Index [31] were used. Then, according to their experimental groups, three emotionally contrasted dictations [12] from three texts were proposed. During the experiment, working memory span was assessed three times: before reading the text, after reading and after text dictation.

The impact of the emotional induction procedure was evaluated on the number of correctly recalled items for the working memory test and on the number of orthographic errors (grammatical and lexical) on the same target words in the three dictations (see Table 4).

Table 4. Mean number of correctly recalled items for the working memory test according to emotional induction and moment. Mean number of spelling errors according to emotional induction and error categories.

		Emotional induction		
		Neutral	Positive	Negative
WM test correctly recalled items	Before text reading	24.28 (2.89)	22.75 (4.20)	25.45 (5.75)
	After text reading	26.00 (2.25)	22.33 (2.90)	24.54 (5.41)
	After dictation	27.00 (3.35)	28.50 (4.14)	27.72 (5.06)
Spelling errors	Grammatical	6.40 (0.98)	7.07 (1.32)	6.75 (1.71)
	Lexical	5.80 (1.26)	6.38 (1.75)	7.66 (1.92)

A 3 (emotional induction) \times 3 (moments: before text reading, after text reading, and after dictation) analysis of variance with repeated measures on the last factor was carried out on the number of correctly recalled items from the WM test. There was no significant effect of the emotional induction group ($F(2, 34) < 1$), as shown by the very similar means (neutral: 25.76, positive: 24.52, and negative: 25.90). Conversely, there was a significant effect of the moment ($F(2, 68) = 27.32, p < .0001$). The mean number of correctly recalled items increased with time, due to a learning effect (before text reading: 24.16, after text reading: 24.29, and after dictation: 27.74). The interaction between the two variables was significant ($F(4, 68) = 4.65, p < .01$). After text reading, children in neutral condition improved their performances whereas children's results in negative or positive condition decreased (respectively: $F(1, 34) = 5.70, p < .03$ and $F(1, 34) = 3.95, p < .05$).

Concerning spelling errors, a 3 (emotional induction) x 2 (types of error: grammatical and lexical) analysis of variance with repeated measures on the last factor was carried out on the number of spelling errors. There was no significant effect of the type of errors ($F(1, 34) < 1$), children made as many grammatical (6.73) than lexical (6.44) errors. Conversely, the emotional induction procedure had a significant effect on spelling errors ($F(2, 34) = 4,11, p < .03$). Planned comparisons showed that children in positive (6.79) and negative (7.09) conditions made more spelling errors (both types) than children in neutral (5.87) condition ($F(1, 34) = 7,80, p < .01$). The interaction between the two variables was not significant ($F(2, 34) = 3.00, p < .07$).

7 Conclusion

In this paper, the impact of emotions on young graders' written performance (text production and dictation) was explored. In experiment 1, the effect of emotion on writing fluidity was assessed. In emotional assignments, children seem to have more difficulties to write about negative emotional content than positive or neutral content. However, there was no effect of emotional instruction on spelling errors. These results provide consistent argument for Hayes model's [18], even whether the chosen writing tasks seem to be not enough constraining in terms of rapid resource mobilization to impact orthographic accuracy. Indeed, the time allotted for a writing task is generally longer and children can choose vocabulary, syntax and are able to revise their texts. However, it seems that a relationship between working memory span and spelling capacity exists but disappears when an emotional content is introduced. A decrease in attentional resources during "emotional" writing could explain this result [15], [17], [24-25], but, once again, the participants' emotional state could be not sufficiently strong to induce a lack of WM resources leading to spelling difficulties and thus the orthographic accuracy was not significantly reduced.

As already shown in adults [14], [16], results from experiment 2 showed an effect of the emotional induction procedure on the availability of attentional resources (working memory span) in children. Positive and negative emotions seem to interfere with working memory capacities in reducing the available cognitive resources [24-25]. Moreover, there is an effect of the emotional induction procedure on spelling errors, resulting in an increase in case of positive or negative emotions. This result is comparable to previous results obtained with French children [12].

To conclude, after an emotion induction procedure, young graders' spelling accuracy seems to be less predictable (Experiment 1) and, young graders' working memory span and spelling accuracy decrease (Experiment 2). There is an impact of the emotional state on cognitive resources and sometimes emotion seems to act as a cognitive overload in young graders' working memory. In terms of computational intelligence and learning models, as the emotional state dimension seems to have an impact on the cognitive resources available and on children's abilities, this could lead to important theoretical and practical implications for the construction of experiments, models and also AI models.

References

1. Balas, R., Swekleij, J., Pochwatko, G., Godlewska, M.: On the Influence of Affective States on Intuitive Coherence Judgements. *Cognition and Emotion* 26(2), 312–320 (2012)
2. Bartlett, J.C., Santrock, J.W.: Affect-Dependent Episodic Memory in Young Children. *Child Development* 50, 513–518 (1979)
3. Berninger, V.W., Swanson, H.L.: Modification of the Hayes and Flower Model to Explain Beginning and Developing Writing. In: Butterfield, E. (ed.) *Advances in Cognition and Educational Practice, Children's Writing: Toward a Process Theory of Development of Skilled Writing*, vol. 2, pp. 57–82. JAI Press, Greenwich (1994)
4. Bodenhausen, G.V., Kramer, G.P., Süsler, K.: Happiness and Stereo-Typic Thinking in Social Judgment. *Journal of Personality and Social Psychology* 66, 621–632 (1994)
5. Bower, G.H.: Mood and Memory. *American Psychologist* 36(2), 129–148 (1981)
6. Bryan, T., Bryan, J.: Positive Mood and Math Performance. *Journal of Learning Disabilities* 24, 490–493 (1991)
7. Burkitt, E., Barnett, N.: The Effects of Brief and Elaborate Mood Induction Procedures on the Size of Young Children's Drawings. *Educational Psychology* 26(1), 93–108 (2006)
8. Chaiken, S.: Heuristic Versus Systematic Information and the Use of Source Versus Message Cues in Persuasion. *Journal of Personality and Social Psychology* 39, 752–766 (1980)
9. Chanquoy, L., Alamargot, D.: Mise en Place et Développement des Traitements Rédactionnels: le Rôle de la Mémoire de Travail. *Le Langage et l'Homme* 38(2), 171–190 (2003)
10. Chenoweth, N.A., Hayes, J.R.: The Inner Voice in Writing. *Written Communication* 20, 99–118 (2003)
11. Clore, G.L., Schwarz, N., Conway, M.: Affective Causes and Consequences of Social Information Processing. In: Wyer, R.S., Srull, T.K. (eds.) *Handbook of Social Cognition*, vol. 1, pp. 323–417. Erlbaum, Hillsdale (1994)
12. Cuisinier, F., Sanguin-Bruckert, C., Bruckert, J.P., Clavel, C.: Les Émotions Affectuelles les Performances Orthographiques en Dictées? *L'Année Psychologique* 110(1), 3–48 (2010)
13. Derakshan, N., Eysenck, N.W.: Working Memory Capacity in High Trait-Anxious and Repressor Groups. *Cognition and Emotion* 12(5), 697–713 (1998)
14. Ellis, H.C., Ashbrook, P.W.: Resource Allocation Model of the Effects of Depressed Mood States on Memory. In: Fiedler, K., Forgas, J.P. (eds.) *Affect, Cognition and Social Behaviour*, pp. 25–44. Hogrefe, Göttingen (1988)
15. Ellis, H.C., Moore, B.A.: Mood and Memory. In: Dalgleish, T., Power, M.J. (eds.) *Handbook of Cognition and Emotion*, pp. 193–210. Wiley, Chichester (1999)
16. Eysenck, M.W., Derakshan, N., Santos, R., Calvo, M.G.: Anxiety and Cognitive Performance: Attentional Control Theory. *Emotion* 7, 336–353 (2007)
17. Gotoh, F.: Influence of Affective Valence on Working Memory Processes. *International Journal of Psychology* 43(1), 59–71 (2008)
18. Hayes, J.R.: A New Framework for Understanding Cognition and Affect in Writing. In: Levy, M.C., Ransdell, S. (eds.) *The Science of Writing*, pp. 1–28. Lawrence Erlbaum, Mahwah (1996)
19. Isen, A.M., Daubman, K.A., Nowicki, G.P.: Positive Affect Facilitates Creative Problem Solving. *Journal of Personality and Social Psychology* 52(6), 1122–1131 (1987)

20. Isen, A.M., Labroo, A.A.: Some Ways in Which Positive Affect Facilitates Decision Making and Judgement. In: Schneider, S.L., Shanteau, J. (eds.) *Emerging Perspectives on Judgement and Decision Research*, pp. 365–393. Cambridge University Press, New York (2003)
21. Isen, A.M., Rosenzweig, A.S., Young, J.M.: The Influence of Positive Affect on Clinical Problem Solving. *Medical Decision Making* 11, 221–227 (1991)
22. Khomsi, A.: *Évaluation des Compétences Scolaires: Cycle des Approfondissements*. ECPA, Paris (1998)
23. Masters, J.C., Barden, R.C., Ford, M.E.: Affective States, Expressive Behavior, and Learning in Children. *Journal of Personality and Social Psychology* 37(3), 380–390 (1979)
24. Martin, E.A., Kerns, J.G.: The Influence of Positive Mood on Different Aspects of Cognitive Control. *Cognition and Emotion* 25(2), 265–279 (2011)
25. Oaksford, M., Morris, F., Grainger, B., Williams, J.M.G.: Mood, Reasoning, and Central Executive Processes. *Journal of Experimental Psychology: Learning, Memory and Cognition* 22(2), 476–492 (1996)
26. Phillips, L.H., Smith, L., Gilhooly, K.J.: The Effects of Adult Aging and Induced Positive and Negative Mood on Planning. *Emotion* 2(3), 263–272 (2002)
27. Rader, N., Hughes, E.: The Influence of Affective State on the Performance of a Block Design Task in 6- and 7-Year-Old Children. *Cognition & Emotion* 19(1), 143–150 (2005)
28. Spies, K., Hesse, F.W., Hummitzsch, C.: Mood and Capacity in Baddeley's Model of Human Memory. *Zeitschrift für Psychologie* 204, 367–381 (1996)
29. Teasdale, J.D., Dritschel, B.H., Taylor, M.J., Proctor, L., Lloyd, C., Nimmo-Smith, I., Baddeley, A.D.: Stimulus-Independent Thought Depends on Central Executive Resources. *Memory and Cognition* 23, 551–559 (1995)
30. Weary, G., Jacobsen, J.A.: Causal Uncertainty Beliefs and Diagnostic Information Seeking. *Journal of Personality and Social Psychology* 98, 150–153 (1997)
31. Wechsler, D.: *Échelle d'Intelligence de Wechsler pour Enfants et Adolescents – Quatrième Édition*. ECPA, Paris (2005)

A Preliminary Study of Online Drawings and Dementia Diagnose

Marcos Faundez-Zanuy¹, Enric Sesa-Nogueras¹,
Josep Roure-Alcobe¹, Josep Garre-Olmo², Jiri Mekyska³,
KarmeLe Lopez-de-Ipiña⁴, and Anna Esposito⁵

¹ Escola Universitària Politècnica de Mataró, Tecnocampus
Barcelona, Spain

{faundez,sesa,roure}@tecnocampus.com

² Institut d'assistència sanitària, Salt
Girona, Spain

josep.garre@ias.scs.es

³ SPLab, Department of Telecommunications, FEEC,
Brno University of Technology Brno, Czech Republic

j.mekyska@phd.feec.vutbr.cz

⁴ EHU (Basque Country University)
Donostia, Spain

karmeLe.ipina@ehu.es

⁵ IIASS (International institute for advanced scientific studies)
Vietri sul Mare, Italy

iiass.annaesp@tin.it

Abstract. In this paper we present preliminary results about on-line drawings acquired by means of a digitizing tablet, and performed by control population (left and right hand) as well as pathological subjects using their dominant hand. Experimental results reveal a clear difference between both groups, specially on the on-air movements. Although the acquired samples are not enough to extract significant conclusions we think that this preliminary results encourage the experimentation in this research line. Thus, the main purpose of this paper is to attract the attention of the scientific community.

Keywords: On-line handwriting, Drawings, Dementia.s

1 Introduction

Information and communication technologies are converging to health applications and a great improvement of health diagnose and recover will be possible if the signal processing community collaborates with medical doctors. The authors strongly believe that for the biometric security community this would be quite straightforward, and this paper points out in this direction.

Biometrics has been successfully applied to security applications for some time. However, the extension of other potential applications with the use of

biometric information is a very recent development. This paper summarizes the field of biometrics and investigates the potential of utilizing biometrics beyond the presently limited field of security applications.

The term “biometric” originates from the Greek words Bio (life) and metron (measure), and is defined as the science and technology of measuring and statistically analyzing biological data. Although many people consider biometrics only relevant to security applications, in reality, the relevance of biometrics is much more far reaching. This field has applications relevant to animals, plants and human beings. Some examples are:

- Statistical methods for the analysis of data from agricultural field experiments to compare the yields of different varieties of wheat.
- The analysis of data from human clinical trials evaluating the relative effectiveness of competing disease therapies.
- The analysis of biometric characteristics for animal/human verification or identification.

While some signals can be acquired from both human beings and animals (such as images of iris and retina), others are specific to humans (such as speech, handwriting, etc.).

This paper is focused exclusively on applications which are relevant only to human beings, and more precisely on on-line handwritten drawings. Therefore, we will limit discussion to only human specific signals. The set of these signals can be split into two categories:

1. Behavioral biometrics: this category is based on the measurements and data derived from an action performed by a user, and thus indirectly measures some characteristics of the human body. Signature, gait, gesture and key stroking recognition belong to this category.
2. Physiological biometrics: this category is based on direct measurements of parts of the human body. Fingerprint, face, iris and hand-scanning recognition belong to this category.

The skill level of humans is strongly related to their health state. An important example is the way our cognitive functions are related to the aging process. Cognitive decline is a natural part of the aging process. However, the extent of decline varies across subjects and across functions. For instance, handwriting and speech production is a fine motor control performed by our brain. When these signals are degraded, it is indicative of health problems.

2 On-Line Handwriting

In the past, the analysis of handwriting had to be performed in an offline manner. Only the writing itself (strokes on a paper) were available for analysis. Nowadays, modern capturing devices, such as digitizing tablets and pens (with or without ink) can gather data without losing its temporal dimension. When spatiotemporal information is available, its analysis is referred as online. Modern digitizing

tablets not only gather the x-y coordinates that describe the movement of the writing device as it changes its position, but it can also collect other data, mainly the pressure exerted by the writing device on the writing surface and also the azimuth, the angle of the pen in the horizontal plane, and the altitude, the angle of the pen with respect the vertical axis (Fig. 1).

A very interesting aspect of the modern online analysis of handwriting is that it can take into account information gathered when the writing device was not exerting pressure on the writing surface. Thus, the movements performed by the hand while writing a text can be split into two classes:

- a) *On-surface trajectories* (pen-downs), corresponding to the movements executed while the writing device is touching the writing surface. Each of these trajectories produces a visible stroke.
- b) *In-air trajectories* (pen-ups), corresponding to the movements performed by the hand while transitioning from one stroke to the next. During these movements the writing device exerts no pressure on the surface.




Fig. 1. Azimuth and altitude angles of the pen with respect to the plane of the writing surface

Fig. 2 shows the acquisition of the ten digits from 1 to 0 using an Intuos Wacom digitizing tablet. The tablet acquired 100 samples per second including the spatial coordinates (x , y), the pressure, and a couple of angles (see Fig. 1). The pen-up information is represented in Fig. 1 using “+” while the pen-down is marked with “*”. Fig. 3 shows the temporal evolution of the signals acquired while handwriting the digits in Fig. 2.

Our experiments on the biometric recognition of people reveal that these two kinds of information are complementary [8] and in fact, contain a similar discriminative capability, even when using a database of 370 users [7].

3 On-Line Drawings Applied to Health Analysis

In the medical field, the study of handwriting has proven to be an aid to diagnose and track some diseases of the nervous system. For instance, handwriting skill




Fig. 2. Example of handwritten numerical digits input onto a digitizing tablet. Asterisks (*) represent pen-down information and cross (+) the pen-up.

degradation and Alzheimer's disease (AD) appear to be significantly correlated [3] and some handwriting aspects can be good indicators for its diagnosis [4] or help differentiate between mild Alzheimer's disease and mild cognitive impairment [10]. Also, the analysis of handwriting has proven useful to assess the effects of substances such alcohol [1] [5], marijuana [2] or caffeine [9]. Aided by modern acquisition devices, the field of psychology has also benefitted from the analysis of handwriting. For instance in [6], Rosenblum et al. link the proficiency of the writers to the length of the in-air trajectories of their handwritings.




Fig. 3. Temporal evolution of the acquired parameters when drawing the numbers shown in Fig. 1

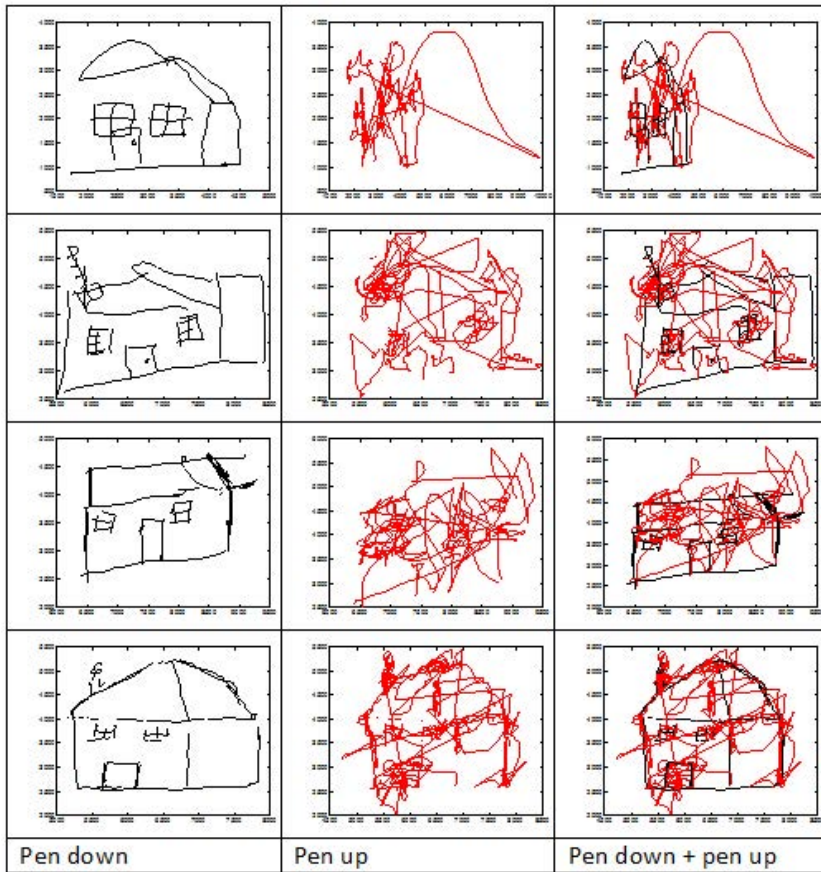


Fig. 4. House drawing performed by four individuals with Alzheimer's disease (one per row). Each column corresponds to pen-down, pen-up and both simultaneously.

Table 1. Statistical analysis/descriptives from the drawings shown in Fig. 4, 5 and 6

Measurement	Control		Pathological
	Dominant hand	Non-dominant hand	Dominant hand
Time in-air	8334	10927	61008
Time on-surface	9680	22177	31521
Total time	18014	33104	92259

In the Fig. 4 we present one complex drawing with three dimensions performed by individuals with AD of different clinical severity. The visual inspection of the pen down image suggest a progressive degree of impairment, where drawing becomes more disorganized and the three dimensions effect is only achieved in the mild case. The visual information provided by the pen up drawing between AD

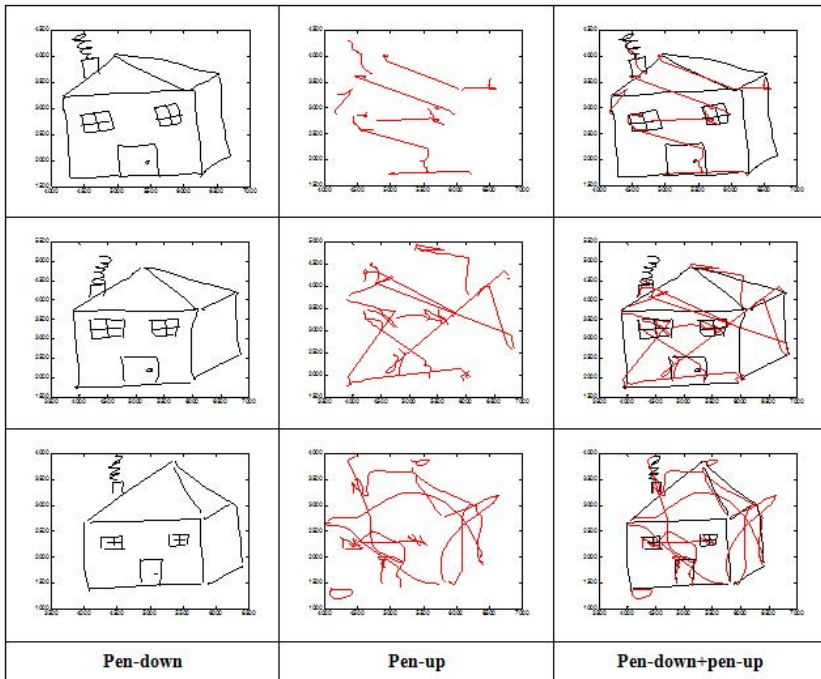


Fig. 5. House drawing performed by three control people (one per row). Each column corresponds to pen-down, pen-up and both simultaneously, performed with the dominant hand.

individuals also indicates a progressive impairment and disorganization when the individuals try to plan the drawing. It is also important to note that the comparison of the pen-up drawing between the mild case of AD and the control (Fig. 5 and 6) also shows important differences. Besides the increased time on air, there is an increased number of hand movements before decide to put the pen in the surface to drawn. We consider that these graphomotor measures applied to the analysis of drawing and writing functions may be a useful alternative to study the precise nature and progression of the drawing and writing disorders associated with several neurodegenerative diseases. Table 1 summarizes some experimental measures of the drawings shown in Fig. 4, 5 and 6.

Looking at the experimental results of Table 1 it is evident the higher time for in-air movements for the AD group, which are around 7 times longer. On the contrary, the time on surface is just around 3 times longer. Thus, there are more differences between control and AD groups when looking at in-air movements.

When comparing the non-dominant movements performed by the control group we obtained a 5.6 ratio and 1.5. Again, the in-air times are significantly higher for the AD group than the control group.

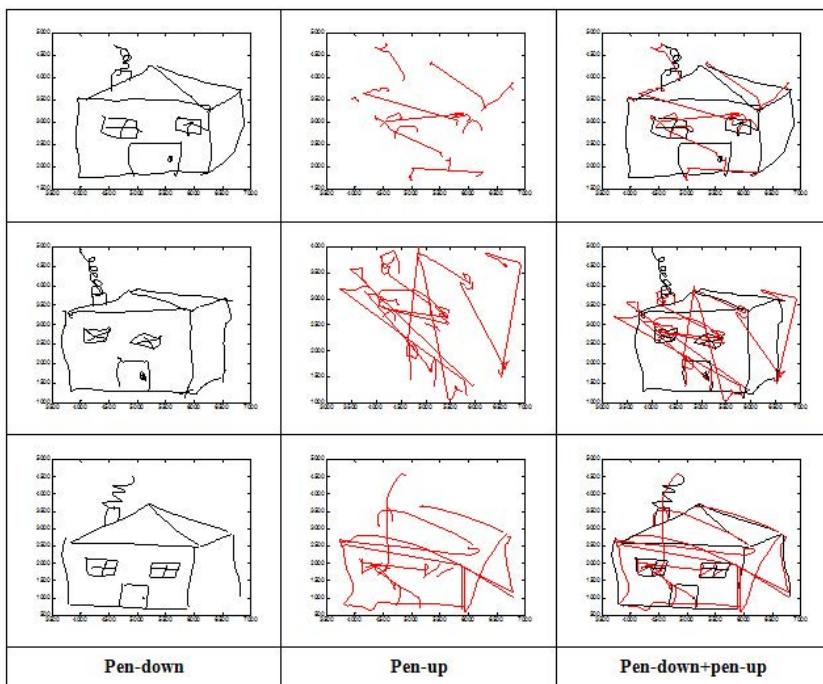


Fig. 6. House drawing performed by three control people (one per row). Each column corresponds to pen-down, pen-up and both simultaneously, performed with the non-dominant hand.

4 Conclusions

Although some pathological drawings may look “normal” according to pen-down information, the pen-up information looks quite entangled and should permit easier diagnose. This observation points out the convenience of online handwriting analysis, which can outperform the classic offline mode, mainly due to the larger amount of available information.

The differences between control and pathological group do not seem to be related to some physical problem, because the control group, even when using the non-dominant hand performs less entangled pen-up movements.

Future work will include a more exhaustive experimental section, with a larger database.

Acknowledgments. This research has been supported by Saiotek, FEDER and MEC, TEC2009-14123-C04-04, KONTAKT-ME 10123, VG20102014033, CZ.1.-07/2.3.00/20.0094, GACR 102/12/1104 and CZ.1.05/2.1.00/03.0072.

References

1. Ascioglu, F., Turan, N.: Handwriting Changes Under the Effect of Alcohol. *Forensic Science International* 132(3), 201–210 (2003)
2. Foley, R.G., Miller, A.L.: The effects of marijuana and alcohol usage on handwriting. *Forensic Science International* 14, 159–164 (1979)
3. Forbes, K.E., Shanks, M.F., Venneri, A.: The evolution of dysgraphia in Alzheimers disease. *Brain Research Bulletin* 63(1), 19–24 (2004)
4. Neils-Strunjas, J., Groves-Wright, K., Mashima, P., Harnish, S.: Dysgraphia in Alzheimer's Disease: A Review for Clinical and Research Purposes. *J. Speech Lang. Hear. Res.* 49(6), 1313–1330 (2006)
5. Phillips, J.G., Ogeil, R.P., Mller, F.: Alcohol Consumption and Handwriting: A Kinematic Analysis. *Human Movement Science* 28(5), 619–632 (2009)
6. Rosenblum, S., Parush, S., Weiss, P.: The In Air Phenomenon: Temporal and Spatial Correlates of the Handwriting Process. *Percept. Mot. Skills* 96, 933–954 (2003)
7. Sesa-Nogueras, E., Faundez-Zanuy, M.: Biometric Recognition Using Online Uppercase Handwritten Text. *Pattern Recogn.* 45(1), 128–144 (2012)
8. Sesa-Nogueras, E., Faundez-Zanuy, M., Mekyska, J.: An Information Analysis of In-Air and On-Surface Trajectories in Online Handwriting. *Cognitive Computation*, 1–11 (2012)
9. Tucha, O., Walitza, S., Mecklinger, L., Stasik, D., Sontag, T.A., Lange, K.W.: The Effect of Caffeine on Handwriting Movements in Skilled Writers. *Human Movement Science* 25(4-5), 523–535 (2006)
10. Werner, P., Rosenblum, S., Bar-On, G., Heinik, J., Korczyn, A.: Handwriting Process Variables Discriminating Mild Alzheimer's Disease and Mild Cognitive Impairment. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences* 61(4), P228–P236 (2006)

Hand-Based Gender Recognition Using Biometric Dispersion Matcher

Xavier Font-Aragones and Marcos Faundez-Zanuy

Escola Universitària Politècnica de Mataró, Tecnocampus
Barcelona, Spain
{font,faundez}@tecnocampus.com

Abstract. This paper presents a novel method for gender recognition through anthropometric hand information. From a visual hand database of a hundred users and distributed in an unbalanced way, contains more men than women. It is designed a simple method to get some length and width measurements from the hand. This information has been passed through a quadratic discriminant classifier called Biometric Dispersion Matcher (BDM) that provides relevant information. In a first step, a discriminative threshold is applied in order to discard those measures which do not have enough information for gender recognition. In a second step, it provides a vector of the main measures. And, finally, it achieves performance rates from 95 %, with a train data set of only 18 men and 9 women, to 98 %, with a higher training data set.

Keywords: Gender recognition, Hand image, Biometric dispersion matcher.

1 Introduction

Gender recognition has not been widely studied in comparison with person identification. The reason becomes clear, as security applications are by far much more interesting for industry, government agencies and business. However, in the context of human computer interaction, gender recognition could play an important role. There is also an interesting field where gender recognition applications can give some advice with intelligent advertising once your gender is detected.

The process of gender recognition is straightforward for humans. We have been trained for that task since we were children. And this task, sometimes made unconsciously, is important for social activities. Actually, the gender recognition using biometric physiological characteristics (such as face, hand, human shape) remains a difficult task. The high variability present within genders, whatever the characteristic, makes us use some of the algorithms and methods available in general recognition applications.

This paper presents a gender recognition algorithm, using the classifier Biometric Dispersion Matcher (BDM) [1]. This method is closely related to Discriminant Analysis [2] and works pretty well for linear and nonlinear problems

as well. The main idea is to build a discriminative function $g(x)$ that makes a decision depending on its sign. The details of the method are explained later in section 3.

The data for gender recognition comes from right hand images. Our database contains 104 users, with 10 samples from each user. So the full database contains 1040 hand images acquired in 5 different acquisition sessions.

The data has to be processed according to two main approaches:

- The first one, considers the whole image as a single high dimensional vector. These methods are called Holistic Methods.
- The other ones are based on anthropometric measurements (i. e. lengths, widths, angles) and are called Geometric Methods.

Our research uses the latter one.

2 Acquisition and Data Extraction

All hand images have been acquired with a simple visual optics camera with a resolution of 640x480 pixels. This data need to be preprocessed in order to obtain the desired anthropometric vector.

In order to get a coherent characteristic vector the hand image has been segmented using Otsu's method [3] (see Fig. 1). Once the image is segmented, another method is used in order to get the key points within the hand.




Fig. 1. Example of right hand image, hand segmentation and edge and point detection

To correctly detect the tips and valleys of the fingers we use a method presented in [5] that uses radial distance with respect to a reference point around the wrist region to detect finger key points. This sequence yields maximum and minimum points that allow us to easily detect the desired key points.

Having located all key points by the previously analysis we compute the desired distances between points. The main distance group computed across hand key points can be seen in the next diagrams (see Fig. 2 and Fig. 3).

Basically the first 20 components (named V1 to V20) correspond to finger length. These lengths are computed from the tip of each of the fingers to the valleys and central part. Two variables in some of the arrows depicted can be observed. The reason is because we use two similar points that are not distant

enough to be appreciate in Fig. 2. One of them is computed as an average point from the two valleys of the finger. The other one is measured across the intersection point between the eigenvector directions defined by the fingers and the two valley points related to finger.

The other group of distances is width (from V21 to V27). The next group (see Fig. 3) of distances are mix of widths and lengths between the center point of the finger (square points) and the center of the hand. These variables range from V28 to V42. Observe that the last ones have not been represented for clarity purposes.

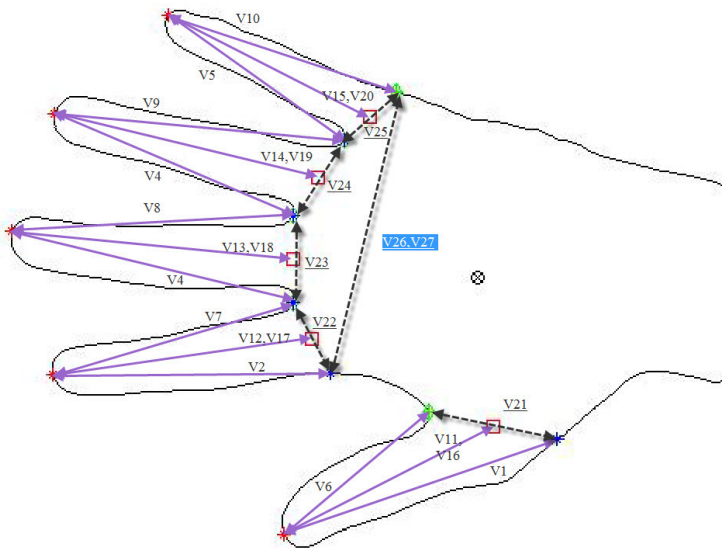


Fig. 2. Hand Measures – I

We have highlighted the variables V26 and V27 because they will be the main variables for gender recognition as we will show in the next results section. It is coherent with what people can usually think, the greater the hand (hand width) the greater the chance to be a man.

In addition, there is a couple of additional computed distances, not yet depicted because they are not significant enough.

3 Biometric Dispersion Method

This method [1] [2] has been proved in recognition biometric applications but not in gender recognition. Its main advantage which is that it defines a selection criteria for all the attributes in order to select those who provide better discrimination power.




Fig. 3. Hand Measures – II

BDM is a simple method linked to Discriminant Analysis that works under the construction of two groups:

- Equal group: where we have pairs of measurements of a characteristic related to the same person.
- Unequal group: where we have pairs of measurements of a characteristic related to different persons.

For a quick understanding of the method, let us suppose we want to use middle finger length (V13 variable). We have different statistics behavior for this measure when computed across the same group than when it is computed across all available data. These differences are used to build the classifier.

These distributions follow a Gaussian distribution with the following parameters:

$$p(x|U) = N(x|0.2(s_p^2 + s_i^2)), \tag{1}$$

$$p(x|E) = N(x|0.2s_i^2), \tag{2}$$

where σ_i is the variance of V13 from the same group (over its samples), and σ_p is the variance of V13 across the whole user population.

The expression $s_i^2 / (s_p^2 + s_i^2)$ is related to the discriminative power of the characteristic component under consideration. BDM uses this expression to select the best feature vector components. It is important to point out the way BDM will conduct the selection criteria. Through a $\sigma_{\text{threshold}}$ all components greater than this threshold will be discarded.

The method performs the classification by applying a discriminative function (3) that let us to define an easy rule to classify:

$$g(x) = \ln \left(\frac{p(x|U)}{p(x|E)} \right) + \ln \left(\frac{p(U)}{p(E)} \right). \quad (3)$$

Then the decision is straightforward: U if $g(x) > 0$, E otherwise.

4 Experimental Results

This section provides the experimental rates of the algorithm proposed. It is important to remark that there is an unbalanced design, because we have more men (68 %) than women (32 %). Fig. 4 shows the distribution of genders.




Fig. 4. Gender Distribution

In order to observe how BDM works, a test increasing the size of the training set will be performed. Starting with a training set of 27 hands (not necessarily from different users) that was distributed between 18 Men and 9 Women. With these initial conditions we train a BDM. Then a test analysis over the remaining data set is applied. It means a total test set of 1013 hands. Then the training set is increased with 36 men and 19 women and so on until we reach a maximum of 40 % of the total data set.

Additionally we want to compare the results of the BDM method with other well known methods like K- nearest neighbor and a logistic classifier. The last two methods take the full characteristic vector as an input in order to performing the gender recognition task.

Fig. 5 shows how the three different methods perform gender recognition through the different number of women present in the training set. Logistic regression does not perform good recognition when the training data set is not




Fig. 5. Performance across different number of Women Present in Training Set

large enough. However, as training set increases, the recognition rate of the logistic classifier improves. K Nearest Neighbor performs quite well across all applied tests and has slightly better results than BDM.

Table 1. BDM Performance Rates

Training Data		BDM	
M	W	dim	Ident.r Error.r
18	9	5	0.9519 0.0481
36	17	26	0.9189 0.0811
54	25	25	0.9116 0.0884
71	33	7	0.9498 0.0502
89	42	7	0.9516 0.0484
107	50	11	0.9434 0.0566
125	58	39	0.9405 0.0595
142	66	20	0.9639 0.0361
160	75	21	0.9652 0.0348
178	83	36	0.9589 0.0411
196	91	36	0.9575 0.0425
214	100	34	0.9656 0.0344
231	108	40	0.9700 0.0300
249	116	42	0.9748 0.0252
267	124	36	0.9707 0.0293
284	132	36	0.9776 0.0224

The main point is that BDM works in each situation presented with the best variables according to its discriminant threshold. Table 1 shows how BDM works across the different training data sets.

One of the advantages to work with BDM is that it is possible to check how each feature component works. Table 2 is presented, which shows the importance of the 15th most relevant features. As a summary:

Table 2. Selected feature components

%Train	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
2.5	27	26	34	35	42	33	22	1	11	4	19	20	10	14	17
5.0	27	26	42	35	1	11	33	22	20	5	10	15	14	34	9
7.5	26	34	27	35	17	7	2	12	5	20	10	19	33	15	14
10.0	26	27	35	34	33	38	1	42	23	20	4	19	5	11	14
12.5	27	26	35	34	42	33	1	11	41	4	19	14	39	23	36
15.0	26	27	35	34	33	42	1	38	11	41	23	20	5	39	10
17.5	26	27	35	34	42	33	1	38	11	4	41	20	23	19	22
20.0	26	27	35	34	42	33	1	41	23	11	22	4	3	19	36
22.5	27	26	42	35	34	33	1	11	23	41	4	38	19	14	5
25.0	27	26	35	34	42	33	1	11	23	5	20	10	4	15	19
27.5	27	26	35	34	1	38	42	33	11	23	20	19	5	10	14
30.0	27	26	35	34	42	33	1	4	11	19	14	23	41	38	3
32.5	27	26	35	34	42	1	11	33	4	20	5	10	19	15	14
35.0	26	27	35	34	42	33	1	41	38	4	23	11	19	14	20
37.5	27	26	35	34	42	1	33	38	11	4	19	14	41	20	5
40.0	26	27	35	42	34	1	33	38	11	23	4	41	19	14	20

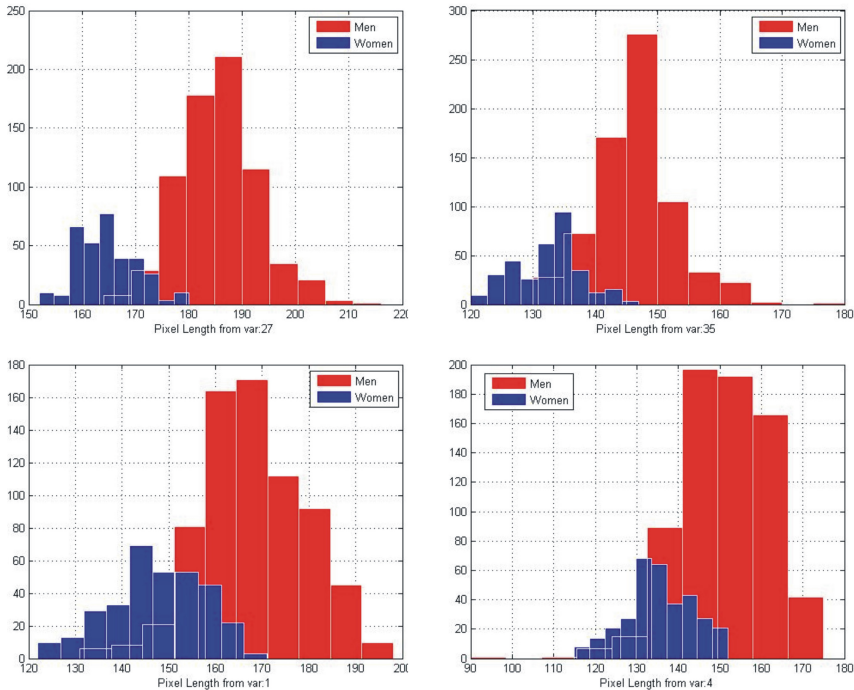


Fig. 6. Distributions with high discrimination

- Components V26, V27 and V35 related to hand width.
- Components V42 and V34 (V42 not depicted in Fig. 3 goes from pinky finger square to the center of the hand), V34 goes from index finger square to the center of the hand.




Fig. 7. Distributions overlapped

- Components V33 and V38 equals to the length between index and middle finger (V33) and between index and pinky finger (V38).
- Components V1, V11 and V4 lengths from thumb finger (V1 and V11) and from index finger (V4).

Finally it is shown how the different feature components are distributed depending on gender. This distribution gives us a clear understanding of the way BDM discard irrelevant data. For instance, components V27 and V35 (see Fig. 6) have a clear distribution. It means it will be easier to classify gender according to these feature components.

That is not the case when both distributions are overlapped as the distribution of V48 shows (see Fig. 7). Across the measurements the mean of men group and women group are the same. Thus, it is not possible to use this feature component (alone) to separate gender.

5 Conclusions

This results have proved the reliability to use anthropometric measures of hand images to recognize gender. Performance results are better than those coming from face analysis [4] and are promising when using a combination with some additional information. By means of a fusion algorithm it will be possible to raise identification rates above 99%.

It has also been proved the suitability of BDM to detect the major feature components. In this context, this classifier contributes to achieve a coherent explanation of which component contain better information for the recognition process.

Acknowledgments. This work has been supported by FEDER and MEC, TEC2009-14123-C04-04.

References

1. Fàbregas, J., Faundez-Zanuy, M.: Biometric Dispersion Matcher. *Pattern Recogn.* 41, 3412–3426 (2008)
2. Fàbregas, J., Faundez-Zanuy, M.: Biometric Dispersion Matcher Versus LDA. *Pattern Recogn.* 42, 1816–1823 (2009)
3. Otsu, N.: A Threshold Selection Method from Gray-level Histograms. *IEEE Transactions on Systems, Man and Cybernetics* 9(1), 62–66 (1979)
4. Smirg, O., Mikulka, J., Faundez-Zanuy, M., Grassi, M., Mekyska, J.: Gender Recognition Using PCA and DCT of Face Images. In: Cabestany, J., Rojas, I., Joya, G. (eds.) *IWANN 2011, Part II. LNCS*, vol. 6692, pp. 220–227. Springer, Heidelberg (2011)
5. Yoruk, E., Konukoglu, E., Sankur, B., Darbon, J.: Shape-based hand recognition. *IEEE Transactions on Image Processing* 15(7), 1803–1815 (2006)

Revisiting AVEC 2011 – An Information Fusion Architecture

Martin Schels, Michael Glodek, Friedhelm Schwenker, and Günther Palm

Institute of Neural Information Processing, Ulm University, Germany
firstname.lastname@uni-ulm.de

Abstract. Combining information from multiple sources is a vivid field of research. The problem of emotion recognition is inherently multi-modal. As automatic recognition of the emotional states is performed imperfectly by the single mode classifiers, its combination is crucial. In this work, the AVEC 2011 corpus is used to evaluate several machine learning techniques in the context of information fusion. In particular temporal integration of intermediate results combined with a reject option based on classifier confidences. The results for the modes are combined using a Markov random field that is designed to be able to tackle failures of individual channels.

1 Introduction

For pattern recognition applications from multiple independent modalities one of the most challenging issues is that often complete sensor information at any time cannot be assumed. It is mandatory for a fusion architecture to deal with missing or uncertain classifications which, in the worst-case, may be even conflicting. Such a scenario is given by the Audio/Visual Emotion Challenge (AVEC) 2011 data collection. The data set was recorded in natural and unrestricted conditions and labeled in four affective dimensions: *arousal*, *expectation*, *power* and *valence* [1]. The challenge was composed of three sub-challenges: audio emotion recognition on audio-visual, prediction based on the video channel for every frame and multi-modal emotion recognition.

The presented work studies a new design of a fusion architecture. Several machine learning strategies such as reject options, temporal integration of intermediate results and an information fusion architecture, which implements the concepts of sparse decision fusion are evaluated with the data at hand. Sparsity in this context means that a prediction from one or more individual classifiers is not always available. This naturally occurs in this context as e.g. audio processing makes only sense when the subject is speaking. As a result the fusion has to cope with scattered decisions over time which need to be interpolated appropriately. Methods on the implementing of the reject option along with a newly developed method of temporal decision fusion using a Markov Random Field (MRF) are presented. The proposed architecture is studied using the AVEC 2011 corpus.

2 Methods

2.1 Classification Using a Reject Option

As automatic classification using statistical models is often performed imperfectly, hence it is crucial to deal with the possibility of misclassification. On the one hand, it is common to define a loss function to minimize the cost of a misclassification. However, it may as well be convenient, especially when the cost of a falsely classified sample is high, not to make any decision for a presented sample. This approach is called “classification with reject option” [5,6]. Rejection in this context means deciding (yes or no) whether a certain confidence level has been achieved or not. A very intuitive way to measure the confidence is to use a classifier with probabilistic output. Various attempts have been made to introduce confidence based rejection criteria. However, the standard approach is based on thresholds determined by heuristics on probabilistic classifier outputs [12].

2.2 Classifier Fusion Using Markov Random Fields

We propose to use Markov Random Field (MRF) [3] for the combination of multiple decisions over time. Figure 1 shows an exemplary layout of the proposed MRF. The variable x corresponds to the estimated decision obtained by the combination of multiple sources. In order to model the dynamics of the classification over time, the variables for each time step x_t are connected to a chain. Within this figure two “streams of decisions” y_1 and y_2 are connected to the chain whenever a classifier was able to give a decision (these time steps are given by the sets \mathcal{L}_1 and \mathcal{L}_2). In case decisions are missing, the information is propagated from the known decisions via the chain to the imperfectly supported nodes.




Fig. 1. Graphical model of the Markov Random Field for classifiers fusion. The estimated decision x_t is influenced by the available decisions $y_{m,t}$ of the modality m and $t \in \mathcal{L}_m$ and the adjacent estimates x_{t-1} and x_{t+1} .

The MRF is defined by two potentials Ψ and Φ . The potential Ψ_m of the modality m forces the estimated outcome x_t to take the same value as the available decisions y_{mt} . It is given by

$$\Psi = \sum_m \Psi_m = \sum_m \sum_{t \in \mathcal{L}_m} k_m (y_{mt} - x_t)^2,$$

where \mathcal{L}_m is the set of time steps in which a decision for m is available and k_m is a parameter describing the strength of its influence. The second potential Φ enforces the lateral similarity and is given by

$$\Phi = \sum_{t=1}^T \sum_{i \in N(t)} w_{t+i} (x_t - x_{t-(1-2i)})^2,$$

where $\mathbf{w} \in \mathbb{R}^{T-1}$ is a vector weighting the smoothness between two adjacent nodes and the set $N(t)$ contains the value 0 in case x_t has a adjacent node x_{t-1} and the value 1 in case x_{t+1} exists. The parameter \mathbf{w} can be set by external knowledge at design time, e.g. to enforce similarity within one period of time.

The joint distribution of the estimated vector \mathbf{x} given the decisions of the modalities is defined by

$$p(\mathbf{x}|\mathbf{y}_1, \dots, \mathbf{y}_M) = \frac{1}{Z} \exp\left(-\frac{1}{2}(\Psi + \Phi)\right).$$

If there one is not interested in the probability itself, but in minimizing the mode of the log-posterior probability, the variable \mathbf{x} can be determined using gradient descent [11].

3 “AVEC 2011” Data Collection

The data used in this study is the one provided within the Audio/Visual Emotion Challenge (AVEC) 2011 in the ACII 2011 workshop [119]. Over-all three sub-challenges were proposed: an audio challenge on word-level, a video challenge on frame-level and an audiovisual also on video frame-level.

The data was recorded in a human computer interaction scenario in which the subjects were instructed to interact with an affectively colored artificial agent. Audio and video material was collected in over-all 63 recordings from 13 different subjects. The recorded data was labeled in four affective dimensions: *arousal*, *expectancy*, *power* and *valence*. The annotations of the raters have been averaged for each dimensions resulting in a real value for each time step. Subsequently, the labels are binarized using a threshold equal to the grand mean of each dimension. Every recording was annotated by two to eight raters. Along with the sensor data and annotation, a word-by-word transcription of the spoken language was provided which partitions the dialog into conversational turns. For the evaluation of the challenge only arousal was taken into account as classification of the other dimensions yielded poor results [1].

3.1 Audio Features

In the following the three feature sets, that are calculated from audio signal are described.

- The *fundamental frequency* values are extracted using the f_0 tracker available in the ESPS/waves+ [2] software package. Besides the f_0 track, the *energy* and the *linear predictive coding* (LPC) of the plain wave signal is extracted [10]. All three features are concatenated to a ten-dimensional early fusion feature vector.

¹ <http://sspnet.eu/avec2011/>

² <http://www.speech.kth.se/software/>

- The *Mel frequency cepstral coefficient* (MFCC) representation is inspired by the biological known perceptual variations in the human auditory system. The perception is modeled using a filter bank with filters linearly spaced in lower frequencies and logarithmically in higher frequencies in order to capture the phonetically important characteristics of speech. The MFCC are extracted as described in [13].
- The *perceptual linear predictive* (PLP) analysis is based on two perceptually and biologically motivated concepts, namely the critical bands, and the equal loudness curves. Frequencies below 1 kHz need higher sound pressure levels than the reference, and sounds between 2 - 5 kHz need less pressure, following the human perception. The critical band filtering is analogous to the MFCC triangular filtering, apart from the fact, that the filters are equally spaced in the Bark scale (not the Mel scale) and the shape of the filters is not triangular, but rather trapezoidal. After the critical band analysis and equal loudness conversion, the subsequent steps required for the relative spectral (RASTA) processing extension, follow the implementation recommendations in [7]. After transforming the spectrum to the logarithmic domain and the application of RASTA filtering, the signal is transformed back using the exponential function.

3.2 Video Features

The features from the video channel were computed using the well-known computer expression recognition toolbox (CERT) [8] which is designed to recognize emotion related facial processes (such as action units or basic emotions) but also general attributes of a subject (e.g. gender, glasses present). The output of modules “Basic Emotions 4.4.3”, “FACS 4.4”, “Unilaterals” and “Smile Detector” are extracted to infer the labels of the AVEC corpus. Over-all a 36-dimensional vector for a frame was obtained. CERT only delivers sound values in case the face of a subject is recognized. Due to the unrestricted settings of the AVEC corpus (subjects may turn away or leave the visual range of the camera) classification results are missing from time to time.

4 Experiments and Results

All experiments are conducted in a strict *subject independent* cross validation, i.e. the subjects that are in the training set do *not* occur in the respective test set.

4.1 Spoken Utterances

To classify the speech, the available data was pooled in a word-level granularity and an approach as described in [2] was conducted to overcome different lengths of the data. The three different feature sets defined in Section 3.1 are separately transformed into an uniform vector by sampling randomly reference sequences. Based on these subsets three hidden Markov models (HMM) are constructed. The final feature vector of a

sample sequence is then the concatenated log-likelihoods of the HMM. The three resulting new feature sets are then put together to a new data matrix. The final classification was performed using a random forest per dimension [4].

The result of this word-wise classification is given in Table 1. The accuracies of the classifications are given together with the F_1 scores for each dimension and its complementary assignment [5]. While all accuracies are over 50% (up to 59.6% for arousal), the F_1 scores are relatively unbalanced for all labels but arousal.

Table 1. Accuracies and F_1 scores for the different emotional categories for the audio classification on a word basis. Also the standard deviation is displayed.

Label	Accuracy	F_1 score variable set	F_1 score variable not set
Arousal	0.591 (0.040)	0.575 (0.220)	0.455 (0.233)
Expectancy	0.572 (0.032)	0.562 (0.196)	0.492 (0.161)
Power	0.528 (0.053)	0.506 (0.235)	0.440 (0.196)
Valence	0.615 (0.051)	0.579 (0.281)	0.418 (0.254)

In order to further improve the results, a temporal integration is evaluated which exploits the turns of the human computer dialog. Since it can be assumed that the label will not change within a turn, the decisions of a turn are combined by computing the average of the probabilistic output of the classifier (i.e. the scores of the individual trees of the random forest).

The result of this turn based classification is shown in Table 2. Here, a turn comprises 1–201 spoken words with an average turn size of 15 words. Generally, the pure accuracies are increased, but unfortunately only for the label arousal the F_1 scores increase symmetrically. For the other categories, the imbalanced values of the F_1 score further intensifies. Furthermore, the standard deviations increase.

Table 2. Accuracies and F_1 scores for the different emotional categories for the audio classification on a *turn* basis

Label	Accuracy	F_1 score variable set	F_1 score variable not set
Arousal	0.661 (0.066)	0.578 (0.358)	0.395 (0.396)
Expectancy	0.597 (0.048)	0.540 (0.330)	0.453 (0.247)
Power	0.550 (0.146)	0.514 (0.348)	0.356 (0.394)
Valence	0.666 (0.076)	0.584 (0.398)	0.325 (0.357)

4.2 Facial Expressions

The feature vector to recognize the emotional dimensions from video were generated by CERT. In about 8% of the video frames the toolbox was unable to locate the subject's

³ The F_1 score is given by the harmonic mean of the recall R and precision P . Since the F_1 score does not take the true negatives into account, the score for the true and the complementary label assignments are listed.

face, such that no features have been generated for these frames. The classification of the four dimensions is conducted using a naive Bayes classifier based on a Gaussian distribution [3].

In this experiment, accuracies up to 60% were achieved (compare Table 3). The F_1 scores show similar behaviors like in Section 4.1: the arousal has a balanced F_1 scores for both, label is set and label is not set.

Table 3. Accuracies and F_1 scores of the frame-wise classification of the facial expression using the video channel

Label	Accuracy	F_1 score variable set	F_1 score variable not set
Arousal	0.568 (0.041)	0.607 (0.048)	0.509 (0.094)
Expectancy	0.549 (0.042)	0.496 (0.096)	0.567 (0.108)
Power	0.559 (0.030)	0.576 (0.112)	0.490 (0.122)
Valence	0.600 (0.075)	0.670 (0.117)	0.439 (0.069)

As the classification for each frame appears to be relatively weak and it is not expected that the label alternates between two consecutive frames, it is intuitive not to make a decision for all presented samples. Hence, a reject option strategy is implemented. The classification of the frame is assessed with respect to the estimated posterior probability.

Table 4 shows the result of this approach when 50% of the test samples are rejected. In general, the accuracy increases when rejecting samples what indicates the reliability of the confidence measure used. On the other hand one can observe already an increase of the imbalances in the F_1 scores. This means that the classifier is more certain when deciding for one case than the other.

Table 4. Accuracies and F_1 scores of the frame-wise classification of the facial expression using the video channel with a reject option of 50% of the frames

Label	Accuracy	F_1 score variable set	F_1 score variable not set
Arousal	0.608 (0.069)	0.678 (0.075)	0.471 (0.132)
Expectancy	0.575 (0.061)	0.509 (0.157)	0.572 (0.157)
Power	0.584 (0.043)	0.593 (0.145)	0.488 (0.157)
Valence	0.667 (0.128)	0.747 (0.146)	0.379 (0.083)

Table 5 summarizes the classification when 99% of the testing samples are rejected. This high portion can be justified as a video channel delivers relatively quick information compared to the observed emotional phenomena. Accuracies of more than 80% can be achieved. However, the rejection also enforces the issue, that emerged before: when the labels arousal or valence are not set it can happen that the rejection mechanism does not allow to classify a single frame into this category.

Table 5. Accuracies and F_1 scores of the frame-wise classification of the facial expression using the video channel with a reject option of 99% of the frames

Label	Accuracy	F_1 score variable set	F_1 score variable not set
Arousal	0.812 (0.056)	0.886 (0.025)	0.224 (0.410)
Expectancy	0.614 (0.077)	0.493 (0.335)	0.455 (0.383)
Power	0.566 (0.140)	0.499 (0.333)	0.315 (0.344)
Valence	0.862 (0.099)	0.924 (0.060)	0.025 (0.050)

Table 6. Accuracies and F_1 scores for the different emotional categories for the fusion using MRF word-level fusion

Label	Accuracy	F_1 score variable set	F_1 score variable not set	Rejection rate
Arousal	0.651 (0.047)	0.696 (0.031)	0.586 (0.089)	99%
Expectancy	0.573 (0.048)	0.438 (0.087)	0.645 (0.075)	0%
Power	0.572 (0.042)	0.624 (0.101)	0.453 (0.115)	0%
Valence	0.607 (0.076)	0.695 (0.091)	0.392 (0.050)	50%

Table 7. Accuracies and F_1 scores for the different emotional categories for the fusion using MRF turn-level fusion

Label	Accuracy	F_1 score variable set	F_1 score variable not set	Rejection rate
Arousal	0.677 (0.067)	0.723 (0.050)	0.608 (0.114)	99%
Expectancy	0.574 (0.050)	0.436 (0.097)	0.644 (0.081)	0%
Power	0.571 (0.044)	0.626 (0.101)	0.446 (0.116)	0%
Valence	0.606 (0.091)	0.712 (0.096)	0.312 (0.081)	50%

4.3 Multi-modal Fusion

There are several challenges for a combiner that aims at returning a decision for the smallest time granularity (i.e. a frame) in this setting. Classifiers may take advantage of a reject option and hand over only confident decisions. The sensor sample rate, feature extraction and classification may lead to significant time offsets between the classifier's decisions or sensors may fail due to various reasons (e.g. technical malfunction, person is not in range etc.). As a result, the fusion needs to handle incomplete inputs over time for every modality. For these circumstances, the proposed MRF (compare Figure 1) is an intuitive selection. Furthermore, the ability to control the dynamics of the classification over time is exploited: the lateral similarity of the over-all decision was enforced within a turn, while changes between turns in the estimate were amplified. This in a sense emulates the setting in Section 4.1 more flexible.

Table 6 and Table 7 show the results of the combination of the audio on word- and turn-level decisions with different results of the video modality using different rejection rates.

Since the results for the individual modalities in Section 4.1 and 4.2 are based on rejection or word and turn-level, it is difficult to compare them directly to the outcome of the fusion step. In comparison to Table 3 (video channel without rejection) all accuracies are outperformed. The increase can be explained by the audio classification results and to the averaging of the decisions by the MRF. The influence of the audio recognition

is reflected by the slight decreases in the F_1 score compared to results on video. Furthermore, it can be seen that the benefits of the reject option and word- and turn-level decisions is strongly dependent on the characteristics of the classifiers. For instance in case of the dimension *arousal*, the audio classification performed well enough such that only confident decisions of the video channel could improve the result. The classifiers of the dimensions *expectancy* and *power* are both very weak, the combination of all available information, however, did mitigate the imbalanced F_1 scores for audio and also improved the performance obtained on the video channel. The dimension *valence* is improved using a compromise. Hereby only 50% of the decisions for the video channel are rejected. However, the audio recognition results along with the smoothing of the estimation of the MRF could improve the performance.

5 Discussion and Conclusion

In this paper, the AVEC 2011 corpus was re-examined by evaluating different machine learning techniques in order to improve classification. We used temporal integration of intermediate results, reject options in individual classifiers based on the respective confidence and a MRF for the combination of the classification of audio and video. As the outcome of the channel shows, the accuracies on this data set are generally relatively low. However, the increase of the performance was possible by conducting a post processing step. The reject options in classifiers was applied such that they do no longer contribute to the error statistics. The sparse decisions over time then were fused using a MRF and it could be shown that the multi classifier system performance could be improved using these approaches.

Combining the different outputs with a MRF is promising as the model is both stable and flexible enough to handle missing values. Furthermore, it is beneficial that expert knowledge is encoded, e.g. knowledge about the dynamics of the labels. Future work will extend the information fusion architecture using a shifting a window approach to handle on-line classification tasks.

Acknowledgment. This paper is based on work done within the the Transregional Collaborative Research Centre SFB/TRR 62 “Companion-Technology for Cognitive Technical Systems” funded by the German Research Foundation (DFG).

References

1. Besag, J.: On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society. Series B (Methodological)* 24(3), 259–302 (1986)
2. Bicego, M., Murino, V., Figueiredo, M.A.T.: Similarity-Based Clustering of Sequences Using Hidden Markov Models. In: Perner, P., Rosenfeld, A. (eds.) *MLDM 2003*. LNCS (LNAI), vol. 2734, pp. 86–95. Springer, Heidelberg (2003)
3. Bishop, C.M.: *Pattern Recognition and Machine Learning*, vol. 4, ch. 8.3, pp. 383–392. Springer (2006)
4. Breiman, L.: Random forests. *Machine Learning* 45(1), 5–32 (2001)

5. Chow, C.: On optimum recognition error and reject tradeoff. *IEEE Transactions on Information Theory* 16(1), 41–46 (1970)
6. De Stefano, C., Sansone, C., Vento, M.: To reject or not to reject: that is the question-an answer in case of neural classifiers. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews* 30(1), 84–94 (2000)
7. Hermansky, H., Morgan, N., Bayya, A., Kohn, P.: RASTA-PLP speech analysis technique. In: *Proc. of the Int. Conf. on Acoustics, Speech, and Signal Processing*, vol. 1, pp. 121–124. IEEE (1992)
8. Littlewort, G., Whitehill, J., Wu, T., Fasel, I., Frank, M., Movellan, J., Bartlett, M.: The computer expression recognition toolbox (CERT). In: *Proc. of the Int. Conf. on Automatic Face & Gesture Recognition and Workshops*, pp. 298–305. IEEE (2011)
9. McKeown, G., Valstar, M.F., Cowie, R., Pantic, M.: The SEMAINE corpus of emotionally coloured character interactions. In: *Proc. of the Int. Conf. on Multimedia and Expo (ICME)*, pp. 1079–1084. IEEE (2010)
10. Rabiner, L., Juang, B.H.: *Fundamentals of speech recognition*. Prentice Hall (1993)
11. Schuller, B., Valstar, M., Eyben, F., McKeown, G., Cowie, R., Pantic, M.: AVEC 2011–The First International Audio/Visual Emotion Challenge. In: D’Mello, S., Graesser, A., Schuller, B., Martin, J.-C. (eds.) *ACII 2011, Part II*. LNCS, vol. 6975, pp. 415–424. Springer, Heidelberg (2011)
12. Thiel, C., Schwenker, F., Palm, G.: Using Dempster-Shafer Theory in MCF Systems to Reject Samples. In: Oza, N.C., Polikar, R., Kittler, J., Roli, F. (eds.) *MCS 2005*. LNCS, vol. 3541, pp. 118–127. Springer, Heidelberg (2005)
13. Zheng, F., Zhang, G., Song, Z.: Comparison of different implementations of MFCC. *Journal of Computer Science and Technology* 16(6), 582–589 (2001)

Discriminating Human vs. Stylized Emotional Faces: Recognition Accuracy in Young Children

Anna Esposito^{1,2}, Maria Teresa Riviello^{1,2}, and Vincenzo Capuano^{1,2}

¹ Seconda Università di Napoli, Department of Psychology, Caserta, Italy

² International Institute for Advanced Scientific Studies (IIASS), Vietri sul Mare, Italy
iiass.annaesp@tin.it, mariateresa.riviello@unina2.it,
vincenzo.capuano85@gmail.com

Abstract. This paper intends to contribute to the research on the perception of emotion with a case study on the recognition of realistic vs. stylized facial emotional expressions in typical developing three and six-year-old children. In particular, it reports on two perceptual experiments aimed at evaluating children's ability in identifying human and stylized male and female facial emotional expressions of happiness, anger and surprise. Results show that six-years-old children are able to recognize facial expressions of happiness and anger exploiting stylized as well as realistic human figures, preferring stylized faces for the identification of surprise. Three-year-old children are not able to recognize surprise and are significantly better in recognizing happiness rather than anger, suggesting that the ability to recognize certain emotional faces emerges through experience. In addition, it also suggests that this ability is not affected by face gender.

Keywords: Facial emotional expression, Recognition accuracy, Typical developing children.

1 Introduction

Emotion understanding plays an important role in many aspects of human communication and interaction. It is considered part of a general socio-cognitive competence, which concerns the ability to understand people's intentions, states of mind, beliefs, desires, and emotions [7], in order to respond adequately to social stimuli and communicate efficiently. Since facial expressions represent an important information of the individual's inner state [5], children's ability to decode them is crucial for successful social interaction and functioning [13].

As emotion identification is an important predictor of the health in the development of children's social competences [6], a consistent part of the research has been devoted to their ability to perceive and interpret facial expressions of emotion. Focus has been set on the association among anomalies in emotion recognition and socio-cognitive disorders, learning and medical disabilities [2, 4], and psychiatric disorders, including autism and schizophrenia [1, 3]. Recently Pons and

Harris [16] addressed the emotional perception ability in typical developing children in nursery and primary schools, creating an assessment tool which examines children's understanding of emotion: the Test of Emotion Comprehension (TEC) [15]. The TEC consists of a picture book with a cartoon scenario on the top of each page. Below each scenario, are four emotional results represented as drawing figures of facial expressions. This instrument evaluates each child's ability to: 1) recognize facial emotional expressions; 2) comprehend external causes; 3) understand desire-based emotions; 4) understand belief-based emotions; 5) understand the influence of reminders; 6) comprehend the regulation of an experienced emotion; 7) understand the possibility of hiding an emotion; 8) understand mixed emotions; and 9) understand moral emotions. The TEC allows the measurement of children's understanding of these different components between the ages of 3 and 12 years using a series of tests on forced choice items.

The study presented in this article focuses on the first component of the emotion understanding and starts with the premise that children are capable of recognizing emotions through facial expressions [15-16]. In particular it investigates the ability to recognize facial emotional expressions in typical three and six-year-old children, exploiting both stylized and human faces in order to gauge possible variations in the recognition of emotions when more realistic expressions are examined. In addition, it explores possible developmental differences in the recognition of such emotional facial expressions that may depend on the emotion category, and finally possible influences of the face gender on the children's performance.

2 Material and Procedure

In this study three and six-year-old children were involved in a recognition task with the goal to assess their ability to identify emotional facial expressions of happiness, anger and surprise (these emotions are considered as basic and universal by several theories on emotions [9-11, 14, 17-19]). The abovementioned emotional states were portrayed by photos of human faces and stylized ones.

The main goal was to investigate possible developmental differences in the recognition performance when realistic/non-realistic emotional facial expressions were exploited. Influences of the face gender on the children's recognition accuracy were also investigated.

2.1 Participants

The study was conducted on two different groups of Italian children, a group of three-year-olds recruited from the 1° Circolo Didattico di Poggiomarino (Na), and a group of six-year-olds all first elementary school students from Maria Santissima di Casaluce in Frattamaggiore, Italy. The three-year-old children were initially 20 (10 males and 10 females). However, not all of them were able to perform the task and the reported data were collected from 14 of them. In order to allow the three-year-olds to assign a label to the faces, a set of 10 questions was defined and presented as a joke

to the children for each facial expressions. The questions did not refer to the emotions under examination, in an attempt to allow the children to assign the correct emotional label through their own words. For each face, the experimenter proposed the questions until the child had given the correct answer, or until all of them had been asked. In this latter case, if the right label was not elicited, it was assumed that the child was not able to decode the emotional face. The six-year-olds were 28 (14 females and 14 males) and the labeling procedure was the same. The questions, always the same for each child, were in Italian (see Appendix) and written on a grid paper where the experimenter also reported the collected answers. All participants were children with not known learning or language disorders problems, as well as no personality or behavioral problems. Family background ranged mainly from working to middle class. The parents gave their informed consent, and the ethic committee of the Second University of Naples granted the permission to conduct the experiments.

2.2 Materials

A total of 12 static facial expressions were used: 6 were photos of human faces, extracted from the Facial Action Coding System (FACS, [8]), and 6 were stylized faces, extracted from a comic book as reported in Esposito [12]. In each group, each emotion under consideration was expressed by a male and a female face.

2.3 Procedure

In order to verify whether the ability to decode facial emotional expressions is affected by the type of expressions exploited (human vs. stylized faces), participants were divided into two groups. The three-year-olds were split into a group of 6 (3 males and 3 females) that was tested on the stylized faces and a group of 8 (6 females and 2 males) that was tested on the human photos. Similarly, the six-year-olds were divided into two groups of 14, each composed of 7 girls and 7 boys and tested on the human and the stylized faces respectively. The faces were randomly presented on a computer screen, and the children were requested to label the emotion expressed by them. In order to familiarize with the children, the experimenter spent two months with them, before starting the data collection.

3 Results for the Six-Year-Olds

Figure 1 illustrates the percentage of correct answers obtained by the two six-year-old groups, tested on the human and stylized facial expressions, respectively. The recognition accuracy was statistically assessed through a χ^2 test in order to investigate for significant differences in the recognition of human vs. stylized emotional faces.

The χ^2 test is a statistic method used to evaluate whether two or more independent groups have the same proportion of categorical responses for one or more independent variables. The test exploits $m \times n$ contingency tables. In our specific case the data are reported in a 2×2 contingency table - 2 populations (1 for the stylized and one for the human faces) \times 2 category levels (correct recognition/no correct

recognition) - and the null hypothesis states that each population has the same proportion of observations.

The results of the test, along with the degrees of freedom (Df), are compared with a standard previously calculated χ^2 table of distributions (χ^2 Critique) in order to verify whether the p-value is less than the significance level α assumed a priori to reject the null hypothesis. To prevent overestimation of statistical significance for small data samples, as in our case, the Yates correction factor can be applied. This correction factor reduces the χ^2 value, and therefore, increases the resulting p-value (for a detailed description of the procedure see [22])

The χ^2 test on our data was computed separately for each of the abovementioned emotion categories. Results are displayed in Table 1.




Fig. 1. Percentage of recognition accuracy in the 6-year-olds for human vs. stylized emotional faces

Table 1. Statistical assessment of the 6-year-old recognition accuracy in discriminating human vs. stylized emotional faces of happiness, anger, and surprise

χ^2 test Human vs. Stylized Faces				
	χ^2	p-value	Yates χ^2	Yates p-value
Happiness	0,583	0,445	0,146	0,702
Anger	0	1	0,269	0,6
Surprise	5,853	0,015	4,625	0,03
	Df=1	χ^2 Critique= 3,84		$\alpha=0,05$

There is a significant effect of the human vs. stylized faces only for surprise, which is better decoded in the stylized faces. In addition, as it is reported in Table 2, the χ^2 test comparing the child recognition accuracy in decoding happiness, anger, and surprise expressed by both human and stylized faces, shows that surprise is also the emotion which is less easily identified by the children.

A χ^2 test comparing the significance of face gender shows that the children’s recognition accuracy is not affected by the gender of the face (male or female) both for the human and stylized stimuli (see Table 3).

Table 2. Statistical assessment of the 6-year-old recognition accuracy in discriminating among happiness, anger, and surprise expressed by both human and stylized faces

χ^2 test Happiness/Anger/ Surprise									
	Human Figures				Stylized figures				
	χ^2	p-value	Yates χ^2	Yates p-value	χ^2	p-value	Yates χ^2	Yates p-value	
Happiness/Anger	1,469	0,225	0,653	0,419	0,22	0,64	0	1	
Happiness/ Surprise	16,258	0	14,163	0	6,09	0,013	4,667	0,0307	
Anger/ Surprise	24,257	0	21,636	0	8,11	0,004	6,411	0,0113	
	Df=1				χ^2 Critique= 3,84				$\alpha=0,05$

Table 3. Statistical assessment of the 6-year-old recognition accuracy in discriminating male and female human/stylized emotional faces of happiness, anger, and surprise

χ^2 test for Face gender									
	Human Figures				Stylized figures				
	χ^2	p-value	Yates χ^2	Yates p-value	χ^2	p-value	Yates χ^2	Yates p-value	
Female/Male Happiness	2,191	0,138	0,974	0,324	0,373	0,5411	0	1	
Female/Male Anger	2,154	0,142	0,538	0,463	2,154	0,142	0,538	0,463	
Female/Male Surprise	0	1	0,175	0,675	1,348	0,245	0,599	0,438	
	Df=1				χ^2 Critique= 3,84				$\alpha=0,05$

4 Results for the Three-Year-Olds

For the three-year-old children we only investigated on their ability to recognize emotions and the possible effects of the human vs. stylized faces, neglecting the gender of the faces given the small number of children involved. Figure 2 illustrates the percentage of the three-year-old recognition accuracy for each of the emotions under examination on both human and stylized faces.

Figure 2 shows that at the age of three children are good at identifying emotional facial expressions of happiness both in human and stylized faces, less good in recognizing anger and not at all able to recognize emotional facial expressions of surprise. Moreover it appears clear that, at the age of three, children are able to decode anger expressed in human faces significantly better. This is confirmed by the χ^2 test comparing the child performance for human vs. stylized faces (see Table 4), and showing that anger, when the Yates correction is applied, is similarly recognized in human and stylized faces, whereas happiness is decoded equally well in both types of stimuli.




Fig. 2. Percentage of the 3-year-old recognition accuracy for human vs. stylized emotional faces of happiness, anger, and surprise

Table 4. Statistical assessment of the 3-year-old recognition accuracy in discriminating human vs. stylized emotional faces of happiness, anger, and surprise

χ^2 test Human vs. Stylized Faces				
	χ^2	p-value	Yates χ^2	Yates p-value
Happiness	0,128	0,724	0,07	0,791
Anger	4,72	0,0297	3,066	0,07
Surprise	1,375	0,239	0,022	0,882
	Df=1 χ^2 Critique= 3,84			$\alpha=0,05$

In addition, as it is reported for the six-year olds also for the three-year-olds the χ^2 test comparing (see Table 5) their recognition accuracy in decoding happiness, anger, and surprise expressed by both human and stylized faces, shows that surprise is significantly less better identified than happiness in both human and stylized faces, and in human but not in stylized faces for anger, since at the aged of three anger is less poorly recognized in stylized faces (see Figure 2 also).

Table 5. Statistical assessment of the 3-year-old recognition accuracy in discriminating among happiness, anger, and surprise expressed by both human and stylized faces.

χ^2 test Happiness/Anger/Surprise								
	Human Figures				Stylized figures			
	χ^2	p-value	Yates χ^2	Yates p-value	χ^2	p-value	Yates χ^2	Yates p-value
Happiness/Anger	0	1	0,286	0,592	5,042	0,024	3,227	0,072
Happiness/Surprise	24,889	0	21,46	0	16,667	0	13,5	0
Anger/Surprise	24,889	0	21,46	0	5,042	0,024	3,227	0,072
	Df=1				χ^2 Critique= 3,84			$\alpha=0,05$

5 Conclusions

The present study investigates the ability of typical developing three and six-year-old children to recognize human vs. stylized emotional facial expressions. It was found that at the age of three children are not able to recognize surprise, and for stylized emotional faces their ability to recognize anger is less accurate than in recognizing happiness. In addition, three-year-old children are significantly better at recognizing anger in human rather than in stylized faces. At the age of six children are quite accurate at recognizing emotional expressions of happiness and anger in stylized as well as human faces. This is not the case for their recognition accuracy of surprise which significantly differs according to the type (human versus stylized) but not the gender (males versus females) of the stimulus. In the latter case children prefer stylized faces. It is worth noticing that surprise is also less consistently recognized, and among the proposed emotions, its recognition accuracy is significantly lower.

In general, surprise is a difficult emotion to recognize both for adults and children. In further published experiments [12] surprise obtained a lower recognition accuracy with respect to the other emotion categories, debating in part its role of basic emotion.

Results also show that the emotion recognition accuracy in three and six-year-old children is not affected by the gender (male or female) of both stylized and human faces.

Given the importance of emotional faces in human-human interaction the reported data may support the idea that specialization in recognizing human faces and in particular emotional faces emerges during development through experience (the experience hypothesis) allowing the gradual training of dedicated brain areas that may be devoted to both face and emotional face processing [20, 21]. Being the human face a recurrent focal attentional point since birth, children specialize faster on frequently repeated facial emotional stimuli (such as happiness since birth, and anger while growing, often used in approval vs. prohibition caregiver-child interactions). Facial expressions of surprise may be less experienced by children explaining the three-year-olds' inability to decode such emotional facial stimuli and the significant differences in six-year-olds to recognize them with respect to anger and happiness.

These results may have implications for future developments of intelligent avatars and interactive emotional dialogue systems, since they suggest that at the age of six children's ability to recognize emotional stylized faces (as those simulated by virtual agents) is as good (if not better for some emotion category, like surprise in this case) as their ability to recognize realistic ones.

The fact that stylized expressions of surprise are preferred to the human ones can also be due to the nature of the FACS photos that being actor simulated emotional expressions could have been exasperated by the protagonists and may have produced confusion in the children due to the lack of spontaneity. However, since stylized, avatar and robotic faces have proved to be a simplified and preferential tool for teaching children with Autistic Spectrum Disorder to learn and recognize emotional facial expressions [23], the children's preference may also be attributed to the easiness to decode such more predictable stylized expressions.

The data reported in this study are limited to only three basic emotions. It could be interesting to extend the research to more emotional categories and to a larger sample of subjects of different ages, as well as to exploit realistic stimuli of dynamic emotional facial expressions.

Acknowledgements. This work has been supported by the European projects: COST 2102 “Cross Modal Analysis of Verbal and Nonverbal Communication”, <http://cost2102.cs.stir.ac.uk/> and COST ISCH TD0904 “TMELY: Time in MEntal activitY” (http://w3.cost.eu/index.php?id=233&action_number=TD0904). Acknowledgements go to an unknown reviewer for the useful comments and suggestions and to Miss Tina Marcella Nappi for her editorial help.

References

1. Baron-Cohen, S.: The extreme male brain theory of autism. *Trends Cogn. Sci.* 6, 248–254 (2002)
2. Boni, L.C., Brown, R.T., Davis, P.C., Hsu, L., Hopkins, K.: Social information processing and magnetic resonance imaging in children with sickle cell disease. *Journal of Pediatr. Psychol.* 26, 309–319 (2001), doi: 10.1093/jpepsy/26.5.309
3. Brüne, M.: Emotion recognition, “theory of mind,” and social behavior in schizophrenia. *Psychiatry Res.* 133, 135–147 (2005)
4. Celani, G., Battachi, M.W., Arcidianoco, L.: The understanding of emotional meaning of facial expressions in people with autism. *Journal of Autism and Developmental Disorders* 29(1), 57–66 (1999), doi:10.1023/A:1025970600181
5. Davis, W.: Expression of Emotion. *American Philosophical Quarterly* 25, 279–291 (1998)
6. Denham, S.A.: Emotional development in young children. Guilford Press, New York (1998) ISBN 1572303522
7. de Rosnay, M., Hughes, C.: Conversation and theory of mind: Do children talk their way to socio-cognitive understanding? *Journal of Developmental Psychology* 24(1), 7–37 (2006)
8. Ekman, P., Friesen, W.V., Hager, J.C.: The facial action coding system, 2nd edn. Research Nexus eBook, Weidenfeld & Nicolson, Salt Lake City, London (2002)
9. Ekman, P.: An argument for basic emotions. *Cognition and Emotion* 6, 169–200 (1992)
10. Ekman, P.: The argument and evidence about universals in facial expressions of emotion. In: Wagner, H., Manstead, A. (eds.) *Handbook of Social Psychophysiology*, pp. 143–164. Wiley, Chichester (1989)
11. Ekman, P.: Expression and the nature of emotion. In: Scherer, K., Ekman, P. (eds.) *Approaches to Emotion*, pp. 319–343. Lawrence Erlbaum, Hillsdale (1984)
12. Esposito, A.: The Perceptual and Cognitive Role of Visual and Auditory Channels in Conveying Emotional Information. *Cognitive Computation Journal* 1(2), 268–278 (2009)
13. Herba, C., Phillips, M.: Annotation: Development of facial expression recognition from childhood to adolescence: behavioural and neurological perspectives. *Journal of Child Psychology and Psychiatry* 45(7), 1185–1198 (2004), doi:10.1111/j.1469-7610.2004.00316.x
14. Izard, C.E.: Basic emotions, relations among emotions, and emotion–cognition relations. *Psychological Review* 99, 561–565 (1992)

15. Pons, F., Lawson, J., Harris, P., de Rosnay, M.: Individual differences in children's emotion understanding: Effects of age and language. *Scandinavian Journal of Psychology* 44(4), 347–353 (2003)
16. Pons, F., Harris, P.: *Test of emotion comprehension: TEC*. Oxford University, Oxford (2000)
17. Scherer, K.R.: The role of culture in emotion-antecedent appraisal. *Journal of Personality and Social Psychology* 73, 902–922 (1997)
18. Scherer, K.R., Banse, R., Wallbott, H.G., Goldbeck, T.: Vocal cues in emotion encoding and decoding. *Motivation and Emotion* 15, 123–148 (1991)
19. Scherer, K.R.: Vocal correlates of emotional arousal and affective disturbance. In: Wagner, H., Manstead, A. (eds.) *Handbook of Social Psychophysiology*, pp. 165–197. Wiley, New York (1989)
20. Johnson, M.H.: Subcortical face processing. *Nature Reviews Neuroscience* 6, 766–773 (2005)
21. Turati, C., Macchi Cassia, V., Simion, V., Leo, I.: Newborns' face recognition: role of inner and outer facial features. *Child Development* 77, 297–311 (2006)
22. Grizzle, J.E.: Continuity Correction in the χ^2 Test for 2x2 Tables. *The American Statistician* 21, 28–32 (1967)
23. Adams, A., Robinson, P.: An Android Head for Social-Emotional Intervention for Children with Autism Spectrum Conditions. In: D'Mello, S., Graesser, A., Schuller, B., Martin, J.-C. (eds.) *ACII 2011, Part II. LNCS*, vol. 6975, pp. 183–190. Springer, Heidelberg (2011)

Appendix

1. Come ti sembra questo/a signore/a? How does this woman/man appear to be?
2. Cosa pensi stia facendo? What you think she/he is doing?
3. Cosa pensi stia provando? How do you think she/he is feeling?
4. Cosa provi quando anche tu ti vedi come lei/lui? How do you feel when you are like her/him?
5. Cosa esprime? Can you guess what she/he is expressing?
6. Ti piace questa faccia? Sì/No? Perché? Do you like this face? Yes/No? Why?
7. Pensi stia giocando? Do you think she/he is playing?
8. Pensi stia litigando? Do you think she/he is arguing?
9. Pensi stia piangendo? Do you think she/he is crying?
10. Pensi che abbia visto qualcosa di inaspettato? Do you think she/he has seen something unexpected?

For the experimenter: Has the child decoded the emotional face correctly?

Emotional Status Determination in HCI Interface for the Paralyzed

Rytis Maskeliunas, Vidas Raudonis, and Paulius Lengvenis

Kaunas University of Technology
rytis.maskeliunas@ktu.lt

Abstract. Authors present parts of the work on the development of multimodal HCI control interface targeted at paralyzed people. In this paper the proposed emotional status determination solution is presented based on the data gathering from eye tracking sensors. The participants were provided with an audiovisual stimulus (slides/videos with sounds) and the emotional feedback was determined by the combination of gaze tracking and artificial neural network processing. Some initial experimental analysis and the data of the recognition accuracy of the emotional state based on the gaze tracking are provided along with the description of implemented algorithms.

Keywords: Emotions, HCI, paralyzed, gaze-tracking.

1 Introduction

A human-computer interaction (further HCI) system in principle could be designed for people with severe disabilities unable to move or control any parts of their bodies except for their eyes [1, 2, 3]. The cause of this is often cases of sever paralysis or diseases such as ALS, multiple sclerosis and similar. Usually the intellect and memory of these people remain not affected, sometimes neither are the senses of sight, taste, smell or hearing and touch, meaning that in principle it would be possible to provide the means to make the disability most “endurable”. A lot of aids exist to help the person himself, to ease the work of physiotherapists and to support disabled people’s family. One of the more successful solutions is a head and eye tracking system capable of capturing the movements of the head and the eyes of the disabled person to establish the necessary HCI schemes or determine the user’s needs or status.

Eye-tracking can be especially useful investigating the behavior of the individuals with physical and psychological disorders. Such studies typically focus on the processing of emotional stimuli, suggesting that eye-tracking techniques have the potential to offer an insight into the downstream difficulties of everyday social interaction and integration which such individuals often experience [4, 5, 6]. Studies such as [7] describe how eye analysis can be used to understand human behavior, the relationship between pupil responses, various social attitudes and behaviors, and how it might be useful for diagnostic and therapeutic purposes.

Human – Computer interaction may be significantly improved by incorporating social and emotional processes [8], especially in those targeted at people with special needs. It is clear that vital emotional information might affect the human behavior in the context of information technology. It is considered that this area is somewhat still in its infancy despite the suggested positive effects of recognizing the emotional cues and further using that information in the development of HCIs. The emotional part closely correlates to the eye based HCIs. For example it is known, that the increase in the size of the pupil of the eye have been found to accompany the viewing of emotionally toned or interesting visual stimuli [9]. Emotional stimuli may also be used to attract the user’s attention and then divert to a control scheme of a HCI. For example, eye tracking can be used to record visual fixation in nearly real-time to investigate whether individuals show a positivity effect in their visual attention to emotional information as in [10]. Different viewing patterns can be detected using strongly stimulant pictures [11].

The text bellow presents parts of our work on the development of multimodal HCI control interface targeted at paralyzed people, concentrating on the recognition of emotions. The proposed ANN based emotional status determination solution is presented based on the data gathering from a gaze tracking system, while providing the participants with an audiovisual stimulus (strongly emotionally “stimulant” slides/videos with sounds). Some initial experiments and the data on the recognition accuracy of the emotional state based on the gaze tracking are provided along with the description of implemented processing algorithms.

2 Our Implementation

The motor apparatus of the eye, the so called muscles of the orbit, precisely and rapidly, controls the eye movements for accurate alignment with fovea, because this area deals with sharp vision. When keeping the gaze on the small object, the eye constantly compensates small head movements to keep the object in the area of sharp vision. Moreover, most eye movement is executed without awareness, i.e., there is no voluntary control. A sufficient amount of studies worldwide prove an interrelation between pupil size, pupil motions and a person's cognitive load or stress. The eye movements and the size of eye pupil strongly depend on the various factors of the environment and mental state of person. To recognize in which emotional state is the person, the authors of this article have developed an eye pupil analysis system that is based on application of gaze tracking and processing using artificial neural networks (ANN).

Our gaze tracking algorithm finds the rough pupil center in the iterative manner and executes the logical indexing on the gray level image using certain threshold of grayness value, which is variable (adaptive). The center of eye image is used as starting position for the region. All other positions are defined by locating a pupil center in the last frame. The values of all pixels which are higher than threshold are equalized to one, otherwise to zero. The threshold Θ is increased or reduced from the default gray level value Θ_0 according to certain conditions which are defined by the current measured diameter r of the object of interest. The threshold Θ is increased by step

$\Delta\Theta$, when current diameter of the object is smaller than the possible limits of the pupil size $[R_{\min}, R_{\max}]$ and otherwise Θ is decreased by step $\Delta\Theta$ if these limits are exceed. The threshold value does not change if current measurement r is between limits. The rough pupil center (coordinates C_x and C_y) is computed in the next step. More about our eye tracking algorithm is published in [12]. The schematic algorithm of the proposed eye tracking algorithm is shown in the table 1.

Table 1. The proposed algorithms for eye pupil tracking and registration

Pupil detection based on adaptive gray level threshold	The accurate detection of the pupil center
<p>The rough center of the pupil is extracted from the grayscale image of the eye $I(u, v)$, where $u \in N, v \in M$.</p> <p>while ($flag = 0$) do</p> <p>step 1 <i>Collect candidate points</i></p> <p style="padding-left: 20px;">$count = count + 1$</p> <p style="padding-left: 20px;">$k = 0$;</p> <p>for $u \in N$ do</p> <p style="padding-left: 20px;">for $v \in M$ do</p> <p style="padding-left: 40px;">if $I(u, v) \leq \Theta$ then</p> <p style="padding-left: 60px;">$k = k + 1$;</p> <p style="padding-left: 60px;">$x_k = u, y_k = v$;</p> <p style="padding-left: 40px;">end if</p> <p style="padding-left: 20px;">end for</p> <p>end for</p> <p>$k \in K$;</p> <p>for $i \in K$ do</p> <p style="padding-left: 20px;">for $j \in K$ do</p> <p style="padding-left: 40px;">$d(i, j) = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}$;</p> <p style="padding-left: 20px;">end for</p> <p>end for</p> <p>step 2 <i>Make measurements of the point cloud</i></p> <p style="padding-left: 20px;">$r = \max(\max_{i \in K} (d(i, j)))$;</p> <p style="padding-left: 20px;">$C_x, C_y \leftarrow x(k), y(k)$</p> <p>step 3 <i>Verify the conditions</i></p> <p>if $R_{\min} \leq r \leq R_{\max}$ then</p> <p style="padding-left: 20px;">$flag = 1$;</p> <p>else if $R_{\min} > r$ then</p> <p style="padding-left: 20px;">$\Theta = \Theta + \Delta\Theta$;</p> <p>else if $R_{\max} < r$ then</p> <p style="padding-left: 20px;">$\Theta = \Theta - \Delta\Theta$;</p> <p>end if</p> <p>if $\Theta_{\min} \leq \Theta \leq \Theta_{\max}$ OR $count > count_{limit}$ then</p> <p style="padding-left: 20px;">$flag = 1, \Theta = \Theta_0, r = \infty$;</p> <p>end if</p> <p>end while</p>	<p>The explicit pupil center $C' = (C'_x, C'_y)$ of the point cloud x_k, y_k is computed using nonlinear least squares approach.</p> <p>if $r \neq \infty$ then</p> <p>step 1 <i>Perform data filtration on the candidate points</i></p> <p style="padding-left: 20px;">$d_i \leftarrow \sqrt{(\bar{x} - x_i)^2 + (\bar{y} - y_i)^2}, i \in K$;</p> <p style="padding-left: 20px;">$\bar{d}, \sigma_d \leftarrow d(i), i \in K$;</p> <p style="padding-left: 20px;">$n = 0$;</p> <p>for $i \in K$ do</p> <p style="padding-left: 20px;">if $\bar{d} - 3\sigma_d \leq d(i) \leq \bar{d} + 3\sigma_d$ then</p> <p style="padding-left: 40px;">$n = n + 1$;</p> <p style="padding-left: 40px;">$x'_n = x_i, y'_n = y_i$;</p> <p style="padding-left: 20px;">end if</p> <p>end for</p> <p>step 2 <i>Fit smallest surrounding circle to the filtrated data</i></p> <p style="padding-left: 20px;">$\mathbf{u} = (C'_x, C'_y, r')$, $\mathbf{X} = \begin{pmatrix} x'_1 & \dots & x'_n \\ y'_1 & \dots & y'_n \end{pmatrix}$, $\mathbf{X} \in R^2$</p> <p style="padding-left: 20px;">;</p> <p style="padding-left: 20px;">$f_j(C'_x, C'_y, r') = \ \mathbf{C}' - \mathbf{X}_j\ ^2 - r'^2$;</p> <p style="padding-left: 20px;">$C'_x, C'_y, r' \leftarrow \sum_{j=1}^m f_j(\bar{\mathbf{u}})^2 = \min_u \sum_{j=1}^m f_j(\mathbf{u})^2$;</p> <p>end if</p>

The following symbols are used in the above algorithms: N, M – are columns and rows of pixels of the analyzed eye image (“photo”); $d(i,j)$ is the Euclid distance between the pixel candidates; σ is the dispersion of Euclid distance used for the filtration of data, x and y are the coordinates of pixel candidates, flag – is the “flag” used to stop the algorithm (works together with variable count). Count variable was used to limit the scanning of a photo in various grey scale levels (in total 255, in 8 bit color code table).

Our experimental emotion detection system is illustrated in figure 1. In addition to the proprietary gaze tracking hardware/software the system runs the proprietary real-time emotion analysis toolkit based on an Artificial Neural Networks.

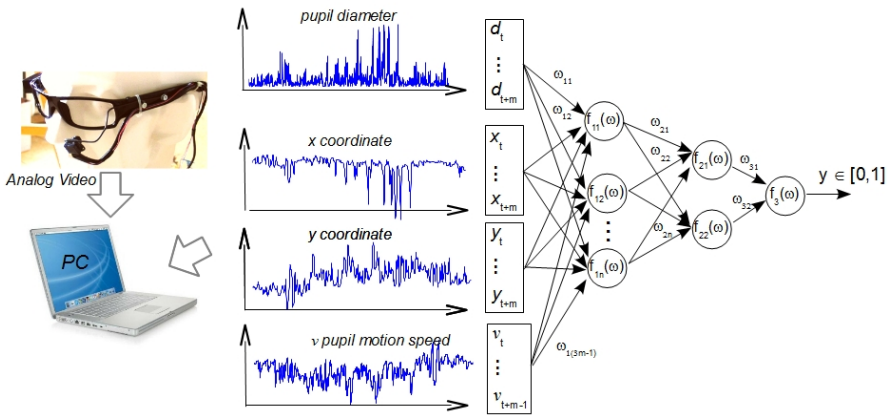


Fig. 1. The ANN model and hardware implementation of experimental setup

We have implemented a 3 layer ANN: the first layer is made from of 8 neurons, the second of 3 neurons and the output layer of 1 neuron. ANN networks have a variable input number and are trained based on 3 features: the size of the pupil, ant the position of the pupil (coordinates x, y) and speed v . The artificial neural network can be described using the following formulas:

$$X = ([d_t \dots d_{t+m}][x_t \dots x_{t+m}][y_t \dots y_{t+m}][v_t \dots v_{t+m-1}]); \tag{1}$$

$$y'_j = f(\sum_{i=1}^{3m} X_i \omega_{ji}); \tag{2}$$

$$y''_k = f(\sum_{j=1}^n y'_j \omega_{kj}); \tag{3}$$

$$y = f(\sum_{k=1}^2 y''_k \omega_2); \tag{4}$$

where t – is the current sample, X – the input vector of the artificial neural network, y – output, ω – weights of the neural network, d – is the diameter of the recognized pupil, x_t and y_t are the coordinates of the pupil’s center and v_t is the speed of eye pupil motion.

3 Experimental Investigation and Results

At this initial stage of the evaluation we have chosen to analyze 4 very common emotions: 1) neutral (regular, typical state); 2) disgust; 3) funny state; and 4) interest (to be interested in something) state. This prototype emotion analysis system was evaluated on 20 people (15 males, 5 females, age ranging from 24 to 42). All participants were presented with a close-up (field-of-view consisted mostly of the display) Power-Point slideshow consisted of various photographs sorted on the type of emotion they were supposed to provoke (the playback time limit was 3 minutes for each of the emotional photo collection (same number of photos for each emotion)). After each set of emotional pictures there were 30 second pauses in the automated slideshow. During the experiment we have registered the size of the pupil, the coordinates of the center of the pupil in the video frame, as well as movement speed and acceleration.

Figure 2 illustrates the experimental analysis (the fragment of 6 people) on the size variation of eye pupil based on a specific emotion (on the left) and size dispersion of eye pupil based on a specific emotion (on the right). The x axis shows the average normalized diameter of an eye pupil. The size of each participant's eye pupil was different during the perception of each emotion, for example the size of the first person's pupil was ~17 % larger when being in neutral state vs being interested (0.075 pixels vs 0.062 pixels). The chart on the right illustrates the variation of eye pupil size over time frame. The deviation data showed that it is important to note that the size can change quiet reasonably during a time period while still experiencing same or similar feelings or emotion.




Fig. 2. The bar graph of relationship between average pupil size and the emotional reaction of the person (left), relationship between standard deviation of pupil size and emotional reaction (right)

The results only confirmed the known fact that the “jump” from one emotion to another can be determined from the physical state and parameters of human's eye pupil (though according to some data a person can be trained to conceal his emotional state – i.e. the fooling of a polygraph effect). We have notice that the actual eye size during each emotional stage was different for all of the 20 people that participated in our experiment. As people reacted differently to different emotional stimuli, and due

to a nature of a human, the visual data we showed could have been perceived differently. For example, relaxing nature, neutral views were very interesting to one person with photographer's background. The size of the eye pupil also varied due to many other factors, such as environmental lighting, stress, physiological state, emotional status before the experiment, etc. (these were somewhat limited in our laboratory testing, but still important to consider).

Overall according to the data we got, we (at the time of writing, using the above system and algorithms) could not universally and subject independently determine what REAL emotion a human was experiencing based only on his eye pupil's size. The parameter had to be added to a group of other companion parameters in the mathematical model of artificial neural network.

Figure 3 illustrates a variation of the center of eye pupil in the two dimensional space (measurement unit - pixels). Overall registration period of the center of eye pupil is divided into four parts based on the shown emotional stimuli. The movement trajectory of eye pupil is shown in normalized pixel units.

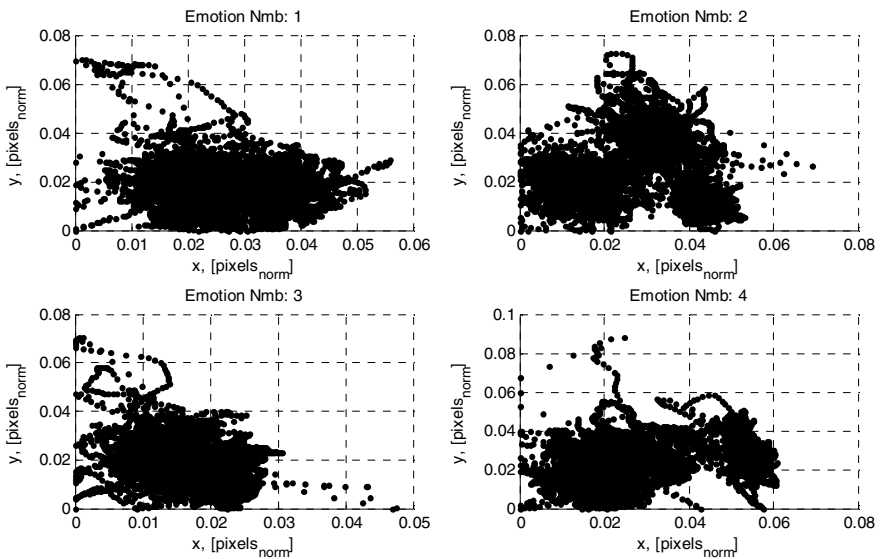


Fig. 3. An illustration of the average variation of the center of eye pupil in a two dimensional space for each emotion analyzed

Analyzing the overall changes of coordinate movements it is clear that the pupil moves quite different during the state of each of the emotion analyzed in the 2D space. The next obvious conclusion was that this data might be usable as one of the features to determine the current emotion. Another important factor noted was that the movement parameter also depended on the context of the emotional information shown on screen as well as the position and motion of the object on screen (especially during videos when the view is concentrated on the subject). For example all of the participants tried not to concentrate on the disgusting pictures, so the average

trajectory covered reasonably large 2D space (figure 3, emotion 2). A similar large overlap of the 2D space was noted when viewing neutral pictures (figure 3, emotion 1) for most of the participants, as most viewed the content (views of nature) quite relaxed and weren't concentrating much. During the viewing of other types of material the movement was somewhat more concentrated. Overall the analysis of movement means that as was in the case of eye pupil's size – this parameter can only be used in combination with other parameters – and in no way only by itself.

The above is also true for the average values of speed and acceleration of eye pupil as shown in figure 4, which illustrates the variation of the typical movement speed of eye pupil depending on the emotional stimuli shown (on the left) and the variation of acceleration of the eye pupil based on the current emotional stimuli (on the right). The speed of the eye pupil's movement was measured in pixels per second (px/s) and the acceleration in pixels per squared second (px/s²). These parameters were also very individual to each of the participants and obviously depended on the initial emotional status before the experiment and the individual subjective interpretation of the data shown. A fragment of 6 people shown clearly illustrates this: the movement speeds were quite consistent and varied only ~10% for the first participant, while the variation was huge (46 %) for the sixth participant. The chart of the acceleration values shows the consistency in the movement of participant's eye pupil. If our test subject perceived "strong" emotions or he was very interested in the information shown during the time-frame on screen, the acceleration increased. For example the acceleration values for the third participant varied more than 50 %.




Fig. 4. The bar graph of relationship between average speed of pupil motion and the emotional stimuli (left), relationship between average acceleration of pupil motion and emotional stimuli (right)

The above initial real-life experimental evaluation confirmed that actual emotions of each of the participants were quite individual so the system must be adapted (trained) to each of the participants in order to reliably recognize the emotional state of the person. In the end to determine the actual effectiveness of the emotional recognizer and to train the ANN we have used the data gathered using 4 sets of emotional PowerPoint slideshows. We have used two for training, and then the remaining two were shown next day for evaluation. Two of the four sets were composed of pictures, the other two of videos, all were the same length.

Figure 5 illustrates the average functional relationship of the number of feature samples and the recognition accuracy of 4 different emotions analyzed. The overall best recognition accuracy performance rate (~90%) was achieved when we used 18 samples per feature. This means that the system can determine the emotion with a 2 second delay with approximately 10 % deviation.




Fig. 5. The functional relationship between recognition accuracy and the number of samples per feature

The overall recognition accuracy for each of the emotion is shown in figure 6. The use of parameters gained by gaze tracking combined with artificial neural network modeling allowed achieving the average subject dependent recognition of 87.35 %. At least in our limited evaluation the joyful emotions were recognized most accurately (90.27 %), and the recognition error for the neutral emotions was the largest (16 %).




Fig. 6. The average recognition accuracy of each of the emotion

4 Conclusions

Our initial real-life experimental evaluation of the gaze tracking data confirmed that the interpretation of actual emotions of each of the participants along with the parameters we measured (the size, the position, the speed, the acceleration of eye pupil) were quite different and dependent on each individual test subject. We think that gaze tracking based emotion recognition system must be adapted (trained) to the individual data of each of the participants in order to reliably recognize the current emotional state of the person.

The implementation of Artificial Neural Networks was used for this purpose. The overall best recognition accuracy performance rate (~90%) was achieved when we used 18 samples per feature. This means that the system can determine the emotion with a 2 second delay with approximately 10 % deviation.

The use of parameters gained by gaze tracking combined with artificial neural network modeling and calculations allowed achieving the average subject dependent recognition of 87.35 %. At least in our limited evaluation the joyful emotions were recognized most accurately (90.27 %), and the recognition error for the neutral emotions was the largest (16 %).

Acknowledgement. Parts of this work were funded by EU SF project "Postdoctoral Fellowship Implementation in Lithuania" (project No.: 20101216-90) within the framework of the Measure for Enhancing Mobility of Scholars and Other Researchers and the Promotion of Student Research (VP1-3.1-ŠMM-01) of the Program of Human Resources Development Action Plan.

References

1. Betke, M., Gips, J., Fleming, P.: The camera mouse: Visual tracking of body features to provide computer access for people with severe disabilities. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 10(1), 1–10 (2002)
2. Miglietta, M.A., Bochicchio, G., Scalea, T.M.: Computer-assisted communication for critically ill patients: a pilot study. *The Journal of TRAUMA Injury, Infection, and Critical Care* 57, 488–493 (2004)
3. Jacob, R.J.K., Karn, K.S.: Eye tracking in human-computer interaction and usability research: Ready to deliver the promises (Section commentary). In: Hyona, J., Radach, R., Deubel, H. (eds.) *The Mind's Eyes: Cognitive and Applied Aspects of Eye Movements*, pp. 573–605 (2003)
4. Luria, S.M., Strauss, M.S.: Comparison of eye movements over faces in photographic positives and negatives. *Perception* 7, 349–358 (1978)
5. Klin, A., Schultz, R., Cohen, D.: Theory of mind in action: developmental perspectives on social neuroscience. In: *Understanding Other Minds: Perspectives from Developmental Neuroscience*, pp. 357–388 (2000)
6. Van Der Geest, J.N., Kemner, C., Verbaten, M.N., Van Engeland, H.: Gaze behavior of children with pervasive developmental disorder toward human faces: a fixation time study. *Journal of Child Psychol. Psychiatry* 43, 669–678 (2002)

7. Hess, E.H.: *The tell-tale eye: How your eyes reveal hidden thoughts and emotions*, 259 p. Van Nostrand Reinhold, Oxford (1975)
8. Kappas, A., Krämer, N.C.: *Face-to-Face Communication over the Internet: Emotions in a Web of Culture, Language, and Technology (Studies in Emotion and Social Interaction)*, 316 p. Cambridge University Press (2011)
9. Hess, E.H., Polt, J.M.: Pupil Size as Related to Interest Value of Visual Stimuli. *Science* 132(3423), 349–350 (1960)
10. Isaacowitz, D.M., et al.: Selective preference in visual fixation away from negative images in old age? An eye-tracking study. *Psychology and Aging* 21(1), 40–48 (2006)
11. Lykins, A.D., Meana, M., Kambe, G.: Detection of Differential Viewing Patterns to Erotic and Non-Erotic Stimuli Using Eye-Tracking Methodology. *Archives of Sexual Behavior* 35(5), 569–575 (2006)
12. Prosevičius, T., et al.: Autoassociative gaze tracking system based on artificial intelligence. *Electronics and Electrical Engineering* 5(101), 67–72 (2010)

Emoticons Signal Expertise in Technical Web Forums

Liliana Mamani Sanchez and Carl Vogel

Computational Linguistics Group, Trinity College Dublin
{mamanisl,vogel}@tcd.ie

Abstract. Past research has demonstrated intercultural differences in emoticon use with effects of the topic of discourse (e.g. science vs. politics) interacting with the culture of online postings (e.g. UK, Italy, Sweden, Germany). The current research focuses within a discourse, and within a lingua franca for communication and attempts to assess whether emoticon use varies as a function of user-type within the online context. The online context is a web user-forum associated with a software technology company. The user categories are determined by a few orthogonal classifications: employees, novice users, and experts; recipients of kudos vs. non-recipients of kudos; etc. As part of a developing theory of presentation of “professional” selves, and perceptions thereof, we test the hypotheses that kudo recipients deploy markedly fewer negative emoticons than comparison categories and that non-employee experts use markedly more emoticons in general than other categories of forum users. Also interactivity across the different group of users and their correlation with emoticon use was explored.

1 Introduction

Community forums are being increasingly used by companies as they provide a channel of communication with consumers. In this paper we explore the forum of a major software company in which participants discuss technical aspects of products and services. Emoticons have been used in sentiment analysis systems as clues for determining sentiment scores ([6]), for training sentiment classifiers with domain independence [7] and for collecting training data for sentiment classification in micro-blogging systems [1]. Cross-cultural analyses of emoticon use are also available [5,8]. Here, we compare emoticons and smilies usage in terms of their relative frequency across three group of forum users. The distinction between emoticons and smilies has emerged over time with the potential for graphical user interfaces to depict them as composed from typographical characters (e.g. “:-)”) or pictorally (e.g. 😊). In an example of synonymy avoidance in natural language [2], the former have become known as “emoticons”, and the

¹ The discussion in the forum analyzed tends to be monolingually English, although, parallel forums exist.

latter as “smilies”; this terminology is adopted in the present work. The forum users share the same language but have different levels of expertise, at least in a nominal dimension (all in a forum may ask or respond to questions that presuppose varying levels of expertise with the subject matter at hand, but are classified in general within the system – a participant may be given the nominal label “guru” without each posting illustrating this status). We are aware that users in forums are likely conscious of how they are perceived, and desire to be deemed as professional and expert as the employees who make contributions. Thus, we anticipate a kind of convergence on linguistic and nonlinguistic features in postings of employees as the amateur contributors participate (cf [4]).

In principle, we distinguish three groups of forum contributors, since they represent distinct aspects of a consumer related forum scenario. Firstly, common users (consumers) approach the forum in search of solutions to technical issues arising from the use of software products [2]. Secondly, gurus are facilitator-users who may have been common users and were given this “badge” name because their degree of contribution to the forum in terms of quality of answers, level of engagement and level of knowledge they share with less expert users. Gurus do not receive any financial reward for their services. The reward for them has more a subjective character for which the main motive for their contribution appears to be the prestige they can obtain, partly via feedback from other users who may reward postings with “kudos” (by clicking on this nominal label in evaluating a posting). Theories from social psychology giving accounts of this kind of dynamic have been broadly described in [3]. Thirdly, employees who are current or past workers: in contrast to the others, they may receive financial reward for their contributions to the forum, i.e. as part of their position duties. Additional details about our method are described in Section 2.

We explore the usage of emoticons and smilies as signals of emotion across users categories. The set of emoticons and smilies will be further described in Section 3 along with the dataset we used. We wanted to know the extent to which emoticon and smilie usage is related to this user classification and whether this usage was influenced by the prestige users obtain and their level of interactivity in the forum. We use the notion of how many kudos (i.e. positive ratings) posts are given as a clue of user prestige. Additional clues for determining level of interactivity we use are: the ratio of the number of posters in a group with relation to the total of posts, and the depth of readership, evidenced by the depth that a post has in a thread. These metrics are detailed in Section 3.2. We present our results in Section 4 and finally we conclude with a discussion about our main findings and future work in Section 5.

2 Method

We decided to split the common users group between ranked and unranked users (aka, “not-ranked”). Rank is assigned to common users who have started to show levels of contribution, but differently from the guru role, they can climb

² We ignore short and long-term lurkers.

a hierarchy of ranks according to certain quantitative thresholds. Gurus on the other hand are assigned this role in a more subjective manner, taking into account quantitative and qualitative parameters. Therefore, the first categorization based in roles includes: {guru, employee, ranked, not-ranked} users. The second categorization is based on whether a user has received kudos: {kudoer user, non-kudoer user}. The third categorization groups posts that received kudos in one category and post that did not receive kudos in another one: {kudoer post, non-kudoer post}. Surveying over the role-based categories allows us to explore emotion expression in groups where the distinction based in levels of expertise is more fine-grained as it is implemented by forum moderators. On the other hand the kudo-based categorizations allows us to work over another dimension of user and post classification, as kudos are given by any kind of user to posts they find useful or outstanding.

We decide to work with emoticons (based on ASCII characters) and smilies (based on pictures) as they both convey emotions and we evaluate our metrics individually or jointly on both types of signals of emotion. The metrics of interactivity we evaluate are related to the volume of posts each user has created. Aspects of volume of posts are the quantity and length. These metrics were inspired by previous research done over use of emoticons based on characters in newsgroups in various languages discussing politics ([58]).

3 Data and Processing

The forum data was obtained from an academic alliance with the R&D department of the software company whose forums we've analyzed. The provided data comprises: actual post content in XML and HTML format (subject and body) and metadata about posts and users. The metadata around posts comprises: posting date, user id, thread id, post id, last edition time, last edition author, kudos received, and views received. Metadata about user comprises: roles and date of promotion to the guru role in case such a role applies. Body of messages comprise the authors' writing, but can also include two elements that are not written by the author: quotes referring to previous posts which are embedded in the text, either in part or entirely, and an edition field.

We used a set of 98 emoticons selected from previous research ([58]). Additionally, we included a set of 45 smilies (e.g. 😊 and 😞) that were provided by the forum management system. Note that the smilies cover the same sorts of expressions of affect as the emoticons (e.g. smiles versus frowns), but also purport to encode additional distinctions as well (such as gender: 😊 vs. 😞).

3.1 Treatment of the Data

The raw forum dataset consisted of 308,274 posts covering all the messages posted between the creation of the forum and 2010-10-12T11:24:16+00:00. We excluded posts and threads where only one or two groups of users could participate, as well as posts from users belonging to the guru group before they

were assigned this role. In related ongoing work, we examine the factors that lead to such individuals ultimately becoming classified as gurus. The working dataset has 208,284 posts. After isolating the text written by a single user, we use regular expressions to count the occurrences of emoticons. Matching true positives proved to be difficult since web text often contains symbols that can be confounded with emoticons resulting in false positives such as the cases shown in Fig. 1. We chose a conservative approach and only matched emoticons well delimited by spaces, tabs or newlines characters.

FAT partition of (C:) and the old (F:) is now (E:)... The storage Destinatioon was a Slave Drive (F:) ...
... 7) Selected Recover My Computer. 8) In the drop-list-box above the list of recovery points, I selected ...
Can add URL or use mask (e.g. with ? or *)

Fig. 1. Cases of false matches for the emoticons :) , 8) and *)

3.2 Forum Data Profile

In this section we explain the metric we used to measure interactivity. Later calculated values will be used to explore a correlation between usage of signals of emotion and level of interactivity. Table 1 shows the number of posts, users and average ratio of posts per individual (APPI) across groups for the three categorisations. The APPI gives a sense of how much interactivity is present in each group: Guru users show the highest level of interactivity, ranked users show the second highest ratio of posts, employees show less interactivity than these two groups, and not-ranked users show the least traffic of posts (almost 1 per user). Kudoer users have a high ratio of posts when compared to non-kudoer users. For the third categorisation, the number of users who authored each kind of post was counted; users from kudoer posts correspond with the number of kudoer users, while the number of users for nonkudoer posts comprises nonkudoer users plus users who have at least one post without kudos. From these figures, it can be seen that 2007 kudoer users have also some nonkudoer posts and 295 users have only kudoer posts. The APPI counts for kudoer and nonkudoer posts are not significantly different, because 2007 of the 2302 kukoer users also contributed posts that received no kudos.

Although illustrative, the APPI is a general measure of interactivity since there is no way to determine whether users have read each other's posts even if the posts belong to the same thread. In some posts, authors included quotes from previous posts, but this is not always the case. However, not including quotes does not mean the poster has not read previous posts since we consider it a default principle every reply to the thread is related to the post that started the thread.³ These caveats understood, we explore two extra measures related to volume of messages per thread: depth of a message and thread creation.

³ Even here, it is possible for users to contribute "off-thread" posts to a thread.

Table 1. Posts per group for the three categorisations

	Role	numposts	numusers	APPI
By role	employee	25,490	400	63.725
	guru	27,489	15	1,832.60
	notranked	50,456	19,462	2.593
	ranked	104,849	2,273	46.128
User Cat.	kudoer	139,164	2,302	60.456
	nonkudoer	69,120	19,848	3.482
Post Cat.	kudoer	18,540	2,302	8.054
	nonkudoer	189,744	21,855	8.682

Table 2. Average depth of postings for the three categorisations

	Role	Mean	sd
By role	employee	12.96	31.11
	guru	9.75	18.48
	notranked	18.49	51.14
	ranked	16.46	34.84
User Cat.	kudoer	14.35	33.15
	nonkudoer	17.94	43.47
Post Cat.	kudoer	15.45	39.26
	nonkudoer	15.34	36.05

The depth of a message in a thread signals levels of interactivity if we follow the logic that a user posting the n -th message in a thread has read the previous $n - 1$ posts. We assigned the first post in a thread a 0-depth, the next 1 and so on. This intuition that a poster of a message at level n has read all the previous $n - 1$ messages may, however, be wrong, especially cases where threads have a high number of posts. Yet, the measure gives us an idea of the latency of threads in terms of interactivity as it shows a user's tendency to participate in extant conversations. Table 2 shows the average and standard deviation of posts' depths, for the three user classifications. The less experienced users: Not-ranked, non-kudoers have the highest average of post's depth, while the depth averages for kudoer posts and non-kudoer posts are not significantly different.

The next measure we explore is whether a post is a new posting (creates a new thread) or a reply (contributes to an existing thread). Posting new messages is less interactive than replying, as a reply-poster responds to at least one previous post, while thread-creators at most invite reply, whether or not their new posts engage with those of separate prior threads. Table 3 shows the number of new posts in contrast to the number of replies and the ratio of new posts to replies (NPRR) across groups.⁴ The numbers show that even posting a high number of replies, non-ranked, non-kudoers and non-kudoer posts' authors write more new posts than the rest of users. Furthermore gurus are less likely to start a thread, but they are highly interactive with 75 replies in average for each new post.

On the basis of these measures, there is every likelihood that the group of gurus is the most interactive one, while for the kudo-categorisations, kudoers and kudoer-post authors are the most interactive ones. Long threads point to high interactivity because they reflect group discussion rather than bipersonal conversations. However is not possible to measure the length of a thread per group in our scenario, since threads are not exclusive to a specific group.

Table 4 shows averages of word-level tokens per post, that could help to find out levels of interactivity if we hypothesise that the most interactive users tend to write more content that could include a high use of emoticons, however this

⁴ In this case, we use a version of the dataset that includes posts made by gurus before their promotion, for sake of integrity in the calculation of NPRR.

Table 3. Ratio of new posts to replies per role

	Role	NewPosts	Replies	NPRR
By role	employee	1,954	23,563	0.08
	guru	637	48,109	0.01
	notranked	16,729	33,730	0.50
	ranked	11,868	93,186	0.13
User Cat.	kudoer	11,990	150,477	0.08
	nonkudoer	19,198	48,111	0.40
Post Cat.	kudoer	1,599	20,254	0.08
	nonkudoer	29,589	178,334	0.17

Table 4. Average of tokens per post

	Roles	Tokens
By role	employee	58.63
	guru	70.70
	notranked	89.23
	ranked	74.03
User Cat.	kudoer	70.24
	nonkudoer	85.75
Post Cat.	kudoer	99.38
	nonkudoer	73.04

intuition is not correct in all the cases since sometimes users write their posts copying/pasting logs from software tools. As we pointed in Sect. 3.1 we tried to identify these elements to reduce them to one token, but it was not possible in every case. Still from these tables we can deduce that not-ranked, non-kudoer users and kudoer-post authors are the most prolific regarding post size.

That is, in this section we have profiled user categories according to proxy measures of interactivity and forum-related expertise. We conjecture that use of signals of affect, whether with emoticons or smilies, is a function of user-interactivity and expertise. We expect the more expert users to lean towards signals of positive affect, and greater levels of affect signalling with greater levels of interactivity. In next section we relate the general findings in this section to frequencies of signals of emotion usage.

4 Results



























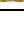
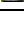
A table of the 10 most frequently used emoticons and smilies in each of the three categories of affect considered is provided in Table 5.

4.1 Usage of Signals of Emotion Across Groups

We explored how frequently signals of emotion are used in relation to the amount of messages posted by each group. The number of posts with and without signals of emotion and the proportion of use of signals are shown in Table 6. In general the use of either emoticon or smilies is low in our dataset ($\sim 6\%$) across the three categorizations, while the combined counts (emoticons plus smilies) show that at most 10% of the posts use at least one signal of emotion when the authors are ranked users. Combined proportions from kudoers and non-kudoers are not very different, although kudoers use signals of emotion more frequently.

Posts from gurus, not-ranked, non-kudoer users and non-kudoer-posts authors include more emoticons than smilies when compared to the ones authored by employees and kudoers. More than double the number of employees' posts containing emoticons contain smilies, kudoers use almost the same amount of

Table 5. Most frequently used signals of affect via emoticons and smilies

Rank	Positive				Negative				Neutral			
	E+	N	S+	N	E-	N	S-	N	E?	N	S?	N
1	:)	2,692		3620	!!!	3,272		1057	()	30		851
2	:~)	1,157		2336	???	2,410		463	\$\$\$	21		11
3	;)	779		1462	: (597		428	(=	18		10
4	:D	588		421	!?!?	252		40	\$\$	13		10
5	:P	223		37	:- (123		25	(D)	11		0
6	;~)	208		31	:/	48		18	<=	5		0
7	=)	186		27	;- (6		18	:>	4		0
8	8)	102		21	>;	3		18	(8x)	1		0
9	:-D	99		18	:-((3		11	:-<	1	-	-
10	=>	98		17	:X	2		5	I	1	-	-

emoticons as smilies. The ratio of usage of emoticons to smilies by gurus and non-kudoer-posts’ authors is 1.12 and 1.25 respectively. In not-ranked and non-kudoers’ posts this ratio raises to 2.6. Another observation is the disjoint use of emoticons and smilies in posts, the combined count of posts using either emoticons or smilies or both almost equates to the addition of the individual counts of emoticons and smilies in all the groups (%C-(%E+%S) ≈ 0).

As non-employees, gurus, ranked and not-ranked users can be less formal than users belonging to the company; it is the case that employees use the forum to make formal announcements where the use of emoticons and smilies would be less obviously appropriate than in more engaging posts. This observation is supported by the fact that employees use less signals of emotion in their posts when compared to the users from the other two role-based groups. Reading Table 6 vertically shows that not-ranked, non-kudoers users and non-kudoer-post authors are the ones who use emoticons with more frequency than their counterparts in all the cases of emoticons; kudoers, ranked users and kudoer-posts’ authors are the ones who more use smilies and combined signals of emotion in their respective categorisations.

We show the average frequency and standard deviation of signals of emotion per post for the three categorizations in Table 7, and token average per post containing at least one signal of emotion in the last two columns.

The ratio of frequency of signals of emotion type {positive(+), negative(-), neutral(?)} to overall frequency of signals of emotion per group is shown in Table 8. More than half of the emoticons used by ranked and not-ranked users were negative, although the difference between ratio of positive and negative emoticons for ranked is not as big as for not-ranked users. There is no significant difference between the use of positive and negative emoticons by ranked users in comparison to other categories ($p = .3039$). Employees and gurus used a significant amount of positive emoticons (more than 80%) compared to negative emoticons. Their use of emoticons is not significantly different ($p = .0887$ for

Table 6. Number of posts with and without (E)moticons, (S)milies and (C)ombined, with percentage of posts containing these signals of emotion for three categorisations

	Categories	E> 0	No E	%E> 0	S> 0	S= 0	%S> 0	C> 0	C= 0	%C> 0
By role	employee	244	25,246	0.96	572	24,918	2.24	810	24,680	3.18
	guru	1,170	26,319	4.26	1,041	26,448	3.79	2,176	25,313	7.92
	notranked	3,214	47,242	6.37	1,196	49,260	2.37	4310	46,146	8.54
	ranked	5,795	99,054	5.53	5,896	98,953	5.62	11,334	93,515	10.81
User Cat.	kudoer	5,993	133,171	4.31	7,018	132,146	5.04	12,670	126,494	9.10
	nonkudoer	4,430	64,690	6.41	1,687	67,433	2.44	5,960	63,160	8.62
Post. Cat.	kudoer	811	17,729	4.37	1,025	17,515	5.53	1,778	16,762	9.59
	nonkudoer	9,612	180,132	5.07	7,680	182,064	4.05	16,852	172,892	8.88

Table 7. Average of signals of emotion per post and tokens per post containing at least one signal of emotion across the three categorisations

	Category	Emoticons		Smilies		Combined		Tokens	
		$\mu(E)$	$\sigma(E)$	$\mu(S)$	$\sigma(S)$	$\mu(C)$	$\sigma(C)$	Mean	sd
By role	employee	1.09	0.351	1.11	0.423	1.112	0.414	78.048	100.481
	guru	1.418	1.076	1.336	0.771	1.402	1	76.622	91.991
	notranked	1.231	0.654	1.144	0.53	1.235	0.673	115.004	143.601
	ranked	1.269	1.619	1.292	1.15	1.321	1.446	92.002	128.743
User Cat.	kudoer	1.273	1.204	1.294	1.094	1.319	1.190	87.88	118.54
	nonkudoer	1.265	1.449	1.143	0.517	1.264	1.303	109.88	145.45
Post. Cat.	kudoer	1.29	0.82	1.25	0.98	1.31	0.98	120.83	145.34
	nonkudoer	1.27	1.35	1.27	1.01	1.30	1.25	92.19	125.93

positive and $p = .1879$ for negative). Neutral emoticons are rarely used in this dataset by any user category. Non-kudoers use of negative emoticons is as much as the double of positive, kudoers in the other hand used more positive emoticons but not in the same magnitude as in the role-categorisation. Kudoer posts show more positive emoticon use than non-kudoer posts.

Positive smilies are more used than negative smilies across all groups, with employees using the greatest quantity of positive smilies and not-ranked and non-kudoers, the least. Also employees usage of negative smilies is not significantly different from guru’s case ($p = .1129$). Although the use of negative smilies is small in each group, not-ranked and non-kudoer users have the highest ratio of negative smilies, more than double than for gurus. Neutral smilies were rarely used in this dataset.

According to combined signals of emotion ratios, the usage of positive signals reaches more than 56% across employees, gurus, ranked, kudoer users and kudoer and non-kudoer posts, being particularly high in the employees and gurus; their use of positive signals is not significantly different($p = .1271$). Negative signals usage is high in not-ranked and non-kudoer users compared to the other groups. For ranked user there is no significant difference in their use of positive and negative emoticons compared to other categories ($p = .306$). If not stated

Table 8. Ratio of signals of emotion type to total of signals per group

	Category	E(+)	E (-)	E(?)	S(+)	S(-)	S(?)	C(+)	C(-)	C(?)
By role	employee	0.808	0.180	0.011	0.913	0.035	0.052	0.882	0.078	0.040
	guru	0.851	0.147	0.002	0.875	0.052	0.073	0.862	0.104	0.034
	notranked	0.316	0.673	0.011	0.618	0.334	0.048	0.394	0.586	0.020
	ranked	0.480	0.512	0.008	0.709	0.202	0.090	0.596	0.354	0.049
User Cat.	kudoer	0.596	0.396	0.008	0.753	0.161	0.087	0.681	0.268	0.051
	nonkudoer	0.332	0.660	0.009	0.625	0.326	0.050	0.407	0.574	0.019
Post Cat.	kudoer	0.611	0.378	0.011	0.806	0.118	0.076	0.719	0.235	0.046
	nonkudoer	0.473	0.519	0.008	0.720	0.199	0.081	0.583	0.377	0.040

otherwise, all the differences reported with relation to Table 8 are significant ($p < 0.05$).

Table 9 shows the ratio of signals of emotion relativized by the number of posts per group. Usage of positive signals of emotion is high in the group of gurus, followed by ranked and kudoer users. Ratios are very small in the group of employees due to their little usage of emoticons (cf. Table 6). Not-ranked and non-kudoer users have the biggest ratio of negative emoticons compared to positive emoticons (almost double). Positive and negative emoticons are used almost in the same proportion by ranked users. Neutral emoticons usage is very marginal when relativized to number of postings.

Positive smilies are mostly used across all the groups, this is particularly high in guru, ranked, kudoer users and kudoer posts. Neutral emoticons usage is marginal. Combined counts also show a high use of positive signals of emotion than of negative signals of emotion by all groups but not-ranked and non-kudoers where negative signals are mostly used. Nonetheless, the usage of negative emoticons by ranked users is higher than in employees and gurus groups.

Table 9. Ratio of signal of emotion type to total of posts per group

	Category	E(+)	E (-)	E(?)	S(+)	S(-)	S(?)	C(+)	C(-)	C(?)
By role	employee	0.008	0.002	0.000	0.023	0.001	0.001	0.031	0.003	0.001
	guru	0.051	0.009	0.000	0.044	0.003	0.004	0.096	0.011	0.004
	notranked	0.025	0.053	0.001	0.017	0.009	0.001	0.042	0.062	0.002
	ranked	0.034	0.036	0.001	0.051	0.015	0.007	0.085	0.051	0.007
User Cat.	kudoer	0.033	0.022	0.000	0.049	0.010	0.006	0.082	0.032	0.006
	nonkudoer	0.027	0.054	0.001	0.017	0.009	0.001	0.044	0.063	0.002
Post Cat.	kudoer	0.034	0.021	0.001	0.056	0.008	0.005	0.090	0.029	0.006
	nonkudoer	0.030	0.033	0.001	0.037	0.010	0.004	0.067	0.044	0.005

4.2 Correlation with Levels of Interactivity

As seen in Sect. 3.2, one metric of interactivity is the average of postings per individual (APPI). Gurus and kudoers are significantly more involved in posting than employees, not-ranked and non-kudoer users. The APPI values for not-ranked and non-kudoer users are very low, the closest less active group (employees) differs from these values by factor of ~ 20 . Recall that more than 80% of the emoticons and smilies used by gurus are positive. The use of negative signals per not-ranked users is always bigger than their counterparts (e.g. 59.7% in comparison to 8.1 and 12.7% for employees and gurus respectively on combined signals of emotion). Also, non-kudoer users show high use of negative signals (58.6%) compared to kudoer users.

There is little correlation between number of positive emoticons and post length ($r = .05$, $p = .00018$). The correlation of number of negative emoticons and post length was positive but very weak: $r = .28$ ($p < .0001$, two-tailed), while the correlation of number of neutral emoticons and post length was weakly positive ($r = .24$) but significant ($p = .02754$), maybe due to small number of posts that use neutral emoticons ($N = 84$). The correlation of number of positive signals is small ($r = .15$, $p < .0001$). There is a weak positive correlation between number of negative signals of emotion and post length ($r = .30$, $p < .0001$) and for neutral emoticons the correlation is marginal ($r = .15$, $p < .0001$).

A horizontal forum is the one with more new posts than replies to old ones, for instance newsgroups discussing politics (cf. 5). The forum we explore here is very vertical (more replies than questions) due to its nature: providing support to customers; non-answered threads does not contribute to the image of the forum as a support service mechanism.

5 Concluding Remarks

This exploration of emoticons and smilies has shown some trends according to expertise-oriented user classifications in a technical community forum. Our hypotheses are confirmed with regards to the positive sentiment shown by kudos-receivers. This is directly related to the interactivity variable, since kudos are given as a reward for valuable contributions in the forum. Differences in the usage of two kinds of signals of emotion were found across the different user groups. Negative emoticons are mostly used by common users that often use the forum to expose their technical issues. The findings from this research will be used to explore automatic identification of users deserving promotion.

We plan to explore additional metadata about users and their posts such as: number of views, kudos given by each user, frequency of posting across time and amount of message editing. Another dimension is to explore some intersections between groups, for instance there may be not-ranked users who receive kudos. Also, we have not explored in depth the ranked and not ranked groups as they are assigned ranks according to a hierarchy with promotion based on merit.

As indicated above, we have excluded from this analysis assessment of which posters will eventually be promoted and the extent to which use of signals of affect provide reliable predictors of ultimate promotion.

Acknowledgements. This research is supported by the Trinity College Research Scholarship Program and Science Foundation Ireland (Grant 07/CE/I1142) as part of the Centre for Next Generation Localisation (www.cngl.ie) at Trinity College Dublin. We thank Mikhail Timofeev for providing us with the forum metadata and information about user categories.

References

1. Bifet, A., Holmes, G., Pfahringer, B.: Moa-tweetreader: real-time analysis in twitter streaming data. In: Elomaa, T., Hollmén, J., Mannila, H. (eds.) DS 2011. LNCS, vol. 6926, pp. 46–60. Springer, Heidelberg (2011), <http://dl.acm.org/citation.cfm?id=2050236.2050243>
2. Carstairs-McCarthy, A.: Synonymy avoidance, phonology and the origin of syntax. In: Hurford, J.R., Studdert-Kennedy, M., Knight, C. (eds.) *Approaches to the Evolution of Language. Social and Cognitive Bases*, pp. 279–296. Cambridge University Press (1998)
3. Cheng, R., Vassileva, J.: User motivation and persuasion strategy for peer-to-peer communities. In: *Proceedings of the 38th Annual Hawaii International Conference on System Sciences, HICSS 2005*, vol. 07, p. 193. IEEE Computer Society, Washington, DC (2005), doi:10.1109/HICSS.2005.653
4. Goffman, E.: *The Presentation of Self in Everyday Life*. Doubleday, New York (1956)
5. Janssen, J., Vogel, C.: Politics makes the Swedish:-) and the Italians:-(. In: *Proceedings of LREC – EMOT 2008, Sentiment Analysis: Emotion, Metaphor, Ontology and Terminology*, pp. 53–61 (2008)
6. Kucuktunc, O., Cambazoglu, B.B., Weber, I., Ferhatosmanoglu, H.: A large-scale sentiment analysis for Yahoo! answers. In: *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining, WSDM 2012*, pp. 633–642. ACM, New York (2012), doi:10.1145/2124295.2124371
7. Read, J.: Using emoticons to reduce dependency in machine learning techniques for sentiment classification. In: *Proceedings of the ACL Student Research Workshop, ACLstudent 2005*, pp. 43–48. Association for Computational Linguistics, Stroudsburg (2005), <http://dl.acm.org/citation.cfm?id=1628960.1628969>
8. Vogel, C., Janssen, J.F.: Emoticonsciousness. In: Esposito, A., Hussain, A., Marinaro, M., Martone, R. (eds.) *COST Action 2102. LNCS*, vol. 5398, pp. 271–287. Springer, Heidelberg (2009)

Machine Learning and Soft Computing Methodologies for Music Emotion Recognition

Angelo Ciaramella¹ and Giuseppe Vettigli²

¹ Dept. of Applied Sciences, University of Naples “Parthenope”,
Isola C4, Centro Direzionale I-80143, Napoli, Italy
angelo.ciaramella@uniparthenope.it

² University of Naples “Parthenope”, Napoli, Italy
giuseppe.vettigli@studenti.uniparthenope.it

Abstract. Social interaction is one of the main channels to access reality and information about people. In this last years there is a growing interest in community websites that combine social interaction with music and entertainment exploration. Music is a language of emotions and music emotional recognition has been addressed by different disciplines (psychology, cognitive science and musicology). Aim of this work is to introduce a framework for music emotion recognition based on machine learning and soft computing techniques. First, musical emotional features are extract from audio songs and successively they are elaborated for classification or clustering. One user can submit a target song, representing his conceptual emotion, and to obtain a playlist of audio songs with similar emotional content. In the case of classification, a playlist is obtained from the songs of the same class. In the other case, the playlist is suggested by system exploiting the content of the audio songs and it could also contain songs of different classes. Several experiments are proposed to show the performance of the developed framework.

Keywords: Music Emotion, Emotion Recognition, Audio Features, Machine Learning, Soft Computing.

1 Introduction

The Multimedia content (pictures, music, videos, e-mails, etc.) is one of the main channels through which we access reality and information about people and their social interaction [15]. To represent the content of multimedia material in terms of the social interactions, they portray means to bring information retrieval systems closer to our social intelligence, with potentially high improvements in terms of retrieval performance. Music is a medium that helps to learn and develop the necessary communication and social skills that are essential for an independent adult life [1]. Emotions, usually, play a critical role in rational decision-making, perception, human interaction, and human intelligence [16]. Music has an extraordinary ability to evoke powerful emotions [4]. This ability

is particularly intriguing because, unlike most other stimuli that evoke emotion, such as smell, taste or facial expression, music has no obvious intrinsic biological or survival value. Although, changes in certain physiological processes have been characterized in response to music neural correlates of emotional responses to music, their relation to music perception and their relation to other forms of emotion have not been well studied.

Many issues for music emotion recognition have been addressed by different disciplines such as physiology, psychology, cognitive science and musicology [12]. In many cases, the terms emotion and mood have been used interchangeably. It is noted that, in most psychology related books and papers emotion usually has a short duration (seconds to minutes) while mood has a longer duration (hours or days).

For these reasons we observe that the issue of recognizing music emotions based on acoustic music signal features to build an intelligent system for music emotion recognition is challenging. In [16] the authors present a hierarchical framework to automate the task of mood detection from acoustic music data, by following some music psychological theories in western cultures. The hierarchical framework has the advantage of emphasizing the most suitable features in different detection tasks. Three feature sets, including intensity, timbre, and rhythm are extracted to represent the characteristics of a music clip.

In [23] the authors propose and compare two fuzzy classifiers to determine how likely a song segment belongs to an emotion class. The music features are analyzed considering an *Arousal and Valence* (AV) scheme. While, in [12] the authors focus on a music emotion recognition system based on fuzzy inferences. In particular music features are analyzed and mapped into AV values by a fuzzy inference engine.

Recently, several community websites that combine social interaction with music and entertainment exploration have been proposed. For example *Stereo-mood* [18] is a free emotional internet radio that suggests the music that best suits mood and daily activities of an user. It allows the users to create playlists for different occasions and to share emotions through a manual tagging process.

In this paper we describe a system for music emotion recognition based on *machine learning* and *soft computing* techniques. In our system the user formulates a query providing a target audio song with similar emotions to the ones he wishes to retrieve. The process is realized by using supervised techniques on labeled data or unsupervised techniques on unlabeled data. In our experiments we consider a dataset composed by songs belonging to 4 emotion classes. The emotional classes are a subset of the Thayer model showed in Figure 1. In the first case, the supervised techniques permit to automatically tag new songs using the labels found in the dataset and the playlist is obtained by using a fuzzy similarity measure. In the other case, a clustering approach (fuzzy based) is used to automatically label the data. The vectors provided by the clustering techniques are used to assign a target audio song to one of the clusters. We introduce an appropriate technique, based on a combination of the k-nearest training samples, to build the membership functions. In this case the emotional information

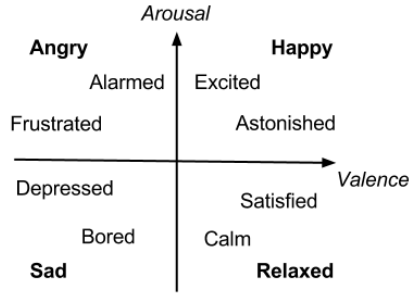


Fig. 1. Two-dimensional emotion representation in Thayers model

is suggested by the system and the approach can be considered a real alternative to human based tagging systems (i.e. stereomood).

The paper is organized as follows. In Section 2 the music emotional features are described. In Section 3 the overall system and the used techniques are described. In Section 4 and Section 5 several experimental results and considerations are proposed, respectively.

2 Emotional Factors from Music

To classify music emotions 5 musical features are considered. In particular, the features extracted from the acoustic music signals are: *intensity*, *rhythm*, *key*, *harmony* and *spectral centroid*. In detail, we have that intensity is represented by loudness and its regularity, rhythm by tempo and standard deviation of beat durations, scale by key, mode and tonality and harmony is represented by the harmonic distribution [16,19,12,23]. In the following we describe the main aspects of these techniques.

2.1 Intensity

The intensity of the sound sensation is related to the amplitude of sound waves [10]. Usually, low intensity is associated to *sadness*, *melancholy* but also *tenderness* or *peacefulness*. High intensity is associated to positive emotions like *joy*, *excitement* or *triumph*. While very high intensity with many variation during the time could be associate to *anger* or *fear*.

To characterize the intensity of the sound we consider the regularity of the volume in the song. First, we compute the mean energy of the signal as follows

$$AE(x) = \frac{1}{N} \sum_{t=0}^N |x(t)|^2, \quad (1)$$

where $x(t)$ is the value of the amplitude at time t and N is the length of signal. And then the standard deviation of AE according to the following equation

$$\sigma(AE(x)) = \sqrt{\frac{1}{N} \sum_{t=0}^N (AE(x) - x(t))^2}. \quad (2)$$

The value of $\sigma(AE(x))$ reflects the regularity of the volume in the song. That is, if the volume is high AE also has a high volume. Moreover, equation 2 computes the regularity of loudness. The more frequently the loudness appears, the higher this value will become.

2.2 Rhythm

To describe the rhythm of a song, an analysis of *beat* and *tempo* is considered. The beat is the regularly occurring pattern of rhythmic stresses in music. When we count, tap or clap along with music we are experiencing the beat. Tempo is the speed of the beat, usually expressed in *Beats Per Minute* (BPM). Fast music causes blood pressure, heart and breathing rate to go up, while slow music causes these to drop. Regular beats makes listeners *peaceful* or even *melancholic*, but irregular beats could make some listeners feel *aggressive* or *unsteady*. In our system, we used the algorithm proposed in [8] to track the beats. The beat tracking algorithm returns an estimation of the beat locations. For each location l_i the related tempo is computed as $\tau_i = 60/(l_i * 0.0116)$ where 11.6ms is the resolution of the onset detection function applied in the algorithm. The rhythm is considered as the standard deviation of the values $\tau_{i+1} - \tau_i$. The more the beats are irregular, the higher this values will become. [7].

2.3 Key

A scale is a group of pitches (scale degrees) arranged in ascending order. These pitches span an octave. Scale patterns can be duplicated at any pitch. Rewriting the same scale pattern at a different pitch is called transposition. And a key is the major or minor scale around which a piece of music revolves. To characterize the scale in our system the Key Detection algorithm proposed in [14] is used. The algorithm returns the estimated key for each key change. We consider the key of the song as the key associated with the maximum duration in the song.

2.4 Harmony and Spectral Centroid

Harmonics can be observed perceptually when harmonic musical instruments are performed in a song. Harmony refers to the way chords are constructed and how they follow each other. To characterize the harmony we use an analysis of the overtones. The analysis consists of the evaluation of the following function

$$HS(f) = \sum_{k=1}^M \min(\|X(f)\|^2, \|X(kf)\|^2), \quad (3)$$

where, f is the frequency, X is the Short Time Fourier Transform of the source signal, and M denotes the maximum number of frequencies for which the mean of $X(f)$ is higher than a value θ [1] (only those frequencies are used in the computation). At the end, the standard deviation of $HS(f)$ is obtained.

The analysis of the harmony above does not consider the fundamental pitch of the signal, hence, we also consider the *spectral centroid* [11].

3 System Architecture

The prototype of the emotion recognition system is described in Figure 2. It has a perspective of social interaction and for this reason it has been designed for the Web.

The purpose of the system is to provide a framework to retrieve audio songs in a database using emotional information. It is possible to have two cases. On one hand, the songs into the database are emotional labeled by the users. On the other hand, there is no information about the emotion information of the songs. In the first case, a supervised approach is applied and a classifier is trained. In the second case an unsupervised approach is applied and the emotional information is automatically extracted by clustering the songs.

The query engine takes a target audio song as input and returns a playlist of similar songs. However the query engine has a different behavior for the two cases. In the first, the classifier is used to identify the class of the target song and the results are the most similar songs in the same class. To identify the most similar songs a fuzzy similarity measure based on the Łukasiewicz product, is used [20,21,15]. In the second, the membership of each song, computed by the clustering algorithm, are compared to select the results. Thereafter, an appropriate algorithm is proposed to compute the fuzzy memberships of a target audio song.

We considered three techniques to classify the song in the supervised case: Multi-Layer Perceptron (MLP), Support Vector Machine (SVM) and Bayesian Network (BN) [3,9]. While, we considered Fuzzy C-Means (FCM) and Rough Fuzzy C-Means (RFCM) for the clustering task [2,22]. Following we only describe the techniques and the methodologies adopted in the clustering process to clarify the assignment process of the fuzzy memberships.

3.1 Fuzzy and Rough Fuzzy C-Means

The *Fuzzy C-Means* (FCM) is a fuzzification of the C-Means algorithm, and it was proposed by Bezdek [2]. It partitions a set of N patterns $\{\mathbf{x}_k\}$ into c clusters by minimizing the objective function

$$J_{\text{FCM}} = \sum_{k=1}^N \sum_{i=1}^c (\mu_{ik})^m \|\mathbf{x}_k - \mathbf{v}_i\|^2 \quad (4)$$

¹ In the experiment we used $\theta = 10 \times 10^{-3}$.

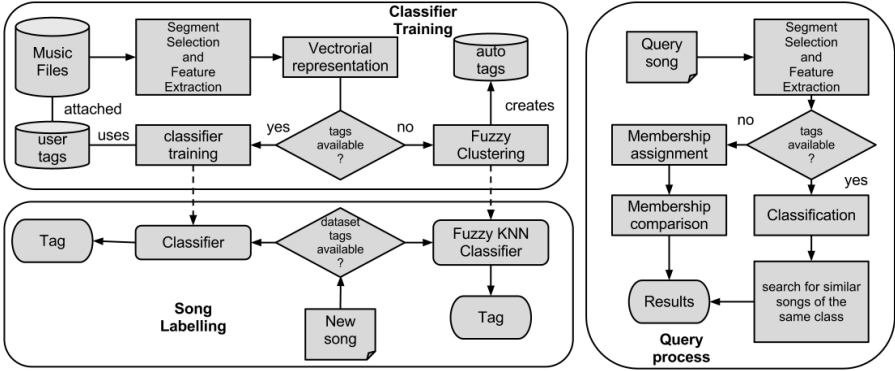


Fig. 2. System architecture

where $1 \leq m < \infty$ is the fuzzifier, \mathbf{v}_i is the i -th cluster center, $\mu_{ik} \in [0, 1]$ is the membership of the k -th pattern to it, and $\|\cdot\|$ is a distance between the patterns, such that

$$\mathbf{v}_i = \frac{\sum_{k=1}^N (\mu_{ik})^m \mathbf{x}_k}{\sum_{k=1}^N (\mu_{ik})^m} \quad (5)$$

and

$$\mu_{ik} = \frac{1}{\sum_{j=1}^c \left(\frac{d_{ik}}{d_{jk}}\right)^{\frac{2}{m-1}}} \quad (6)$$

with $d_{ik} = \|\mathbf{x}_k - \mathbf{v}_i\|^2$, subject to $\sum_{i=1}^c \mu_{ik} = 1, \forall k$. The algorithm to calculate these quantities proceeds iteratively [2].

Based on the lower and upper approximations of *rough set*, the *Rough Fuzzy C-Means* (RFCM) clustering algorithm makes the distribution of membership function become more reasonable [22]. Moreover, the time complexity of the RFCM clustering algorithm is lower by compared with the traditional FCM clustering algorithm. Let $X = \{x_1, x_2, \dots, x_n\}$ be a set of objects to be classified, the i -th class be denoted by w_i , its centroid be v_i , and the number of class be k . Define $\underline{R}w_i = \{x_j | x_j \in w_i\}$, $\overline{R}w_i = \{x_j | \|x_j - v_i\| \leq A_i, A_i > 0\}$, we have

1. if $x_j \in \underline{R}w_i$, then $\forall l \in \{1, \dots, k\}, l \neq j, x_j \in \overline{R}w_l, x_j \in \overline{R}w_l$
2. if $x_j \in \overline{R}w_i$, then at least exist $l \in \{1, \dots, k\}$, make $x_j \in \underline{R}w_l$

Provided that A_i is called the upper approximate limit, which characterizes the border of all possible objects possibly belonged to the i -th class. If some objects do not belong to the range which is defined by the upper approximate limit, then it belongs to negative domain of this class, namely, it does not belong to this class. The objective function of RFCM clustering algorithm is:

$$J_{\text{RFCM}} = \sum_{k=1}^N \sum_{i=1, \mathbf{x}_k \in \overline{R}w_i}^c (\mu_{ik})^m \|\mathbf{x}_k - \mathbf{v}_i\|^2 \quad (7)$$

where the constraints are $0 \leq \sum_{j=1}^n \mu_{ij} \leq N$, $\sum_{i=1, \mathbf{x}_k \in \bar{R}w_i}^c \mu_{ik} = 1$. We can also get the membership formula of RFCM algorithm as follows

$$\mathbf{v}_i = \frac{\sum_{k=1}^N (\mu_{ik})^m \mathbf{x}_k}{\sum_{k=1}^N (\mu_{ik})^m} \tag{8}$$

and

$$\mu_{ik} = \frac{1}{\sum_{l=1, \mathbf{x}_k \in \bar{R}w_i}^c \left(\frac{d_{ik}}{d_{lk}}\right)^{\frac{2}{m-1}}} \tag{9}$$

Also in this case the algorithm proceeds iteratively.

3.2 Fuzzy Memberships

After the FCM (or FRCM) process, we have that the i -th object in the c class has a membership μ_{ic} . In fuzzy classification, we assign a fuzzy membership μ_{uc} for a target input x_u to each class c (on C total classes) as a linear combination of the fuzzy vectors of k -nearest training samples:

$$\mu_{uc} = \frac{\sum_{i=1}^k w_i \mu_{ic}}{\sum_{i=1}^k w_i} \tag{10}$$

where μ_{ic} is the fuzzy membership of a training sample x_i in class c , x_i is one of the k -nearest samples, and w_i is the weight inversely proportional to the distance d_{iu} between x_i and x_u , $w_i = d_{iu}^{-2}$. With equation 10 we get the $C \times 1$ fuzzy vector μ_u indicating music emotion strength of the input sample: $\mu_u = \{\mu_{u1}, \dots, \mu_{uC}\}$ such that $\sum_{c=1}^C \mu_{uc} = 1$. The corresponding class is obtained considering the maximum of μ_u .

4 Experimental Results

In this section we report some experimental results. We consider, for simplicity, a dataset of 100 audio songs belonging to 4 different classes: *Angry*, *Happy*, *Relax*, *Sad*.

For the classification phase we compared 3 machine learning approaches: MLP, SVM and BN. For this kind of data we obtain, with a 10-fold cross-validation methodology [3], the results visualized in Table II. We observe that the methodologies are comparable. In the case of clustering by using FCM and RFCM approaches, we obtain, making a mean on 100 iterations, 61.36% and 66.14% ($A = 0.5$) of perfect classification, respectively. For each iteration, the label of the class is assigned by voting the samples belonging to it and a song is considered perfect classified if it is assigned to the right class. These results are not comparable with those obtained with the supervised techniques. However we highlight that in the first case the playlist is obtained from songs contained in the same class. In the second case the emotional information is suggested by the system.

For example, we propose a query that has as target the “*La domenica delle salme*” song of “*De Andrè*”, classified as Sad, and the results of the first 4 similar songs are:

- “*De Andrè* - *Il suonatore Jones*” (Sad)
- “*Yann Tiersen* - *Comptine d’un autre été*” (Relax)
- “*BJork* - *Vespertine*” (Relax)
- “*De Andrè* - *Un blasfemo*” (Sad).

We can consider that these songs are musically similar but they not belong to the same class.

In a new experiment we used the “*Musclemuseum*” song of “*Muse*” labeled as Angry. In this case the 4 similar songs are:

- “*Helloween* - *My Life for One More Day*” (Angry)
- “*Negramaro* - *Mentre tutto scorre*” (Angry)
- “*Muse* - *Space Dementia*” (Angry)
- “*Muse* - *Hysteria*” (Angry).

In this case we obtained 4 songs of the same class but they are not of the same author and they are of different genre.

In the last experiment we query the song “*L’angioletto in blue jeans*” of the “*Zecchino d’oro*” classified as Happy. The 4 similar songs are:

- “*Zecchino d’oro* - *La sveglia biricchina*” (Happy)
- “*Per Cortese* - *Non capirò mai*” (Happy)
- “*Yann Tiersen* - *La Valse Des Vieux Os*” (Relax)
- “*Zecchino d’oro* - *Il topo zorro*” (Happy).

Also in this last case we obtain songs of different authors and classes but we find that they are similar from emotional point of view.

Table 1. Results for 10-fold cross-validation with three different supervised machine learning approaches considered for the automatic song labeling task (TP, True Positive; FP, False Positive)

Classifier	TP Rate	FP Rate	Precision	Recall
Bayes	65%	11%	66%	65%
SVM	72%	9%	73%	72%
MLP	70%	9 %	70%	70%

5 Conclusions

In this work we introduced a music emotion recognition system based on machine learning and soft computing techniques. The considered musical features were intensity, rhythm, scale, harmony and spectral centroid.

The purpose of the system is to provide a tool to search songs in a database using emotional information. The query engine takes in input a target audio song provided by the user and returns a playlist of most similar songs. From one hand a classifier is used to identify the class of the target song and the results are the most similar songs in the same class. This is obtained by using a fuzzy similarity measure based on the Lukasiewicz product. On the other hand, a clustering algorithm is used to automatically label the data. In this case an appropriate algorithm is proposed to compute the fuzzy memberships of the new songs.

The obtained results with clustering are not comparable with those obtained with the supervised techniques. We, however, stress that in the first case the playlist is obtained by songs contained in the same class and in the second case the emotional information is suggested by the system. The approach can be considered a real alternative to human based classifier systems (i.e. stereomood).

In the next future the authors will focus the attention on a greater database of songs, further musical features and the use of semi-supervised approaches. Moreover they will experiment new approaches as the Fuzzy Relational Neural Network [6], that permits to extract automatically memberships and IF-THEN reasoning rules.

References

1. Barrow-Moore, J.L.: The Effects of Music Therapy on the Social Behavior of Children with Autism, Master of Arts in Education College of Education California State University San Marcos (November 2007)
2. Bezdek, J.C.: Pattern Recognition with Fuzzy Objective Function Algorithms. Plenum Press, New York (1981)
3. Bishop, C.M.: Pattern Recognition and Machine Learning. Springer (2006)
4. Blood, A.J., Zatorre, R.J., Bermudez, P., Evans, A.C.: Emotional responses to pleasant and unpleasant music correlate with activity in paralimbic brain regions. *Nature Neuroscience* 2, 382–387 (1999)
5. Ciaramella, A., Riccio, A., Galmarini, S., Giunta, G., Potempski, S.: Comparison of Dispersion Models by Using Fuzzy Similarity Relations. In: Pirrone, R., Sorbello, F. (eds.) *AI*IA 2011. LNCS*, vol. 6934, pp. 57–67. Springer, Heidelberg (2011)
6. Ciaramella, A., Tagliaferri, R., Pedrycz, W., Di Nola, A.: Fuzzy Relational Neural Network. *International Journal of Approximate Reasoning* 41, 146–163 (2006)
7. Bello, J.P., Daudet, L., Abdallah, S., Duxbury, C., Davies, M., Sandler, M.B.: A Tutorial on Onset Detection in Music Signals. *IEEE Transactions on Speech and Audio Processing*, 1035–1047 (2005)
8. Davies, M.E.P., Plumbley, M.D.: Context-dependent beat tracking of musical audio. *IEEE Transactions on Audio, Speech and Language Processing* 15(3), 1009–1020 (2007)

9. Duda, R.O., Hart, P.E., Stork, D.G.: Pattern Classification. Wiley-Interscience (2000)
10. Revesz, G.: Introduction to the psychology of music. Courier Dover Publications (2001)
11. Grey, J.M., Gordon, J.W.: Perceptual effects of spectral modifications on musical timbres. *Journal of the Acoustical Society of America* 63(5), 1493–1500 (1978)
12. Jun, S., Rho, S., Han, B.-J., Hwang, E.: A Fuzzy Inference-based Music Emotion Recognition System. In: 5th International Conference on Visual Information Engineering, VIE 2008 (2008)
13. Keller, J.M., Gray, M.R., Givens, J.A.: A Fuzzy k-Nearest Neighbor Algorithm. *IEEE Transaction on Systems, Man and Cybernetics* 15(4), 580–585 (1985)
14. Noland, K., Sandler, M.: Signal Processing Parameters for Tonality Estimation. In: Proceedings of Audio Engineering Society 122nd Convention, Vienna (2007)
15. Vinciarelli, A., Pantic, M., Heylen, D., Pelachaud, C., Poggi, I., D’Errico, F., Schroeder, M.: Bridging the Gap Between Social Animal and Unsocial Machine: A Survey of Social Signal Processing. *IEEE Transactions on Affective Computing* (2011)
16. Lu, L., Liu, D., Zhang, H.-J.: Automatic Mood Detection and Tracking of Music Audio Signals. *IEEE Transaction on Audiom Speech, and Language Processing* 14(1) (2006)
17. Pauws, S.: Musical key extraction from audio. In: Fifth Int. Conf. on Music Inf. Retr. Proc., paper 142 (2004)
18. Stereomood website, <http://www.stereomood.com>
19. Thayer, R.E.: The Biopsichology of Mood and Arousal. Oxford University Press, New York (1989)
20. Turunen, E.: Mathematics Behind Fuzzy Logic. *Advances in Soft Computing*. Springer (1999)
21. Sessa, S., Tagliaferri, R., Longo, G., Ciaramella, A., Staiano, A.: Fuzzy Similarities in Stars/Galaxies Classification. In: Proceedings of IEEE International Conference on Systems, Man and Cybernetics, pp. 494–496 (2003)
22. Wang, D., Wu, M.D.: Rough fuzzy c-means clustering algorithm and its application to image. *Journal of National University of Defense Technology* 29(2), 76–80 (2007)
23. Yang, Y.-H., Liu, C.-C., Chen, H.H.: Music Emotion Classification: A Fuzzy Approach. In: Proceedings of ACM Multimedia 2006, pp. 81–84 (2006)

Homo-Machina Visual Metaphors, Representations of Consciousness and Scientific Thinking

Mauro Maldonato and Ilaria Anzoise

Università degli Studi della Basilicata, Potenza, Italy
mauro.maldonato@unibas.it

Abstract. Since the ancient past, philosophers and scientists have developed a particular kind of representations aiming at describing the human nature. Many of these images could be considered as visual metaphors, so it is possible to study them using two different approaches. The first one aims at analysing the connection existing between these images, the scientific paradigms, and other elements of the cultural context within which they emerged. The second approach strives to underline some specific elements that characterise the inherent power of these visual metaphors, ensuring their success in scientific and cultural context.

Keywords: visual metaphors, consciousness, machine metaphor, computer metaphor.

1 Visual Metaphors

Since the ancient past, philosophers and scientists, in different cultural contexts, have developed a particular kind of representations aiming at describing the human nature, in particular the relationship between mind and body, and the idea of consciousness. Many of these images could be considered as visual metaphors. The linguistic research stated that, since language is a sophisticated cognitive instrument through which human beings both communicate and elaborate the world around them, metaphors are representations that they produce in order to understand their inner and outer reality.

In this way if the visual language, as production of images, could be considered a cognitive strategy, metaphors could be considered cognitive structures via which the less clearly delineated (and usually less concrete) concepts are partially understood in terms of the more clearly delineated (and usually more concrete) concepts, which are directly grounded in our experience [1] Since our ordinary conceptual system, in terms of which we both think and act, is fundamentally metaphorical in nature [2] the conceptual metaphors are cognitive structures via which we can understand and shape our whole reality. So, the use of conceptual metaphors ends up modeling all our world, the physical and the psychological one (see Fig. 1, Fig. 2 and Fig. 3). In other words representation reveals how the human brain carries out its work of transforming sensory knowing into conceptual knowing so that since concepts are mental forms, it follows that the form that knowledge assumes depends on the type of modeling used [3].




Fig. 1. Fritz Kahn, *Das Leben des Menschen*, 1929. The machine as a visual and a conceptual metaphor of the human body.



Fig. 2. Giuseppe Arcimboldo, *Rudolph II, Holy Roman Emperor, painted as the God of the seasons Vertumnus*, 1590. The nature as a visual and a conceptual metaphor of the human body.



Fig. 3. Giovan Battista Della Porta, *De Humana Physiognomonia libri IIII*, 1586. The animal as a metaphor for the human nature.

Since the representations of human nature, developed within the philosophical and scientific fields, are devised to understand both physical and abstract experiences, as the interaction between mind and body is both physical and abstract, it is possible to study these visual metaphors using two different approaches. The first one, based on the historical perspective, aims at analysing the connection existing between these

images, the scientific paradigms, and other elements of the cultural context within which they emerged. The second approach strives to underline, from the perspective of visual semiotic, some specific elements that characterise the inherent power of these visual metaphors, ensuring their success in scientific and cultural context.

Following the first approach, we will analyse the relationship between two particular kinds of representations: the machine-metaphor and its relationship with the scientific revolution of the 16th century, and the computer-metaphor related to the cybernetic revolution of the 20th century. Adopting the second approach we will analyse some data from a pilot study, based on the Jungian theory of the archetypal representations, on visual metaphors of consciousness. The study was intended to identify some specific iconographic elements that play an important role both in cognitive and in self-representation processes.

Since a metaphor could be considered both a cultural product and a cognitive structure, it is necessary to study them following both approaches. Analysing these visual metaphors, by means of the two approaches, we will have the opportunity to highlight the role that both cultural and psychological elements play in collective representations.

2 Machine Metaphor

The machine metaphor was certainly not an invention of 17th century's scholarship. The Greek scientific thinking of the 3rd century BC already used, frequently, mechanical metaphors to describe human anatomy and physiology. Two anatomists, Herophilos (335-280 BC) and Erasistratus (304-250 BC), applied a lot of scientific terms borrowed from coeval Greek technology, especially from the fields of architecture and ship construction, describing the whole body as a mechanism activated by the most important principles of Alexandrian technology: torsion-spring artillery and pneumatic devices. It is very interesting to notice that there is a closer relationship between the machine metaphor and the progress of anatomical knowledge in those historical periods when anatomical studies had a central role in medical knowledge, namely in the Alexandrian Era and in 17th century. This suggests that the figurative requirement developed from the anatomical approach to the human body finds a useful means of expression in the machine metaphor¹.

The scientific revolution ideally begun in 1543 with Copernicus's *De revolutionibus orbium coelestium* and Vesalius's *De humani corporis fabrica*, and ended with the publication of Newton's *Philosophiae naturalis principia mathematica* in 1687. Philosophers and scientists, in about 150 years, completed the elaboration of classical physics and mechanics producing the premises of a mechanical description of Nature and of Man. However, the machine metaphor began to look attractive to collective imagination well before the emergence of quantitative science. Indeed,

¹ Indeed the figurative needs encouraged the encounter, in the 16th century, between artists and anatomists. Lucio Russo stated that, in that period, the anatomical knowledge's improvement was simplified by three strategic elements: the retrieval of ancient texts, the transformation of anatomical practice, and the introduction of illustration [4].

during the second period of the 16th century the machine metaphor was used to describe many aspects of human nature and life. For example, the relationship between God and the world was depicted as the relationship between a Great Clockmaker and his Clock, while the absolute monarchy, as the relationship between the King, the State, and the subjects, was illustrated as a Great Clock in which all the parts collaborate to the general functioning. In that period the production of automata grew significantly thanks to technological advances and the recovery of the ancient treatises on mechanics.



Fig. 4. Automaton figure of a monk, probably attributable to Giovanni Torriano. South Germany or Spain, c. 1550. This automaton walk, strikes his chest with the right arm, rises and lowers a cross and a rosary with the left arm and also turn his head rolling the eyes and moving his mouth.

It was also strongly encouraged by the sovereigns of Europe, as they saw in the automata the perfect metaphor for the subject. It is well known, in fact, that the emperor Charles V (1516-1556), in the last years of his life loved to spend time contemplating the automatons constructed by his clockmaker Giovanni Torriano (1450-1510?). The automatons, which attracted the curiosity of courts all over Europe, were not only a perfect metaphor for the subject or for the government of the State: through these mechanical puppets men began to look at themselves as in a mirror. An example of this kind of thinking is a treatise written by Bernardino Baldi (1553-1617), an Italian humanist who translated Heron's treatise on the automata². Baldi stated that the automatons are not only an ingenious toy, useful to understand the principles of mathematics and mechanics. He argued that humans were spending a lot of time, motionless, watching these simulacra of life moving, because the automatons were a kind of mechanical mirror of human nature. Meaning that watching the mechanical puppets, men were meditating on themselves (see Fig. 4). So, long before the scientific revolution was accomplished, the automata became the

² Baldi, B.: *Discorso di chi traduce sopra le macchine semoventi*. In: Erone Alessandrino, *De Gli Automati, ovvero Macchine se moventi*. pp. 4-13. Girolamo Porro, Venezia (1589).

mechanical self-portrait of the man. It seems that the collective imagination developed, through the figure of the automaton, the new figurative needs of human reality, well before science had furnished the complete theoretical basis.

A decisive step towards the development of the human-machine metaphor was made in 1633 when Descartes finished writing the treatise *Le Monde* of which *L'Homme*, the most carefully written and sophisticated part of Cartesian biology, was the conclusion. The news of Galileo's sentence persuaded the French philosopher to give up the publication of *L'Homme* that was published posthumous, in 1662 in Latin and in French, first in 1664 and then in 1677. Through this treatise, Descartes realized what Carlo Sini has called the first mental experiment in the history of modern science: the experiment of the automaton, aimed at explaining the human body as if it were a machine. According to Sini this is the very first formulation of Descartes's great intuition, an intuition to which Lamettrie in his *Homme machine* reverted one hundred years later. Descartes's main aim was to free the emerging modern science, and the research on human nature, from the restrictions imposed until then by theology and the whole apparatus of superstitions that Christianity inherited from the classical world and further developed. With such an image, originated in order to provide an interpretation of the human body inspired exclusively by the principles of Galilean and Cartesian mechanics, the Automaton and the Clockwork became the metaphors essential to the development of medicine [5]. The automaton of Descartes is a theoretical model, a metaphor of the human nature in every sense: it is as a means by which to understand through what is well-known (the working of a machine) what is obscure (the working of the human body)³. The idea of Man expressed in the *Traité* via that mental experiment of the automaton, first extraordinary intuition of something cybernetic, is the result not only of his studies on mechanics but, above all, of the experience Descartes had in anatomy. Writing *L'Homme* Descartes aimed at two objectives: to free the soul, placing it on the metaphysical level, from all the superstitions connected with the body, and to explain the whole functioning of the human body through sensory and motorial phenomena. That is, through movement and sensible perception which are, still nowadays, the two core questions of every speculation about the automaton.

Descartes was not the only promoter of that mechanist view of the human body to which we owe the great development of medicine in the 17th and 18th centuries. However, thanks to his extraordinary clarity the theoretic machine-body model proposed by the French philosopher had an exceptional influence especially in the field of medicine. The unavoidable reference, in this case, is to the development of iatromechanics and to the work of Giovanni Alfonso Borelli (1608-1679). The mechanist approach, which is the basis of his treatise *De motu animalium* (1680-81), allowed him to mathematize the working of the body by means of quantifying precisely the forces involved in movement. The idea of a machine-body emerges, as Restelli underlines, not only through lexical choices but also through the illustrative tables (see Fig. 5) in the treatise which Borelli carried out with the help of the

³ It is very interesting, regarding the use of the machine body metaphor, to notice that at the end of the treatise Descartes said what was implicit from the very beginning: it is not the machine which artificially imitates man, it is the human body which is made exactly as a machine shows [6].




Fig. 5. Giovanni Alfonso Borelli, *De motu animalium*, t. I, plate. IV, 1680-81

physician Lorenzo Bellini (1643-1703). Within *Discorsi di anatomia* we can find an extraordinary evidence of his idea that automatons may act as models for the concept of the human body [7]. Through the description of a series of sequences such as the perfectly coordinated movements of a galley's rowers, or that of the interaction between singers and instrumentalists within an orchestra, and eventually the one of some hydraulic automatons in a garden, Bellini intended to explain by analogy the working of some functions in the human body system. The machine-body metaphor, thus, proved to be in this period the most effective one to support the efforts of medical science, as we can notice by considering the approach used by Giorgio Baglivi (1668-1707) in his work *De praxi medica* (1696). In it he enunciated the reasons why he agreed with the mechanist idea of the body and, arguing that the body is just a combination of chemical and mechanical movements, he made up analogies between bodily organs and levers, bellows, and water pipes just as Descartes had done in his *Traité de l'homme*.

Therefore in the 18th century finally emerged a new notion of Nature and Man, especially thanks to observations in medical science. An ever increasing number of acquisitions in medicine showed how not only perceptions and emotions, but also intellect faculties such as memory, imagination, and intelligence were closely connected with certain organs, to their more or less pathologic states, to their anatomic structure. The materialism of the 18th century reached its peak with La Mettrie's and d'Holbach's philosophies. They rejected the metaphysical idea of a *res cogitans* separated from the *res extensa* and claimed the close causal dependence of the mind's activity from matter. The notion of the human body as a machine which winds up by itself the springs that make it move, living image of the perpetual motion [8] was so coherent with the shared scientific pattern that the greatest automaton makers of that time such as Jacques de Vaucasson (1709-1782) and Pierre Jaquet-Droz (1721-1790) were often encouraged by physicians and surgeons who believed their work to be useful to improve the research upon human biology. From then on, and for a very long time, the Machine was considered as something *Other from Man*, something to which to relate, something through which to define reality and the

human identity itself: thus machine became metaphor of human nature (see Fig. 6). According to Descartes the automaton is a mental experiment, a theoretical model which can explain the assumptions connected with the body and its principles, while La Mettrie stated that Man is, by all means, an organic machine.



Fig. 6. A picture from the movie *Metropolis* by Fritz Lang, 1927. The machine as the *Other from the Man*, the alter-ego of the man.

3 Computer Metaphor

The idea, suggested by La Mettrie, that all the faculties of the soul depend so much obviously on the peculiar organization of the brain and of the whole body that they associate themselves with it, was further developed during the revolution in the field of the philosophy of science which took place in the 1940s. In this period the research upon Nature and the physical basis of human intelligence through the *alterity* of machine was accomplished through the birth of cybernetics within that research program, started by von Neumann (*The Computer and the Brain*, 1954) nowadays known as artificial intelligence. This could be considered the extreme radicalisation of Cartesian dualism: the Cartesian *res cogitans/res extensa* binomial was transformed in the hardware/software binomial where the hardware stands for the brain, while the software stands for the mind.

The notion of feedback, introduced by the mathematician Norbert Wiener in the 1940s, soon shifted from its use in computer science to use in the psychological field, and in particular it became the core of cognitive science. Hence we can infer how deeply the philosophy of science based upon cybernetics has influenced the notion of consciousness. The best example of how the cybernetic research has greatly influenced the evolution of the studies of the human mind is the Turing machine. This theoretical device is supposed to manipulate symbols and in performing this supposed activity, it apparently simulates that typically human ability to operate with signs and upon signs which represents the action as well as the human realm of mind [9]. However, the most revealing image is the one posed by the neuronal metaphor

introduced by von Neumann. The parallel between the neuron and the working of logical circuits in computers is implied in the title of his work: *The Computer and the Brain*. The comparison is based upon a symmetry: the computer is a kind of brain and the brain is a kind of computer, one is the simulation of the other. Such a statement is so powerful that even today, as Porush underlines, we usually describe the neuron as a simple on/off switch [10].

The mechanical portrait of the nature of individual neurons: their conception as mere binary devices (so much similar to those logic switches conceived by the scientist for the realization of ENIAC and EDVAC calculators), recalls in many respects the portrait of the human body explained using bellows, water pipes, and springs in the machine-man imagined by Descartes.




Fig. 7. Fritz Kahn, *Das Leben des Menschen*, vol. 4, plate VIII, 1929. A metaphor of the cognitive processes.




Fig. 8. A picture of a cyborg from the novel *The planet of the Double Sun* by N. R. Jones, 1967.

Another interesting case of the interaction between the cybernetic paradigm and psychological research is Eric Berne's transactional analysis. Berne's work was largely influenced by the development of both the communication theory and the game theory. In fact, in his article *Concerning the nature of communication*, published in 1953 on *The Psychiatric Quarterly*, he analysed the relationship between psychiatry and cybernetics as well as the meaning of communication in both fields [11]. In that article Berne also illustrated, through analogy, the role of *noise* and *information* in the study of communication from both the psychological and the cybernetic point of view. He stated that, as opposed to what happens for the machines, in human communication *noise* is of more value than *information* since it is the *noise*, and not the *information*, which signals the state of the communicants. The analogy with the mathematical approach to the study of the communication is carried out by means of both explicit references to cybernetics, and the use of such words as *receiver*, *signal*, *transmission*. In that article Berne also put the bases of the psychological games that will further develop in *Games People Play* published in 1964 [12].

So, after the fascination generated by the exceptional achievements in the technological field, the idea of the human mind as an epiphenomenon of the organic structure of the brain was developed within the collective imagination as well as

among scientists, neurologists and psychologists (see Fig. 7). The icon of the new scientific pattern is the image of the cyborg: a man-machine hybrid that, since 1930s, populated the collective imagination, offering a wide variety of versions. As in the case of the automaton's metaphor, the collective imagination seems to perceive by intuition, through the cyborg's visual metaphor, the subsequent development of scientific research. Indeed the cyborg, as picture of collective imagery, appears in literature and figurative arts around the 1930's, though the acronym for this cross-breed man-machine was created only in 1960 by two physicians, Clynes and Kline, engaged in advanced studies aimed at the exploration of Space. The word *cyborg* means cybernetic organism, and being an acronym reminds us of the dualism mind-body. The image itself of a cyborg reintroduces this dualism, suggesting the analogy between the mind-machine binomial against the body-organism one. It is an image which reintroduces in a radical form the *res cogitans/res extensa* Cartesian dualism, following the interpretative pattern joined to the cybernetic perspective. In this case, too, we are facing a representation of the human mind; indeed the cyborg is an image of consciousness coherent with the current scientific pattern. However, analysing the cyborg metaphor we have detected some elements, common to other mythological figures, much older than the cyborg. Many of these elements suggest that there is more than the mere relation with the cultural and scientific context that produced it, behind the function of this figure. The analysis of these elements could provide an answer to the acute observation of Domenico Gallo: for centuries man has built machines and then wondered what they were [13].

4 A Pilot Study on Consciousness Visual Metaphors

So far we analysed how elements related to both the scientific and cultural contexts interact in the development of a collective representation of human nature. These representations synthesise Mankind's identity in the common perception and are translated, by the collective visual language, into particular images. However, there are other elements playing a role in the development of this type of representations. Many of these elements have long been the subject of the psychoanalytic debate. Visual language, shaping the collective representations, could translate some elements related to the psychological experience in iconographical data. Some of these iconographic elements end up to integrate the collective representation of human nature, allowing the interaction between cultural, scientific, and psychological instances.



Fig. 9. The Uroboros from Theodoros Pelecanos, *Synesius*, 1478

Jungian psychology focused on analysing a particular group of images whose main characteristics are being common to many different cultures and carrying, constantly, similar meanings. Most of these images are related to mythological and religious contexts but, according to Jungian theories, their significance does not allude only to religious contents. We are generally led to think of myth as something archaic, linked to old forms of religion, at times to superstition. However several studies have shown that myth and myth-making are something different from a folkloric and primitive attempt to explain natural phenomena. One definition of the myth could be a narrative-iconographic structure, whose elements have great symbolic value. Since these elements allow, because of their symbolic value, the projection of certain contents, it is possible to consider the myth as a metaphorical network. So *myth-making* could have an elaborating function aiming at understanding universal contents, via metaphorical images working as cognitive structures. The Jungian psychoanalyst Eric Neumann stated that the main function of the myth is to codify psychic dynamics in the form of representation typical of the unconscious: an aggregation of symbols that circumscribes what must be explained, understood and interpreted thus allowing a glance at what the symbols point at or try to express [14].




Fig. 10. The “Venus of Willendorf” statuette of a female figure, maybe a Great Goddess, Willendorf, Austria, 24,000-22,000 B.C

Studying the evolution of certain myths and archetypal figures, he argued that they might be referred to the representation of consciousness in the different phases of its evolution. In *The Origins and History of Consciousness*, Neumann maintains that different types of representations, both iconographic and narrative, correspond to different phases of individual consciousness's development. In other words, certain myths, whose protagonists are some specific archetypal figures, are metaphors of those psychic dynamics, belonging to all human beings, through which to the consciousness is developed. Neumann identifies four macro stages in the development of consciousness. Specific representations, iconographic themes and narrative structures of mythological character correspond to these stages.



Fig. 11. Gautier de Moap, Livre de Messire Lancelot du Lac, 1470. Sir Lancelot could be considered one of the best examples of the archetype of the hero.

The 1st stage is the *uroboric stage* (approximately age 1 to 5) is the first phase of man's evolving ego consciousness. Human consciousness feels itself embryonic, for the ego feels fully contained in the primordial symbol [15]. Corresponding to this stage are the following themes: the uroboros: the ancient symbol depicting a snake swallowing its own tail and forming a circle: it represents the total identity between consciousness and unconscious; the primordial egg rising from primordial water (see Fig. 9). The 2nd stage is the *Great Mother stage* (approximately age 5 to puberty) in that stage the rising consciousness feels itself separated from the unconscious, but not yet independent: like the young child, whose life still depends on his mother. Corresponding iconographic and narrative mythological motifs are: Great Mother-type goddess (see Fig. 10); Motherhood theme; Paredrotype characters. The 3rd stage is the *Divine Child stage* (or the Hero stage) (from youth onwards) when the ego-consciousness needs, to reach its complete self-determination, to accomplish the emancipation from the relationship with the unconscious: this is the state of self-consciousness. Corresponding narrative and iconographic themes are: the Hero of the light characters, dragons, snake-like monsters, and the fight against the monster (see Fig. 11).

The 4th stage was defined by Neumann the *Great Individual* (age undefined not everyone achieves this level of development). Consciousness have outrun the fear of the loss of its autonomy and its individuality, now it can assert its own subjectivity. This category of Great Individuals serves as a model for the development of individuality in humankind generally. The individual fate of the hero – and the creative Great Individual is indeed a hero – may be exception, but he is also the exemplar of a process which subsequently affects all individuals in varying degree [16]. The best example of the Archetype of the Great Individual, in Neumann's opinion, is the myth of the Egyptian God Osiris. However, there are many motifs, both iconographical and narrative, related to the Archetype of the Great Individual that characterise the cyborg figure.



Fig. 12. The Egyptian God Osiris represented as a spine (the Djed pillar) and the spine motif in some cyborg representations



Fig. 13. The eye motif in a Osiris representation and in some cyborg representations

The most relevant are: the body-prosthesis, the motifs of the overcoming of death and of the disintegration of the body, the accentuation of iconographic details of the head, and eyes as a metaphor of a superior intelligence (see Fig. 12 and Fig. 13). Some of these iconographical elements were included in a questionnaire aimed at knowing how people, using visual metaphors, employ archetypal images to describe the concept of *consciousness* as self-awareness.




Fig. 14. Consciousness visual metaphors quest.: scores in percentage for the 2nd part.

In a pilot study on consciousness's visual metaphors based on Neumann theory, a questionnaire consisting of 15 sets of images, each one composed of 5 images, was administered to 30 students, 24 female and 6 male, aged between 22 and 43 years. Each set of images was composed by images related to the different representations of the stages of consciousness, according to Neumann's theory.

The images selected were divided for typology, so that each set presented iconographic elements related to the different archetypes organized by topic: for example the topic of the 2nd set was *face's parts*, the topic of the 6th set was *famous people*, the topic of the 9th set was *stylised representations* etc. The questionnaire was divided in three parts: the first part was dedicated to the verbal metaphors, while the second and the third to visual metaphors. In the second part, constituted by 10 sets of




Fig. 15. Consciousness visual metaphors quest.: scores in percentage for the 3rd part.

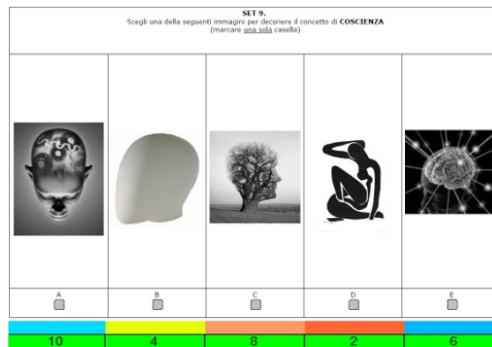


Fig. 16. A page of the questionnaire (part 2, set 9) with the score for each picture

images, the subject had to chose an image, only 1 out of 5, that best depicted, in his opinion, the idea of consciousness. In the third part, constituted by 5 sets of images, the subject had to put the images in order of relevance to the concept of "consciousness" giving a value from 1 to 5 (5 for the image that better depicted the concept of consciousness, 1 for the one that represented it less). The main result of the second part of the study was that images related to the 3rd and 4th stages are predominant: 21,7% for the first stage; 102.34% for the 2nd stage (48,16% Great Mother-type images + 54,18% Young Goddess-type images); 83% for the third stage, and 90 % for the fourth stage. The main result of the third part of the questionnaire was the predominance of the visual metaphors related to the 3rd and the 4th stage: 18% for the first stage; 40% for the 2nd stage (21% Great Mother-type images + 19% Young Goddess-type images); 20% for the third stage, and 22% for the fourth stage (see Fig. 14 and Fig. 15). However, the most significant data pertain to some images of the 7th, 9th, and 14th sets. It is very interesting, in particular, to notice that the first image of the 9th set: a human head in which, instead of the brain there are depicted the wheels of a complex gear, has obtained more preferences than the others (see Fig. 16).

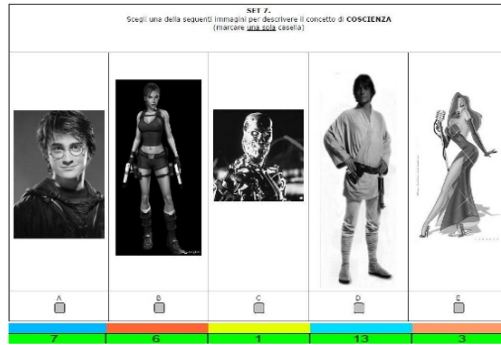


Fig. 17. A page of the questionnaire (part 2, set 7) with the score for each picture

While the third image of the 7th set: the picture of Terminator, a humanoid robot, has obtained only one preference (see Fig. 17). These results are even more interesting when we consider the entire set of images. Indeed, in the 9th set 10 students chose the brain-machine picture and 6 chose the plain brain one, while in the 7th set only 1 student, out of 30, selected the picture of a robot. This observation can be confirmed by the results of the 14th set that show how the picture of the cyborg Darth Vader, presented in comparison with those of other characters, obtained the lowest score (see Fig. 18). This suggests that the analogy with the machine-brain metaphor does not work when other factors enhancing the meaning of the metaphor are involved in the image. So, we can deduce that the machine metaphor can be used in different ways within the same cultural context. Indeed, it is frequently used, in an absolute sense, to describe concepts such as *mental process* or *mental activity*. While if it is accompanied by other elements that extend its meaning, it appears to be insufficient to describe the complexity of human nature.

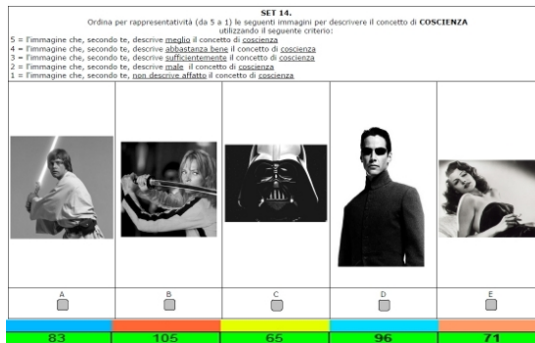


Fig. 18. A page of the questionnaire (part 3, set 14) with the score for each picture

Thus, apparently men create their mechanical self-portraits only to prove that they do not resemble to them. Therefore, apart from being images of man seen through the mirror of science and technology, these *artificial others* seem to have also a reassurance function for the human identity.

References

1. Lakoff, G., Johnson, M.: *Metaphors we live by*, p. 109. The University of Chicago Press, Chicago (2003)
2. Lakoff, G., Johnson, M.: *Metaphors we live by*, p. 3. The University of Chicago Press, Chicago (2003)
3. Sebeok, T., Danesi, M.: *The forms of meaning: modeling systems theory and semiotic analysis*, pp. 6–8. Walter de Gruyter, New York (2000)
4. Russo, L., Santoni, E.: *Ingegni minuti. Una storia della scienza in Italia*, p. 81. Feltrinelli, Milano (2010)
5. Sini, C.: *L'uomo, la macchina, l'automa. Lavoro e conoscenza tra futuro prossimo e passato remoto*, p. 16. Bollati Boringhieri, Torino (2009)
6. Sini, C.: *L'uomo, la macchina, l'automa. Lavoro e conoscenza tra futuro prossimo e passato remoto*, p. 15. Bollati Boringhieri, Torino (2009)
7. Restelli, A.: *Androidi musicali del XVIII secolo. Tra scienza ed attrazione*. In: *De Musica*, XII (2008), <http://users.unimi.it/~gpiana/demus.html> (accessed April 3, 2012)
8. La Mettrie, J.O.: *L'homme machine*. Elie Luzac, Leyde (1747)
9. Sini, C.: *L'uomo, la macchina, l'automa. Lavoro e conoscenza tra futuro prossimo e passato remoto*, p. 82. Bollati Boringhieri, Torino (2009)
10. Porush, D.: *The rise of cyborg culture or the bomb was a cyborg*. *Surfaces* 4(205), 1–32 (1994)
11. Berne, E.: *Concerning the Nature of Communication*. *The Psychiatric Quarterly* 27, 185–198 (1953)
12. Berne, E.: *Games People Play*. Random House, New York (1964)
13. Gallo, D.: *Lunga vita alla nuova carne*. *Psiche* II 119-129, 119 (2007)
14. Neumann, E.: *The origins and history of consciousness*, translated by Hull, R.F.C., p. 7. Routledge, London (1999)
15. Neumann, E.: *The origins and history of consciousness*, translated by Hull, R.F.C., p. 14. Routledge, London (1999)
16. Neumann, E.: *The origins and history of consciousness*, translated by Hull, R.F.C., p. 426. Routledge, London (1999)

Author Index

- Albano, Giuseppina 227
Alberto, Diego 139
Alessandroni, Silvio 279
Anzoise, Ilaria 437
- Bacciu, Davide 41
Barsocchi, Paolo 41
Bertoni, Alberto 179
Biella, Gabriele E.M. 189
Bloise, Elena 169
Boaro, Matteo 311
Bonanno, Lilla 121
Borghese, Alberto 11
Boucenna, Sofiane 345
Bramanti, Alessia 121
Bramanti, Placido 121
Brescia, Massimo 29
Bulotta, Simon 81
- Camastra, Francesco 159
Campo, Matteo 337
Capraro, Wiliam 11
Capuano, Vincenzo 395
Cardin, Marta 217
Cavalletti, Matteo 289
Cavallo, Alberto 3
Cavuoti, Stefano 29
Chanquoy, Lucile 357
Chessa, Stefano 41
Chetouani, Mohamed 345
Ciabattoni, Lucio 289
Ciaramella, Angelo 427
Cohen, David 345
Comminiello, Danilo 93, 101
- Corazza, Marco 217
Crisostomi, Emanuele 271
- D'Agostino, Maria Nicoletta 169
D'Amato, Valeria 201, 209
D'Angelo, Antonietta 169
D'Auria, Luca 111
De Angelis, Francesco 311
De Felice, Domenico 159
Delaherche, Emilie 345
Di Taranto, Maria Donata 169
- Esposito, Anna 367, 395
Esposito, Antonietta M. 111
- Fagiani, Marco 69
Falvo, Maria Carmen 259
Fartoukh, Michaël 357
Faundez-Zanuy, Marcos 367, 375
Ferone, Alessio 129
Filippone, Maurizio 337
Font-Aragones, Xavier 375
Fortunato, Giuliana 169
Frasca, Marco 179
Funari, Stefania 217
Fuselli, Danilo 311
- Gallicchio, Claudio 41
Garofalo, Mauro 29
Garre-Olmo, Josep 367
Gelsomino, Giuliana 189
Gentile, Marco 169
Giove, Silvio 217
Giudicepietro, Flora 111
Glodek, Michael 323, 385

- Gori, Marco 51
 Grassi, Marco 301
- Iannuzzi, Arcangelo 169
 Inuso, Giuseppina 59
 Ippoliti, Gianluca 289
- Jossa, Fabrizio 169
- Labate, Domenico 59
 La Foresta, Fabio 59
 Lanzafame, Pietro 121
 La Rocca, Michele 227
 Lengvenis, Paulius 405
 Lippi, Marco 51
 Longhi, Sauro 289
 Longo, Giuseppe 29
 Lopez-de-Ipiña, Karmele 367
- Maddalena, Lucia 129
 Mahmoud, Hassan 81
 Malchiodi, Dario 21
 Maldonato, Mauro 437
 Marotta, Gennaro 169
 Martini, Marcello 111
 Martirano, Luigi 259
 Mascioli, Fabio Massimo Frattale 279
 Maskeliunas, Rytis 405
 Masulli, Francesco 81
 Mekyska, Jiri 367
 Menconi, Lorenzo 51
 Mesin, Luca 139
 Micheli, Alessio 41
 Morabito, Francesco C. 59
- Nencini, Sara 189
 Nucci, Michele 301
- Occhiuto, Gianluigi 59
- Palm, Günther 323, 385
 Palmieri, Francesco A.N. 3
 Palummeri, Ernesto 81
 Parisi, Raffaele 93, 101
 Paschero, Maurizio 279
 Pasero, Eros 139, 147
 Perna, Cira 227
- Pescapè, Antonio 29
 Petrone, Giovanni 239
 Petrosino, Alfredo 129
 Piazza, Francesco 69, 301, 311
 Piolat, Annie 357
 Pirro, Matteo 289
 Piscopo, Gabriella 201, 209
 Polychroniou, Anna 337
 Possemato, Francesca 279
 Principi, Emanuele 69
- Raudonis, Vidas 405
 Raugi, Marco 271
 Riviello, Maria Teresa 395
 Rizzi, Antonello 279
 Ros, Paolo Motto 147
 Roure-Alcobe, Josep 367
 Rovetta, Stefano 81
 Rubba, Paolo 169
 Russolillo, Maria 201, 209
- Salamin, Hugues 337
 Sanchez, Liliانا Mamani 415
 Sánchez Pacheco, Francisco Jose 239
 Sbordone, Danilo 259
 Scardapane, Simone 93
 Scarpiniti, Michele 93, 101
 Schels, Martin 385
 Schwenker, Friedhelm 385
 Sesa-Nogueras, Enric 367
 Sion, Andrea 179
 Spagnuolo, Giovanni 239
 Squartini, Stefano 69, 311
 Staiano, Antonino 169
 Storti, Gian Luca 279
- Tucci, Mauro 271
- Uncini, Aurelio 93, 101
- Ventre, Giorgio 29
 Vettigli, Giuseppe 427
 Vinciarelli, Alessandro 337
 Vogel, Carl 415
- Zippo, Antonio G. 189