

IFIP AICT 331

Max Bramer
(Ed.)

Artificial Intelligence in Theory and Practice III

Third IFIP TC 12 International Conference
on Artificial Intelligence, IFIP AI 2010
Held as Part of WCC 2010
Brisbane, Australia, September 2010
Proceedings

 Springer



World
Computer
Congress
Brisbane 2010
International Federation for Information Processing

Editor-in-Chief

A. Joe Turner, Seneca, SC, USA

Editorial Board

Foundations of Computer Science

Mike Hinchey, Lero, Limerick, Ireland

Software: Theory and Practice

Bertrand Meyer, ETH Zurich, Switzerland

Education

Bernard Cornu, CNED-EIFAD, Poitiers, France

Information Technology Applications

Ronald Waxman, EDA Standards Consulting, Beachwood, OH, USA

Communication Systems

Guy Leduc, Université de Liège, Belgium

System Modeling and Optimization

Jacques Henry, Université de Bordeaux, France

Information Systems

Barbara Pernici, Politecnico di Milano, Italy

Relationship between Computers and Society

Chrisanthi Avgerou, London School of Economics, UK

Computer Systems Technology

Paolo Prinetto, Politecnico di Torino, Italy

Security and Privacy Protection in Information Processing Systems

Kai Rannenber, Goethe University Frankfurt, Germany

Artificial Intelligence

Max A. Bramer, University of Portsmouth, UK

Human-Computer Interaction

Annelise Mark Pejtersen, Center of Cognitive Systems Engineering, Denmark

Entertainment Computing

Ryohei Nakatsu, National University of Singapore

IFIP – The International Federation for Information Processing

IFIP was founded in 1960 under the auspices of UNESCO, following the First World Computer Congress held in Paris the previous year. An umbrella organization for societies working in information processing, IFIP's aim is two-fold: to support information processing within its member countries and to encourage technology transfer to developing nations. As its mission statement clearly states,

IFIP's mission is to be the leading, truly international, apolitical organization which encourages and assists in the development, exploitation and application of information technology for the benefit of all people.

IFIP is a non-profitmaking organization, run almost solely by 2500 volunteers. It operates through a number of technical committees, which organize events and publications. IFIP's events range from an international congress to local seminars, but the most important are:

- The IFIP World Computer Congress, held every second year;
- Open conferences;
- Working conferences.

The flagship event is the IFIP World Computer Congress, at which both invited and contributed papers are presented. Contributed papers are rigorously refereed and the rejection rate is high.

As with the Congress, participation in the open conferences is open to all and papers may be invited or submitted. Again, submitted papers are stringently refereed.

The working conferences are structured differently. They are usually run by a working group and attendance is small and by invitation only. Their purpose is to create an atmosphere conducive to innovation and development. Refereeing is less rigorous and papers are subjected to extensive group discussion.

Publications arising from IFIP events vary. The papers presented at the IFIP World Computer Congress and at open conferences are published as conference proceedings, while the results of the working conferences are often published as collections of selected and edited papers.

Any national society whose primary activity is in information may apply to become a full member of IFIP, although full membership is restricted to one society per country. Full members are entitled to vote at the annual General Assembly. National societies preferring a less committed involvement may apply for associate or corresponding membership. Associate members enjoy the same benefits as full members, but without voting rights. Corresponding members are not represented in IFIP bodies. Affiliated membership is open to non-national societies, and individual and honorary membership schemes are also offered.

Max Bramer (Ed.)

Artificial Intelligence in Theory and Practice III

Third IFIP TC 12 International Conference
on Artificial Intelligence, IFIP AI 2010
Held as Part of WCC 2010
Brisbane, Australia, September 20-23, 2010
Proceedings

Volume Editor

Max Bramer
University of Portsmouth, School of Computing
Lion Terrace, Portsmouth, Hants, PO1 3HE, UK
E-mail: Max.Bramer@port.ac.uk

Library of Congress Control Number: 2010932510

CR Subject Classification (1998): I.2, H.3, H.4, H.2.8, F.1, H.5

ISSN 1868-4238
ISBN-10 3-642-15285-6 Springer Berlin Heidelberg New York
ISBN-13 978-3-642-15285-6 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

springer.com

© International Federation for Information Processing 2010
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper 06/3180

IFIP World Computer Congress 2010 (WCC 2010)

Message from the Chairs

Every two years, the International Federation for Information Processing (IFIP) hosts a major event which showcases the scientific endeavors of its over one hundred technical committees and working groups. On the occasion of IFIP's 50th anniversary, 2010 saw the 21st IFIP World Computer Congress (WCC 2010) take place in Australia for the third time, at the Brisbane Convention and Exhibition Centre, Brisbane, Queensland, September 20–23, 2010.

The congress was hosted by the Australian Computer Society, ACS. It was run as a federation of co-located conferences offered by the different IFIP technical committees, working groups and special interest groups, under the coordination of the International Program Committee.

The event was larger than ever before, consisting of 17 parallel conferences, focusing on topics ranging from artificial intelligence to entertainment computing, human choice and computers, security, networks of the future and theoretical computer science. The conference History of Computing was a valuable contribution to IFIP's 50th anniversary, as it specifically addressed IT developments during those years. The conference e-Health was organized jointly with the International Medical Informatics Association (IMIA), which evolved from IFIP Technical Committee TC-4 "Medical Informatics".

Some of these were established conferences that run at regular intervals, e.g., annually, and some represented new, groundbreaking areas of computing. Each conference had a call for papers, an International Program Committee of experts and a thorough peer reviewing process of full papers. The congress received 642 papers for the 17 conferences, and selected 319 from those, representing an acceptance rate of 49.69% (averaged over all conferences). To support interoperation between events, conferences were grouped into 8 areas: Deliver IT, Govern IT, Learn IT, Play IT, Sustain IT, Treat IT, Trust IT, and Value IT.

This volume is one of 13 volumes associated with the 17 scientific conferences. Each volume covers a specific topic and separately or together they form a valuable record of the state of computing research in the world in 2010. Each volume was prepared for publication in the Springer IFIP Advances in Information and Communication Technology series by the conference's volume editors. The overall Publications Chair for all volumes published for this congress is Mike Hinchey.

For full details of the World Computer Congress, please refer to the webpage at <http://www.ifip.org>.

June 2010 Augusto Casaca, Portugal, Chair, International Program Committee
Phillip Nyssen, Australia, Co-chair, International Program Committee
Nick Tate, Australia, Chair, Organizing Committee
Mike Hinchey, Ireland, Publications Chair
Klaus Brunnstein, Germany, General Congress Chair

Preface

The papers in this volume comprise the refereed proceedings of the conference Artificial Intelligence in Theory and Practice (IFIP AI 2010), which formed part of the 21st World Computer Congress of IFIP, the International Federation for Information Processing (WCC-2010), in Brisbane, Australia in September 2010.

The conference was organized by the IFIP Technical Committee on Artificial Intelligence (Technical Committee 12) and its Working Group 12.5 (Artificial Intelligence Applications).

All papers were reviewed by at least two members of our Program Committee. Final decisions were made by the Executive Program Committee, which comprised John Debenham (University of Technology, Sydney, Australia), Ilias Maglogiannis (University of Central Greece, Lamia, Greece), Eunika Mercier-Laurent (KIM, France) and myself. The best papers were selected for the conference, either as long papers (maximum 10 pages) or as short papers (maximum 5 pages) and are included in this volume. The international nature of IFIP is amply reflected in the large number of countries represented here.

I should like to thank the Conference Chair, Tharam Dillon, for all his efforts and the members of our Program Committee for reviewing papers under a very tight deadline.

This is the latest in a series of conferences organized by IFIP Technical Committee 12 dedicated to the techniques of artificial intelligence and their real-world applications. The wide range and importance of these applications is clearly indicated by the papers in this volume. Further information about TC12 can be found on our website at <http://www.ifiptc12.org>.

Max Bramer

Organization

Organizing Committee

General Chair

Tharam Dillon Curtin University of Technology, Australia

Program Chair

Max Bramer University of Portsmouth, UK

Executive Program Committee

Max Bramer	University of Portsmouth, UK
John Debenham	University of Technology, Sydney, Australia
Ilias Maglogiannis	University of Central Greece, Lamia, Greece
Eunika Mercier-Laurent	KIM, France

Program Committee

Adel Alimi	University of Sfax, Tunisia
Fabrizio Angiulli	DEIS, University of Calabria, Italy
Stefania Bandini	University of Milan, Italy
Cristian Barrué	Technical University of Catalonia, Barcelona, Spain
Daniel Berrar	University of Ulster, UK
Alina Bogan-Marta	University of Oradea, Romania
Max Bramer	University of Portsmouth, UK
Maricela Bravo	Centro de Investigacion y de Estudios Avanzados del IPN, Mexico
Per Bjarne Bro	Chile
Kryisia Broda	Imperial College, London, UK
Luigia Carlucci Aiello	Università di Roma La Sapienza, Italy
Ana Casali	FCEIA - UNR, Argentina
Fabrice Colas	UCLA, USA
John Debenham	University of Technology, Sydney, Australia
Yves Demazeau	CNRS - LIG Grenoble, France
Vladan Devedzic	University of Belgrade, Serbia
Tharam Dillon	Curtin University of Technology, Australia
Gracaliz Pereira Dimuro	Universidade Federal do Rio Grande, Brazil

Gintautas Dzemyda	Institute of Mathematics and Informatics, Lithuania
Henrik Ericsson	Linköping University, Sweden
Enrique Ferreira	Universidad Católica del Uruguay
Anabel Fraga	Spain
Matjaz Gams	Slovenia
Ana Garcia-Serrano	Universidad Nacional de Educación a Distancia, Madrid, Spain
Daniela Godoy	ISISTAN Research Institute, Argentina
Enrique González	Pontificia Universidad Javeriana, Colombia
Kostas Karpouzis	National Technical University of Athens, Greece
Nik Kasabov	Auckland University of Technology, New Zealand
Nicolas Kemper Valverde	Universidad Nacional Autónoma de México
Joost N. Kok	Leiden University, The Netherlands
Jasna Kuljis	Brunel University, UK
Daoliang Li	China Agricultural University, Beijing, P. R. China
Aurelio Lopez-Lopez	Instituto Nacional de Astrofisica, Optica y Electronica, Mexico
Ramon Lopez de Mantaras	Spanish National Research Council
Suresh Manandhar	University of York, UK
Edson Takashi Matsubara	University of Sao Paulo (USP), Brazil
Brian Mayoh	University of Aarhus, Denmark
Eunika Mercier-Laurent	KIM, France
Antonija Mitrovic	University of Canterbury, Christchurch, New Zealand
Riichiro Mizoguchi	Osaka University, Japan
Pavol Navrat	Slovak University of Technology in Bratislava, Slovakia
Adolfo Neto	Federal University of Technogy - Parana (UTFPR) Brazil
Erich J. Neuhold	University of Vienna, Austria
Bernd Neumann	University of Hamburg, Germany
Toyoaki Nishida	Kyoto University, Japan
Daniel O'Leary	University of Southern California, USA
Andrea Omicini	Alma Mater Studiorum - Università di Bologna, Italy
Mihaela Oprea	University of Ploiesti, Romania
Cristian Rusu	Pontificia Universidad Católica de Valparaíso, Chile
Abdel-Badeeh M. Salem	Faculty of Computer and Information Sciences, Ain Shams University, Egypt
Demetrios Sampson	University of Piraeus, Greece
Silvia Schiaffino	ISISTAN Research Institute, Argentina
Valery Sklyarov	Portugal
Mauricio Solar	Universidad Tecnica Federico Santa Maria, Chile
Georgina Stegmayer	CIDISI Research Institute-CONICET, Argentina
Ahmed Tawfik	French University in Egypt
Eric Tsui	The Hong Kong Polytechnic University

Wiebe van der Hoek
Jin-Mao Wei
Ting Yu
Zdenek Zdrahal
Alejandro Zunino

University of Liverpool, UK
Nankai University, China
University of Sydney, Australia
The Open University, UK
ISISIAN, UNICEN, Argentina

Table of Contents

Part I: Agents

Ecologically Rational Agency	3
<i>John Debenham and Carles Sierra</i>	
Involving the Human User in the Control Architecture of an Autonomous Agent	13
<i>Roberto Micalizio, Giancarlo Nuzzolo, Enrico Scala, and Pietro Torasso</i>	

Part II: Knowledge Acquisition

Exer-Learning Games: Transferring Hopscotch from the Schoolyard to the Classroom	25
<i>Martina Lucht, Steffi Domagk, and Martin Mohring</i>	
Social Relationships as a Means for Identifying an Individual in Large Information Spaces	35
<i>Katarína Kostková, Michal Barla, and Mária Bielíková</i>	

Part III: Data Mining and Machine Learning

J-PMCRI: A Methodology for Inducing Pre-pruned Modular Classification Rules	47
<i>Frederic Stahl, Max Bramer, and Mo Adda</i>	
Enhancement of Infrequent Purchased Product Recommendation Using Data Mining Techniques	57
<i>Noraswaliza Abdullah, Yue Xu, Shlomo Geva, and Mark Looi</i>	
A Machine Learning Approach to Predicting Winning Patterns in Track Cycling Omnium	67
<i>Bahadorreza Ofoghi, John Zeleznikow, Clare MacMahon, and Dan Dwyer</i>	

Part IV: Neural Nets and Fuzzy Logic

Learning Motor Control by Dancing YMCA	79
<i>Rikke Amilde Løvlid</i>	
Analysis and Comparison of Probability Transformations for Fusing Sensors with Uncertain Detection Performance	89
<i>Edwin El-Mahassni, Samuel Davey, and Jonathan Legg</i>	

Part V: Case-Based Reasoning/Planning and Scheduling

A Case-Based Approach to Business Process Monitoring 101
Stefania Montani and Giorgio Leonardi

A Survey on the Dynamic Scheduling Problem in Astronomical
 Observations 111
Matias Mora and Mauricio Solar

Part VI: Intelligent Systems Engineering

Combining Determinism and Intuition through Univariate Decision
 Strategies for Target Detection from Multi-sensors 123
Edwin El-Mahassni

A UML Profile Oriented to the Requirements Modeling in Intelligent
 Tutoring Systems Projects 133
Gilleanes Thorwald Araujo Guedes and Rosa Maria Vicari

Learning by Collaboration in Intelligent Autonomous Systems 143
Jorge Ierache, Ramón García-Martínez, and Armando De Giusti

Part VII: AI Techniques

Context-Based Probabilistic Scene Interpretation 155
Bernd Neumann and Kasim Terzic

Full Text Search Engine as Scalable K-Nearest Neighbor
 Recommendation System 165
Ján Suchal and Pavol Návrat

Developing Intelligent Environments with OSGi and JADE 174
Davide Carneiro, Paulo Novais, Ricardo Costa, and José Neves

Part VIII: AI Applications

Bee Hive At Work: Following a Developing Story on the Web 187
Pavol Navrat and Anna Bou Ezzeddine

Computer-Aided Estimation for the Risk of Development of Gastric
 Cancer by Image Processing 197
*Yoshihiro Sasaki, Ryukichi Hada, Tetsuro Yoshimura,
 Norihiro Hanabata, Tatsuya Mikami, and Shinsaku Fukuda*

Intelligent Hybrid Architecture for Tourism Services 205
Ana Almeida, Bruno Coelho, and Constantino Martins

Knowledge-Based Geo-risk Assessment for an Intelligent Measurement System	215
<i>T. Vicovac, A. Reiterer, U. Egly, T. Eiter, and D. Rieke-Zapp</i>	

Part IX: Short Papers

Workflows Recognition through Multi Agents in Surveillance Systems...	227
<i>Manolis Sardis, Vasilis Anagnostopoulos, and Theodora Varvarigou</i>	
Speaker Verification System Using a Hierarchical Adaptive Network-Based Fuzzy Inference Systems (HANFIS)	233
<i>Zohreh Soozanchi-K., Mohammad-R. Akbarzadeh-T., Mahdi Yaghoobi, and Saeed Rahati</i>	
Case-Based Decision Support in Time Dependent Medical Domains	238
<i>Stefania Montani</i>	
Solving the Capacitated Vehicle Routing Problem and the Split Delivery Using GRASP Metaheuristic	243
<i>Joseph Gallart Suárez and Manuel Tupia Anticona</i>	
Author Index	251

PART I
Agents

Ecologically Rational Agency

John Debenham¹ and Carles Sierra²

¹ Centre for Quantum Computation & Intelligent Systems,
University of Technology, Sydney, Australia

² IIIA, CSIC, Campus UAB, 08193 Bellaterra, Catalonia, Spain

Abstract. The concept of dual rationality in human agents is well understood. It is significant in the work of the economist and Nobel Laureate Friedrich Hayek. In psychology Sigmund Freud described ‘dual information processing’. More recently in cognitive science Daniel Levine has extensively studied the dual mechanisms in the human brain that handle these two forms of rationality. It is perhaps surprising that work in artificial intelligence has focussed heavily on Cartesian constructivist rationality (particularly on game theory, decision theory and logic) as humans deliberate using models based on reason for only a minute proportion of the decisions that they make. This paper aims to open discussion on ecological rationality in artificial, computerised agents, and to examine some of the issues in designing ecologically rational agents.

1 Introduction

This paper is concerned with *bounded rationality* that dates back to David Hume and more recently to the early work of Herbert Simon. Bounded rationality refers to systems that are not founded on Cartesian rationalism; it has been widely addressed in economics [9], and is discussed in all good books on artificial intelligence, e.g. [10]. For over fifty years artificial intelligence research has spawned countless theories and systems that are *not* founded on Cartesian rationalism; one classic contribution being Rodney Brooks’ work reported in his ‘Computers and Thought’ award-winning paper [2]. Despite these advances, work in multiagent systems has been heavily influenced by game theory, decision theory and logic [11].

In economics two forms of rationality are described by the two Nobel Laureates Friedrich Hayek [7] and Vernon Smith [14] as being within ‘two worlds’. Hayek and Smith identify; *constructivist rationality* that underpins rational predictive models of decision making; and, *ecological rationality* founded on the concept of “spontaneous order” that refers to social institutions and practices that *emerge* from the history of an agent’s interactions and are *not* pre-designed. For computerised, intelligent agents the predominant logical distinction is between *deliberative* and *reactive* logic. Hayek and Smith’s two rationalities relate directly to two distinct forms of deliberation, and have

¹ The term ‘order’ refers to: traditions, customs, norms, rules and guidelines. An agent may belong to a number of normative systems (or, electronic institutions [1]) whose norms may be shared with, or in conflict with, those of other systems. The ‘extended order’ includes the whole show. If a multiagent system interacts with human society then its norms will respect the rules and laws that apply to society as a whole.

little to do with autonomic reactivity that typically overrides other processes in both the human neuropsychological system and in intelligent agents.

In psychology ‘dual information processing’ is a key concept as Seymour Epstein observed [4]: “There is no dearth of evidence in every day life that people apprehend reality in two fundamentally different ways, one variously labeled intuitive, automatic, natural, nonverbal, narrative, and experiential, and the other analytical, deliberative, verbal, and rational, (p. 710)”. These ideas date back to Sigmund Freud [5]. More recently, the different brain processes that support dual rationality have been extensively studied in cognitive science by Daniel Levine [8]. But, in artificial intelligence the term ‘rationality’ appears to apply exclusively to constructivist rationality.

We assume that the following is true of all real systems and environments: *There are always facts that an agent does not know that may effect the behaviours of other agents or states of the environment.* In particular: *Agent α expects to know less about agent β 's commitments than β knows.* This is *not* about pervasive uncertainty, it is about a limit to that which is known — a necessary ignorance that is distinct from, but related to, Murphy’s law “Anything that can go wrong will go wrong” — we don’t and can’t know the extent of the “anything”.

We also assume that the following is true of all real systems and environments: *Every fact that an agent knows is uncertain.* This *is* about pervasive uncertainty. When an observation is made it may be momentarily assumed to be true but as time progresses its truth will become more and more uncertain.

The term *rational* means “showing clear thought or reason”, and so to be rational in real environments means: to show clear thought or reason in the presence of necessary ignorance and pervasive uncertainty. Two foundations for rational behaviour are:

Cartesian constructivism. Cartesian constructivism deals with the unknown by assuming a theory that: (i) specifies what facts are required to be known, and (ii) makes assumptions about *some* of the unknowns. See Section 3.1. Game theory is a well known example of Cartesian constructivism, and requires that an agent knows its *actions, consequences, a consequence function* and a *preference relation*; and assumes that each agent will act so as to maximise its expected utility.

Ecological rationality. This is deliberation that uses past experience and contextual triggers to build action sequences from experiential memory. Past experience is a precursor to ecological rationality. For example, as we have described them previously, trust and honour [12] and reputation [13], are purely ecological concepts [3]. Building action sequences from experiential memory involves more than just retrieval. An agent has: to learn to imitate the actions that it believes that others do, to form expectations of the effect of actions, to select actions from a set of candidates, to adapt actions to suit the current norms and state of the environment, and when things don’t work out to learn to experiment with untested actions.

There may be situations in which neither Cartesian constructivism or ecological rationality may be useful, and there may be situations in which they may both be useful.

As the name suggests, ecological rationality is concerned with a richer form of bounded rationality than simplifying the calculation of a theoretically ‘optimal’ action by: rules for simplifying search, rules for terminating search or heuristic decision rules

to select actions from an incomplete set of options as described in [6]. Ecological rationality is taken in the context of the Hayekian view [7] in which agents evolve themselves together with the norms of the systems they inhabit² whilst their environment changes. This all sounds rather Darwinian, but Hayek is careful to distinguish between genetic evolution and cultural evolution [op. cit. page 23].

In previous work we have described some of the tools that an ecologically rational agent will require [3]. In this paper we describe the essence of an ecologically *intelligent* agent. This paper is organised as follows. Various preliminaries are described in Section 2. Section 2.1 discusses the capabilities that rational agents require. Section 3 discusses dual rationality for computerised agents; Section 3.1 cartesian rationality and Section 3.2 ecological rationality. Section 4 describes a basis for the rational deliberative process, and Section 5 concludes with a discussion on what it means for a rational agent to be “intelligent”.

2 Preliminaries

Agent α is in a *multiagent system* with a finite but unknown number of other agents $\mathcal{B} = \{\beta_1, \beta_2, \dots\}$, and a finite number of *information providing agents* $\Theta = \{\theta_1, \theta_2, \dots\}$ that provide the *context* for all events in the system. Θ^t denotes the state of these agents at time t , and \mathcal{I} is the set of all these possible information states. The information providing agents provide freely available data to anybody who wants to listen. This includes, day of the week, time of day, and so on, as well as news for general distribution. In addition, there is an *information search service*³ that proactively seeks information. In so far as the information providing agents are equivalent to an electronic news sheet the search service is equivalent to a search engine, for which, in some cases, a fee may be charged, particularly for confidential information.

The only thing that an agent ‘knows for certain’ is the history of its past communication. Agent α ’s *history* at time t , \mathcal{H}_α^t , is the complete history of: its prior interactions with other agents, observations of information providing agents, and information proactively obtained from the search service. Each *utterance* in the history contains: an illocutionary statement, the sending agent, the receiving agent, the time that the utterance was sent or received. Utterances are organised into dialogues, where a *dialogue* is a finite sequence of related utterances. If a dialogue is not yet complete then it is an *open dialogue*, that is an interaction that an agent is currently “working on”. For convenience we assume that each agent has at most one open dialogue — that is, it participates in at most one “sequence of related utterances” at a time.

An agent acts to satisfy a *need*, v , that are always considered in context, (v, Θ^t) . An agent acts by communicating an utterance, (μ, β) , containing an illocutionary statement, $\mu \in \mathcal{M}$, to another agent, $\beta \in \mathcal{B}$. If an utterance is part of a complete dialogue,

² The evolution of individual agents and component systems are not considered in isolation — the whole ensemble evolves in response to itself and to the environment — they are *complex* systems. For example, in Hayek’s extensive writing there is little mention of ethics as it too evolves.

³ This does not feature in the discussion but it is essential if agents are to proactively seek information and information sources. We just assume it is available.

d , that aimed to satisfy a need then the dialogue is tagged with: the triggering need, v , the prevailing context, Θ' , and a rating $r \in R$ of how satisfactorily the dialogue satisfied the need. So such a *rated dialogue* will have the form: $d = (d, v, \Theta', r) \in \mathcal{H}_\alpha$.

We assume that agents aim to satisfy needs through participating in dialogues that will hopefully achieve a high rating. The *rating* of a dialogue is an *ex post* evaluation of the outcome of the dialogue taking account of the need that triggered it and changes in the world state, or the evaluating agent's state, that occurred during it. A rating acknowledges that the partner agents involved may have taken account of changes in circumstance that occurred during the dialogue, or have even gone “over the odds” and gave more than was expected in some sense. It also includes the possibility that the partner agent may have taken advantage of information asymmetry between the agents, and acted in a way that is contrary to that which α would have desired had he known all the facts. Ratings are on a discrete, finite fuzzy scale.

Ratings are only defined at the *completion* of the dialogue and so it can not play a role *before* or *during* the dialogue — ratings are not ‘utility functions’ in any sense — they are subjective, *ex post* assessments of outcomes of dialogues that account for the changing context during the entire interaction process. An agent may select an interaction partner and direct the subsequent dialogue in the hope that it will achieve a high rating — in reality that is the best he can do. It will only be possible to guarantee a high rating if the agent knows in advance⁴ all the facts that could effect the rating — in reality *this never happens*.

2.1 Capabilities of Rational Agents

An agent should be able to:

- *observe* and *understand* the behaviour of other agents — understanding involves knowing or guessing the other agents’ motivations and *judging* the way in which the behaviours satisfy those motivations. In the special case that the agent receives instruction, it is reasonable to assume that behaviours are presented to the agent with a range of goals or motivations that the behaviours achieve or satisfy. The idea of “teaching or instructing agents” is much simpler than having them observe and understand.
- *imitate* the behaviour of other agents — imitation is more than simple copying what another agent does, it is using another agent’s behaviour to satisfy its own motivations so necessarily involves the *assessment* of those behaviours in its own context.
- *remember* what it does and what it observes — the agent’s memory should be constructed with efficiency of recall in mind. This means that the agent may *classify* observations based on some value judgement about their potential future use. There is no point in remembering rubbish. It may useful to *abstract* observations to some degree to facilitate adaptation and reuse, even if this amounts only to simple “typing”. In any case each behaviour memorised is *rated* for how well it achieved its goal. In the interests of efficiency it may be desirable for the memory to contain

⁴ In game theory, utility is in a sense an *a priori* rating. Utility is a rating of outcomes and is required by the postulates of game theory to be known in advance.

summary measures of trust, honour, reliability, reputation, integrity etc, and for these summaries to be updated each time an observation is made. So in addition to representing the facts of an observation, remembering involves: rating, abstracting, classifying and summarising.

- *recall* relevant past behaviours, trust estimates, reputation estimates etc given a new motivation, need or goal. This may be seen as a complex form of semantic difference.
- *act* rationally — i.e. its acts in a way that it believes will achieve a high *ex post* rating. The agent should be striving to optimise its expected ratings. This does *not* mean that the agent is a game-theoretic agent with a utility function defined by the ratings for the following reasons:
 - ratings are an *ex post* assessment — any prior expectation of their value will be ignorant of circumstances at assessment time and can only be based on wishful thinking and not on any logical model — that is, no method can claim to predict, or to optimise, ratings,
 - as ratings are subjective there is no way that one agent can predict how another agent will rate a dialogue, and
 - an agent may select a non-Cartesian (i.e. ecological) basis for its actions in which case they will not be founded on a model that is capable of forecasting the outcome of performing those actions.

How is an agent to “show clear thought or reason in the presence of necessary ignorance and pervasive uncertainty”? If an agent has not been hard-wired with models of the world or other agents then all that it knows is in its memory. This naturally leads to re-use and adaptation as in case-based reasoning.

3 Dual Rationality for Artificial Agents

In this section we define rationality in both its constructivist and ecological form.

3.1 Cartesian Constructivist Rationality

Given a finite set X , let $\Delta(X)$ denote the set of probability distributions over X . A *constructivist rational* agent α with chosen theory T_α and need v , α will first obtain the data $D_\alpha^t = T_{\alpha D}(\mathcal{H}_\alpha^t, v, \Theta^t)$ that T_α requires given \mathcal{H}_α^t , v and Θ^t . Second, α will act using the lottery:

$$C_\alpha : \mathcal{H} \times \mathcal{N} \times \mathcal{I} \times \mathcal{T} \times \mathcal{D} \rightarrow \Delta(\mathcal{M} \times \mathcal{B}) \quad (1)$$

where \mathcal{H} the set of all history states, \mathcal{N} is the set of need states, \mathcal{I} is the set of information states, \mathcal{T} is the set of theory states⁵, and \mathcal{D} is the set of data states as can be extracted from the search service. As above, \mathcal{M} is the set of illocutionary statements and \mathcal{B} is the set of other agents. C_α is a function that encapsulates α 's interpretation of T_α . For example, if α has an open dialogue then Equation 1 determines (non-deterministically) the next utterance that α makes in that dialogue.

⁵ The elements of \mathcal{T} are theories that may be chosen.

Constructivist rational agents are non-deterministic. This form of non-determinism is similar to mixed strategies in game theory. The action performed by a *deterministic constructivist rational agent* is determined by: $C_\alpha : \mathcal{H} \times \mathcal{N} \times \mathcal{I} \times \mathcal{T} \times \mathcal{D} \rightarrow \mathcal{M} \times \mathcal{B}$; they are similar to pure strategies in game theory.

3.2 Ecological Rationality

An *ecologically rational agent* α with need v in context Θ^t will act using the lottery $E_\alpha(\mathcal{H}_\alpha^t, v, \Theta^t) \in \Delta(\mathcal{M} \times \mathcal{B})$ where:

$$E_\alpha : \mathcal{H} \times \mathcal{N} \times \mathcal{I} \rightarrow \Delta(\mathcal{M} \times \mathcal{B}) \quad (2)$$

where E_α is a function that is *not* founded on an abstraction or theory that models, explains, constrains, describes or prescribes the behaviour of agents or the environment; it encapsulates α 's particular ecological rationality. As above, if α has an open dialogue then Equation 2 determines (non-deterministically) the next utterance that α makes in that dialogue. Ecologically rational agents are non-deterministic. The action performed by a *deterministic ecologically rational agent* is determined by: $E_\alpha : \mathcal{H} \times \mathcal{N} \times \mathcal{I} \rightarrow \mathcal{M} \times \mathcal{B}$. We make two observations:

- E_α , and any other embedded reasoning, should *not* be founded on an abstraction or theory that models, explains, constrains, describes or prescribes the behaviour of agents or the environment. The sum total of the wisdom available to E_α is that which has been observed in \mathcal{H}_α^t .
- Given two ecologically rational agents α and β , suppose $\mathcal{H}_\alpha^t = \mathcal{H}_\beta^t$ there is no expectation that given the same need, v , in the same context, Θ^t , they will act in the same way even if they are deterministic. E_α encapsulates α 's particular ecological rationality, and E_β β 's. This is in contrast to a pair of deterministic constructivist rational agents that would be expected to perform the same actions if they were both expected-utility-optimisers and had both identified a pure dominant strategy⁶ given \mathcal{H}_α^t .

Example. A classical case-based reasoning agent is a deterministic ecologically rational agent.

The shape of the distribution $E_\alpha(\mathcal{H}_\alpha^t, v, \Theta^t)$ enables us to balance exploration and exploitation. The flatter the distribution the greater the degree of exploration.

An agent that operated as in Equation 2 could have been constructivist if the condition that “ E_α is *not* founded on an abstraction or theory that models, explains, constrains, describes or prescribes the behaviour of agents or the environment” had been omitted — the difference between constructivist and ecological lies in the theory T_α . For example, without this condition an agent that operated as in Equation 2 could place expected utility optimising bids in an English auction. It is perhaps more correct to think of E_α as “interpreting the wisdom” in \mathcal{H}_α^t ; just as C_α interprets the wisdom in: T_α , \mathcal{H}_α^t and D_α^t .

⁶ In which case we assume that they will have made assumptions about the behaviour of other agents for example that they are expected-utility-optimisers as well.

The “ecological rationality” in E_α is based on the belief that the wisdom in \mathcal{H}_α^t can somehow tell α how to act rationally. This belief is not an abstraction or “theory” (as described above); it is simply a belief that the wisdom embedded in prior observations are a basis for rational action. In a simple form, this belief may be that prior agent behaviour reliably indicates future behaviour⁷. That is, ecological rationality may be founded on *a sense of trust* that α has in agents, i.e. that they will continue to behave with no less integrity than that which they have displayed in the past. Ecological rationality may also be founded on the reputation that another agent has, or on trust in the institution (i.e. a normative multiagent system to which all the agents belong) to ensure that agents behave in some way. In addition, Ecological rationality may be founded on subtle observations mined from \mathcal{H}_α^t . As a simple example, “Whenever John is running a marketing promotion Carles invariably gives excellent service”. In general, ecologically rational behaviour will be derived from \mathcal{H}_α^t using data mining techniques. Ecological rationality, like trust, is an experience-based phenomenon — it can not exist without experience. At the “top end”, ecological rationality embodies all the models that have been described from information-based agents including the important notion of integrity⁸.

3.3 Selecting a Rational Basis for Deliberation

We now address the question of when both rationalities may be useful which should be chosen? That is, α will require a mechanism, A_α , to determine which form of deliberation, r , to select in a given circumstance:

$$r = A_\alpha(\mathcal{H}_\alpha^t, v, \Theta^t) \in \{C_\alpha, E_\alpha\} \quad (3)$$

What really matters is an agent’s *ex post* rating of a dialogue. E.g. an agent may have used game theory to obtain a shatteringly good deal only to hear that the person he was negotiating with (also a good friend) was sacked for agreeing to such good a deal. This may cause the agent not to rate the dialogue highly. One approach is to rate all dialogues and to use the method in Section 4 to select a dialogue and then: if the closest high-rated dialogue in the history to it was generated by a Cartesian method then re-use that method as long as the data is available otherwise use Section 4 to drive ecological deliberation [3].

4 A Basis for Rational Deliberation

An ecologically rational α has only the dialogues in \mathcal{H}_α to work with, and, if α is concerned with ratings, maybe only the rated dialogues. No two dialogues are identical — they will differ at least by the time at which they occurred — at best they will be ‘similar’. This does not necessarily mean that α ’s beliefs will be uncertain about how

⁷ Such a belief may incorporate context. For example, “John is most reliable *except* on Mondays”.

⁸ The extent to which a partner agent will take advantage of his private information when enacting his commitments. E.g. “I haven’t got the strawberries you ordered because they were all rain damaged.”

another agent will behave: “Carles always looks after me — no question about it”. But if the history lacks consistency in a sequence of the dialogue ratings then it may well mean that α will be uncertain about how another agent will behave. One way to approach this uncertainty is to measure the similarity between a current, open dialogue, (d, v, Θ^t) , and a rated dialogue $(d', v', \Theta^t, r') \in \mathcal{H}_\alpha$. Such a measure, $\text{Sim}((d, v, \Theta^t), (d', v', \Theta^t, r'))$, may be useful in forming a necessarily uncertain belief about the eventual rating of the open dialogue. If the ratings in the history are totally chaotic then it may mean that the best that an ecologically rational agent can do is to act at random — i.e. E_α will be the uniform, maximum entropy distribution.

Devising a satisfactory definition for $\text{Sim}(\cdot, \cdot)$ is not simple. It would be very convenient if such a definition could be found so that:

$$\mathbb{P}(\text{rating}(d, v, \Theta^t) = r) = \frac{\sum_{d' \in \mathcal{H}_\alpha^t} (\text{Sim}((d, v, \Theta^t), (d', v', \Theta^t, r')) \mid r' = r)}{\sum_{d' \in \mathcal{H}_\alpha^t} \text{Sim}((d, v, \Theta^t), (d', v', \Theta^t, r'))}$$

for any open dialogue d , v and Θ^t . $\mathbb{P}(\text{rating}(d, v, \Theta^t) = r)$ is the expectation of the rating of the open dialogue (d, v, Θ^t) . The following assumes that $\mathbb{P}(\text{rating}(d, v, \Theta^t) = r)$ may be estimated somehow. \mathcal{H}_α^t can be pruned⁹.

If the rating scale has a linear order it is then a simple matter to estimate:

$$\mathbb{P}(\text{rating}(d, v, \Theta^t) > \text{rating}(d', v', \Theta^t))$$

and in particular:

$$\mathbb{P}(\text{rating}(d, v, \Theta^t) > \text{rating}(d', v, \Theta^t))$$

Then in the simple case that we can select an entire dialogue “up front”¹⁰ given v and Θ^t , one simple strategy is for E_α to select dialogue d with a probability:

$$\mathbb{P}(\forall_{d'} (\text{rating}(d, v, \Theta^t) > \text{rating}(d', v, \Theta^t)))$$

that can be estimated approximately from the above with suitable dependence assumptions. In natural language: “we select a dialogue with a probability equal to the probability that it is the best choice”. The dialogue selected by this method is not necessarily in \mathcal{H}_α^t .

⁹ Pruning \mathcal{H}_α^t . Given a newly triggered need, v , and prevailing context, Θ^t , then let the candidate dialogues for (v, Θ^t) be:

$$\text{Can}_\alpha(v, \Theta^t) = \{d_i \in \mathcal{H}_\alpha^t \mid \text{Sim}((v, \Theta^t), (v_i, \Theta_i^t)) < \sigma(v, \Theta^t)\}$$

where $\text{Sim}(\cdot)$ is a similarity function and $\sigma(v, \Theta^t)$ is a threshold constant such that:

$$\|\text{Can}_\alpha(v, \Theta^t)\| > n \text{ and } \|\{d_i \in \text{Can}_\alpha(v, \Theta^t) \mid r_i \geq \text{‘excellent’}\}\| > m$$

where n and m are chosen constants. That is, $\text{Can}_\alpha(v, \Theta^t)$ contains at least n nearest neighbours of (v, Θ^t) and contains at least m with an ‘excellent’ rating¹⁰. The idea then is to use $\text{Can}_\alpha(v, \Theta^t)$ instead of \mathcal{H}_α^t .

¹¹ This is not unrealistic. The dialogue could be “Please send 1kg tomatoes”, “1kg tomatoes delivered”, “Please send €1”, “Completed”

Now consider a more general situation, given an open dialogue d , α wishes to know what utterance to communicate next, where an utterance, u consists of an illocution, m , sent to an agent, β , $u = (m, \beta)$. Let $d \oplus u$ denote the dialogue consisting of d followed by u . Then following the reasoning above we can estimate:

$$\mathbb{P}(\forall_{u'} (\text{rating}(d \oplus u, v, \Theta^t) > \text{rating}(d \oplus u', v, \Theta^t)))$$

that is the probability for any u that it will be the highest rated choice. Then, as above, we select u with a probability equal to the probability that it is the highest rated choice.

We clarify what it means to “re-use” a dialogue. For Cartesian dialogues we simply reuse the theory and acquire current data to feed into it. For ecological dialogues things are more subtle. No two dialogues are the same, but when we say to our butcher “I need 2kg of prime beef” we are probably reusing an established “dialogue type”, or *pattern*, with that agent (the butcher). In this case the pattern could be something like: “state your need and see what is proposed”. The fact that I may never have purchased 2kg of prime beef previously does not mean that I am not reusing something. An abstraction of dialogues to patterns is required to deal with this. Now given that we understand the sense in which an Cartesian or ecological dialogue is being reused, it is reasonable to suggest that we would be able to estimate the prior distribution $\text{rating}(u, v, \Theta^t)$ for initial utterance u . For example, u could be “I need 2kg of prime beef”. It is then a matter of monitoring how the dialogue, d , is proceeding so that we can update the prior estimate of the distribution to obtain $\text{rating}(d, v, \Theta^t)$. The abstraction process is tricky but quite intuitive.

The above argument “hides a multitude of sins” by assuming that:

$$\mathbb{P}(\text{rating}(d, v, \Theta^t) = r)$$

can be estimated for any d . In particular it enables us to “pluck dialogues out of thin air”. The problem that is neatly hidden is: where does \mathcal{H}_α come from?

For human agents, \mathcal{H}_α is built up under guidance from parents, school teachers and from their own experience as they as young adults gradually interact with the world in richer and richer ways. Computerised agents are orphans, and can only initially rely on the ideas of their creator; these may well include theories such as game theory. Cartesian deliberation generates dialogues that may be rated and will be represented in \mathcal{H}_α^t , but they may not provide suitable fodder for ecological deliberation because the rationality in former Cartesian dialogues will be grounded in the particular theory chosen, T_α , and not in the agent’s experiential history.

For computerised agents the challenge for (ecological) deliberation is to strategically shape, develop and maintain its experience base so that it contains sufficient wisdom for the (unknown) future. This is what smart business men do. It is what smart social operators do. Computerised agents need to be able to “observe, understand, evaluate, imitate, remember and recall”.

5 Discussion: What Is an *Intelligent* Rational Agent?

It seems appropriate to end this discussion on rationality by considering the question “What is an *intelligent* rational agent?”. Intelligence is a comparative and not an absolute concept. The question “Is α intelligent?” is meaningless, whereas “Is α more

intelligent than β ?" may make sense. Intelligence has nothing to do with knowing scientific theorems or methods that are reported in the public domain. If a software agent was endowed with every known scientific theorem then that in itself would not make the agent intelligent. *Intelligence* is all about the following four questions:

- information about the world and about other agents — how do you source it, how do you rate its integrity in time?
- how to build trusted relationships with other agents that enable you to benefit from the present and, equally important, to cope with the unknown future?
- how to design and manage an agent's memory to benefit from the present and, equally important, to cope with the unknown future?
- how to select a basis for rational action from all the stuff an agent knows — in plain language, how to select a method to achieve a particular goal?

What makes one agent more intelligent than another is the way it does these four things.

References

1. Arcos, J.L., Esteva, M., Noriega, P., Rodríguez, J.A., Sierra, C.: Environment engineering for multiagent systems. *Journal on Engineering Applications of Artificial Intelligence* 18 (2005)
2. Brooks, R.A.: Intelligence without reason. In: Myopoulos, R., Reiter, J. (eds.) *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, Sydney, Australia, pp. 569–595. Morgan Kaufmann, San Francisco (August 1991)
3. Debenham, J., Sierra, C.: Dual rationality and deliberative agents. In: Bramer, M., Ellis, R., Petridis, M. (eds.) *Research and Development in Intelligent Systems XXVI*, pp. 79–92. Springer, Heidelberg (2009)
4. Epstein, S.: Integration of the cognitive and the psychodynamic unconscious. *American Psychologist* 49(8), 709–724 (1994)
5. Freud, S.: *Formulation of the two principles of mental functioning*, Hogarth, London (1958) (originally published in 1911)
6. Gigerenzer, G., Selten, R. (eds.): *Bounded rationality — the adaptive toolbox*. MIT Press, Cambridge (2002)
7. Hayek, F.A.: *The fatal conceit: The errors of socialism*. University of Chicago Press, Chicago (1991)
8. Levine, D.S.: Brain mechanisms for making, breaking, and changing rules. In: Huang, D.-S., Wunsch II, D.C., Levine, D.S., Jo, K.-H. (eds.) *ICIC 2008*. CCIS, vol. 15, pp. 345–355. Springer, Heidelberg (2008)
9. Rubinstein, A.: *Modeling bounded rationality*. MIT Press, Cambridge (1998)
10. Russell, S., Norvig, P.: *Artificial intelligence: A modern approach*, 2nd edn. Prentice Hall, Englewood Cliffs (2002)
11. Russell, S.J.: Rationality and intelligence. *Artificial Intelligence* 94(1-2), 57–77 (1997)
12. Sierra, C., Debenham, J.: Trust and honour in information-based agency. In: Stone, P., Weiss, G. (eds.) *Proceedings Fifth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2006*, Hakodate, Japan, pp. 1225–1232. ACM Press, New York (2006)
13. Sierra, C., Debenham, J.: Information-based reputation. In: Paolucci, M. (ed.) *First International Conference on Reputation: Theory and Technology (ICORE 2009)*, Gargonza, Italy, pp. 5–19 (2009)
14. Smith, V.L.: *Rationality in economics: Constructivist and ecological forms*. Cambridge University Press, Cambridge (2007)

Involving the Human User in the Control Architecture of an Autonomous Agent

Roberto Micalizio, Giancarlo Nuzzolo, Enrico Scala, and Pietro Torasso

Dipartimento di Informatica, Università di Torino, Torino, Italy
{micalizio,nuzzolo,scala,torasso}@di.unito.it

Abstract. The paper presents an architecture for an autonomous robotic agent, which carries on a plan in a partially observable environment. A Supervisor module is in charge of assuring the correct execution of the plan, possibly by inferring alternative recovery plans when unexpected contingencies occur. In the present paper we describe a control strategy where a human user is directly involved in the control loop, and plays the role of *advisor* by helping the robotic agent both for reducing ambiguity in the robot's observations, and for selecting the preferred recovery plan.

Keywords: Autonomous Agent, Situation Awareness, Plan Execution, Replanning.

1 Introduction

The adoption of robotic agents is becoming critical in many humans' activities ranging from very challenging scenarios of space exploration and disaster rescue to various kinds of service robotics. However, the effective use of service robots in partially observable environments requires that the robots must be (at least partially) autonomous since the execution of a plan must take into account that actions may deviate from their expected effects as a consequence of unexpected contingencies (e.g., unpredictable changes), or when the actual state of the environment is different from the one originally expected.

Anticipating all the possible contingencies by synthesizing conditional or conformant plans is often unmanageable (or even impossible) as the domain complexity prevents to solve the planning problem efficiently. As a consequence, the robotic agent must be able to recognize unexpected contingencies while it is performing its own actions, and properly react to them (e.g., by changing its plan/goals or by interrupting the plan execution when a global failure is detected).

Recently a number of works have faced the problem of plan execution in partially known environments by proposing different control architectures (e.g. [1,2]).

Other works, such as [3,4], have focused the attention on plan execution monitoring and diagnosis as fundamental activities to detect and explain failures. More precisely, [4] proposes a closed loop of control involving, besides action monitoring and diagnosis, a local re-planner that can be invoked to recover from action failures by synthesizing alternative plans on-the-fly.

While these approaches have been conceived in order to endow the robotic agent with as much autonomy as possible, in several domains (such as rescue, space exploration and UAV control) there is evidence that a more robust behavior can be obtained by having a human user involved in the control loop and allowing the system to operate at different levels of autonomy [5,6].

In the simplest view, the agent can be either fully autonomous, or teleoperated; however there are other intermediate cases where the human can play different roles. In the present paper we consider the case where the human plays the role of *advisor* by solving on-the-fly problems arising during the plan execution that the agent is unable to solve autonomously. Humans, in fact, can help the agent in interpretative activities (e.g., object recognition, situation assessment, etc.), that still are not effectively solved by the current generation of robotic agents. However, in order to support the human intervention during plan execution, the control architecture of the robotic agent must be significantly revised. First of all, the robotic agent must be able to interact with the human user by sending her/him help requests and by understanding human's commands. More important, the robotic agent must guarantee the human with a sufficient level of *situational awareness* [7] by providing relevant pieces of information concerning not only its internal status (e.g., position, health status, etc.), but also that portion of the environment relevant for solving the problem.

2 Control Architecture

The system architecture proposed in this paper, shown in Figure 1, has been inspired by the three tiers architecture proposed by Ghallab et al. [1].

While our Functional Layer (FL), the lowest tier, is similar to the one discussed by Ghallab, the main differences w.r.t. the Ghallab's solution concern the Decision Layer (DL) and the Executive Layer (EL). In the Ghallab's architecture, in fact, the role of the human is limited in submitting goals and retrieving results when they have been achieved; on the contrary, in our proposal the Human User (HU) plays an active role during the plan execution by supporting the DL in its deliberative processes. For this reason we added a Human User Layer (HU-Layer) on the top of the DL as shown in Figure 1. Of course, the existence

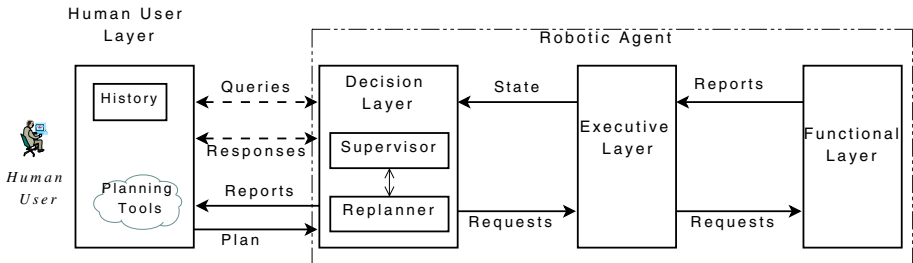


Fig. 1. The extended control architecture

of a HU-Layer forces us to decide when and how the interactions between the HU and the robotic agent RA can occur.

While the proposed architecture supports a Mixed-Initiative approach [8], where both the HU and the RA can start an interaction phase at any time during the execution (see the dashed arrows in Figure 1), for space reasons in the present paper we will focus just on the System Initiative pattern.

In our architecture the HU has the task of providing the RA with an initial plan P (the control architecture does not make any assumption on the way the plan P has been obtained, either manually or via a set of planning tools). The HU-Layer is an intermediate layer between the Human User and the Decision Layer: on one hand it allows the HU to submit the plan P to the RA, on the other hand accepts query and inform messages coming from the DL, and displays these data to the HU so that she/he can answer the requests. The HU-Layer maintains a history of the data received by the DL (i.e. the behavior of the RA and the changes occurred in the environment over a time interval) for increasing the human's situation awareness.

As soon as the DL receives the plan P , DL starts the plan execution phase. Note that the DL includes two modules: a Supervisor and a Replanner; the former implements the control strategy, the latter revises the plan P when unexpected contingencies occur. The Supervisor has a hard task since it has to deal with incomplete knowledge, partial observability and unexpected contingencies. In fact, the plan is synthesized off-line when the actual status of the system is not completely known and assumptions have to be made, but these assumptions could be invalid when the plan is executed. During the execution, the Supervisor has to assess the actual status of the environment via sensing actions. However, the RA has not in general a sufficient amount of information for precisely assessing the actual status of the world (the environment is only partially observable), so it maintains a *belief state* (namely, a set of alternative system states which are consistent with the available observations), rather than a precise state of the environment. The results of sensing actions are used to refine the belief states (i.e. they reduce ambiguity) by pruning off all those states which are inconsistent with them. However, sensing actions are not always so accurate to report a unique interpretation of the environment state and therefore, the result of a sensing action could be still ambiguous. Whenever the actual status of the environment differs from the expected one, the Supervisor invokes the Replanner to repair (when possible) the plan.

The plan. As said above, the user provides the RA with a plan P . We only require that P is a sequence of actions (i.e. a totally ordered plan) where each action is modeled in terms of preconditions and effects. Figure 2 shows an example of such a plan, where the solid edges between actions models causal links; e.g., the edge from action 9 to action 10 is labeled with the service `ONBOARD(R,B1)` means that such a service is an effect of action 9 and a precondition of action 10; dashed edges between actions model precedence relations.




Fig. 2. A segment of the plan P assigned to the robotic agent

3 A Motivating Example

In this section we introduce the scenario we use throughout the paper to exemplify the proposed architecture and control strategy.

We consider an assembly task where a RA has to build a bridge (consisting of two pillars and a beam) by assembling a number of bricks. Each brick has a *length* (short, medium or long), an *orientation* (horizontal or vertical) and a *status* (normal or flawed); only normal bricks can be used to build the bridge. The RA is initially given with a plan P (see Figure 2) to build the two pillars of the bridge one in position T1 the other in position T2. The available bricks are initially located in two repositories S1 and S2; S1 is the nearest repository to the final location of the bridge and contains a sufficient number of bricks to build the bridge. The RA, positioned in S1, can load and unload bricks, move from a location to another and recognize the size of the bricks. Actually, Figure 2 shows just the segment of the plan when the RA loads the brick B1 in S1 and puts it (in vertical orientation) on the top of B2 at the target location T2.

Since at the planning phase the actual status of the resources (i.e., the bricks) is not known, the planner assumes that all the bricks are *normal*. Thus, the plan P is executable as long as these assumptions match with the actual status of the environment. The RA performs appropriate sensing actions to check the validity of the assumptions stated in the preconditions of the actions to be executed. For example, the load action requires that the block to be loaded must be *normal*, therefore, before loading B1 the RA must be sure that B1 is not damaged. However sensing actions are not always so precise to discriminate between *normal* and *flawed*; and hence the response of the sensing action could be $normal(B1) \vee flawed(B1)$. Thus the RA builds an ambiguous belief state where both values are allowed.

When the RA has to deal with such a kind of ambiguous data, it could adopt a very conservative policy by considering as “faulty” any object that is not surely “healthy”; thus the brick B1 would be considered as *flawed*. However, this solution may be not acceptable as the RA could waste resources discarding “healthy” objects. The control architecture discussed in this paper mitigates this problem by allowing the RA to ask the HU intervention for reducing the belief state ambiguity; the aim is to exploit the human’s interpretative skills that usually are better than the interpretative services of the robot. Whenever an action precondition is not satisfied, the RA must repair its original plan, for example by replacing the damaged brick with a normal one. Also in this case the human user may intervene in order to approve the changes to the original plan proposed by the RA.

```

PLANEXECUTIONSUPERVISOR(plan P, status S, goal G)
01 while (hasActionToPerform(P)) {
02   a = getNextAction(P);
03   Spre = activateSensors(pre(a));
04   S = updateCurrentState(Spre,S);
05   executable = checkPreconditions(pre(a), S);
06   if (executable equals unknown)
07     executable = queryHU-checkPreconds(pre(a),S);
08   if (executable) {
09     execute(a);
10     markAsExecuted (a, P);
11     S = updateCurrentState(effects(a),S);
12     if (isRelevant(a)) inform-HU(S, a, succeeded);
13   } else {
14     P = RECOVERY(S,G);
15     if (isEmpty(P))
16       inform-HU(S,a, failed);
17     return failure; } }
18 return success;

```

Fig. 3. The main supervisor function

```

RECOVERY(status S, goal G)
01 plansList = REPLANNER(S,G);
02 if isEmpty(plansList)
03   return NULL;
04 choice = 0;
05 if (size(plansList) > 1)
06   choice = queryHU-selectPlan(plansList,S,G);
07 return getPlanAt(plansList, choice);

```

Fig. 4. The RECOVERY function

4 The Supervision Strategy

In this section we focus on the control strategy adopted by the Supervisor, and on the relations (i.e., flows of information and requests) existing between the Supervisor and the Human User.

Figure 3 reports the high-level description of the strategy, which receives as input the plan P to be executed, the initial status of the system resources \mathcal{S} , and the goal G represented as a set of constraints. The status \mathcal{S} initially represents the *nominal* state of the system resources which has been assumed to synthesize P (e.g., in the scenario described above, \mathcal{S} assumes that all the bricks are *normal*).

Since the plan P has been built under incomplete knowledge, the Supervisor is responsible for checking the preconditions of each action $a \in P$ before executing it. Through `activateSensors` command (line 03), the Supervisor acquires new pieces of information about the environment. In particular, relying on the preconditions of the action a , the function determines which set of sensing actions must be performed. Note that only those preconditions that have been assumed during the planning phase need to be sensed. Whereas, services provided by other actions previously executed can be checked just by inspecting the current status \mathcal{S} . As said above, the result of a sensing action may be ambiguous, therefore the response of `activateSensors` command is in general a set \mathcal{S}_{pre} of alternative states.

The new pieces of information acquired by sensing actions are used to update the current system status \mathcal{S} (line 04). Of course, since \mathcal{S}_{pre} may be ambiguous, \mathcal{S} may be ambiguous in its turn. Therefore, when the Supervisor checks the preconditions of action a in the current status \mathcal{S} (line 05), the response could be *unknown*: \mathcal{S} includes both states where the preconditions of a hold and states where they do not.

The Supervisor tries to disambiguate its belief state \mathcal{S} by asking the HU (line 07) to interpret the collected data. To this end, the Supervisor provides the HU with the preconditions of a still to be checked, and the current (ambiguous)

state \mathcal{S} ; thereby the HU can gain the situation awareness necessary to accomplish her/his task. In general, however, we cannot assume that the HU is always able to disambiguate \mathcal{S} , and hence also the HU may reply *unknown* to the Supervisor (in this case the state \mathcal{S} remains ambiguous).

Only when the action preconditions are satisfied for sure (i.e., when the preconditions are satisfied in each state $s \in \mathcal{S}$), the action is executed (line 09), and the status \mathcal{S} is updated with the action effects (line 11).

In principle, the Supervisor should notify the HU with plan progresses by invoking the `inform-HU` primitive after each action execution. However, in order not to overload the HU with action feedbacks, this happens only when the action a is “relevant” w.r.t. some user-defined criteria. A simple but effective strategy we have adopted in our tests consists in considering as relevant an action a iff at least one of its effects is part of the final goal.

When the action preconditions are not satisfied (or it is impossible to precisely evaluate them), the `RECOVERY` module is activated (line 14). The recovery module (whose main steps are showed in Figure 4) is based on a conformant re-planner, which produces a list of alternative plans reaching the goal G given the current state \mathcal{S} .

Of course, when *plansList* is empty no solution exists and the recovery terminates with a failure. In such a case the detected failure cannot be autonomously recovered by the RA, and a human intervention is required. Conversely, when *plansList* contains a number of plans, the Supervisor asks (line 06, Figure 4) to the HU the preferred plan, the selected plan is used by the Supervisor (line 14, Figure 3) as the new plan P the RA will execute.

Re-planner. The conformant re-planner invoked by the Supervisor during its activity is the same we have discussed in [9]; such a planner has some similarities with the Universal Planner by Veloso et al. [10], since it goes forward from the initial state \mathcal{S} to the goal G , and exploits the symbolic formalism of Ordered Binary Decision Diagrams (OBDDs) to encode action models and \mathcal{S} . The main advantage of our implementation based on OBDDs is that we can avoid backtracking by keeping in just one OBDD (in a compact way) all the possible partial plans built so far: at each step the planning algorithm extends every partial plan so that the search goes on in parallel along different candidate solutions. In this way the extension of the plan of one action is realized in just one shot by exploiting the standard OBDDs operators, and there is no need to iterate over the space of the current partial solutions.

As a consequence, when the goal G is achievable, the re-planner finds all the plans achieving the goal G with the minimal number k of actions (i.e., all these plans are optimal in terms of the number of used actions). The resulting plans are subsequently returned to the `RECOVERY` module, which displays them to the HU.

For the specific purpose of our architecture, we have slightly modified the planner in [9] in order to look for alternative solutions in those cases where just one optimal solution exists. However, these further solutions cannot employ more than $k+\delta$ actions to reach the goal, where δ represents the maximum cost (in terms of number of actions) the HU is willing to accept for a non optimal

solution. The idea is to propose to the user a number of solutions (even not necessarily optimal) that helps her/him in taking a decision. In fact, proposing just one plan, even if optimal, may be too restrictive: the HU could consider the choice of that plan just as a decision made by the Supervisor. Conversely, by proposing also not optimal solutions the HU can envisage different scenarios and take a decision based on different biases.

Note that, since the re-planner is involved in an on-line interaction with the HU, its response time should be compatible with the typical times of Human Computer interaction. The preliminary results we have collected¹ suggest that our planning techniques, implemented in JAVA, can meet such a requirement. In fact, when a solution shorter than 10 actions exists, the re-planner takes less than 1 sec. to find it. In more challenging situations, the re-planner finds conformant plans including up to 20 actions in about 30 sec.; note, however, that it is often undesirable having recovery plans longer than 10 actions: when a recovery plan becomes too long the human intervention may represent a more efficient solution.

Communication primitives. Since the HU is involved in the control of the plan execution, the supervision strategy encompasses some specific communicative acts which assure the HU with the right level of situation awareness. The HU, in fact, can effectively intervene only when she/he knows the relevant aspects of the system status (RA and environment). In order to meet these objectives, the supervision strategy exploits two communication primitives *inform-HU* and *queryHU-...*² The purpose of *inform-HU* is to provide, step by step, the HU with some information to let him aware of the current progress of the plan. The HU-Layer incrementally collects in a history the pieces of data received via *inform-HU* messages. The HU can therefore gain his/her situation awareness relying upon a recent history of system states. Another communicative act is the help request, modeled by the *queryHU-* primitive, issued by the RA when failures or ambiguities arise; it consists in sending the current state \mathcal{S} and the key data concerning the impasse (e.g. alternative plans, or cues about ambiguities arising during the check of the preconditions), and in waiting for the HU's answer.

Note that the *queryHU-* primitive supports both "management by consent" and "management by exception": in the former the RA waits indefinitely a human's reply while in the latter the RA waits for a pre-defined time interval after which it decides on the safest possible solution (e.g., when the agent is unsure about the status of a brick, it assumes that the brick is unusable).

5 Running Example

Let us consider again the segment of plan shown in figure 2, and let us assume that the next action to perform is the load action 9. The action preconditions

¹ The experiments have run on a laptop Intel Core2 2.16 GHz, 2 GB RAM.

² These two primitives resemble the "inform" and "query" primitives defined in the FIPA-ACL standard (www.fipa.org).




Fig. 5. The two alternative plans to resolve the impasse $flawed(B1)$

require that the brick B1 must be *normal*, moreover both B1 and the RA must be within the repository S1. It is worth noting that the RA is in S1 for sure as a consequence of the preceding move action 8 (this atom can be checked by considering the current belief state); whereas both the position of the brick and its *normal* status must be checked via some sensing actions, as these two atoms are assumptions made in the initial state.

Let us assume that, as a result of its sensing actions, the RA Supervisor discovers that the brick B1 is actually located in S1, but it is not normal. The Re-planner module finds two possible solutions, (the recovery plans are reported in Figure 5): the first plan, involving the minimum number of actions, consists in finding another brick, B3, whose dimension is exactly the same as the flawed brick B1; however B3 is located in repository S2, which is far away from the target positions T1 and T2. Therefore, the cost of the move action R11 from S1 to S2 may be relevant. However the planner provides a second solution consisting in assembling a number of small bricks for obtaining one of the desired size. Since the flawed brick B1 is *medium* sized, it can be substituted by composing two *small* bricks B4 and B5, both located in the repository S1.

These two solutions are presented to the HU, who can determine the most suitable possibly by considering not only the data provided by the RA, but also the knowledge she/he has about the whole environment where the RA operates, which in general is not completely captured by the RA.

6 Implementing the Supervision Strategy via PLEXIL

To implement the Supervisor’s strategy we have used the PLEXIL declarative language developed by the NASA [11] (see <http://plexil.wiki.sourceforge.net>). PLEXIL is a declarative language which allows to hierarchically decompose complex tasks into sub-tasks. A PLEXIL model is therefore a tree, where a node may be activated by other nodes previously executed or by the occurrence of exogenous events. PLEXIL helps the programmer in the definition of the complex task to solve as it allows a clear distinction between the task decomposition and the actual implementation of the required services.

For example, the tree in Figure 6 summarizes the supervision strategy previously discussed: while the intermediate nodes define the task decomposition,

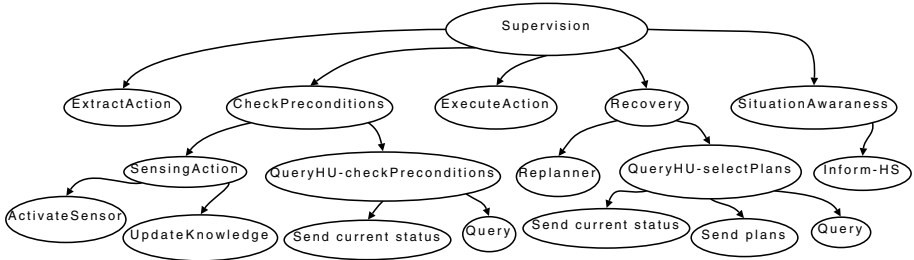


Fig. 6. A graphical representation of the PLEXIL program implementing the Supervisor algorithm

the leaves represent calls to services (such as the Replanner) or to external sub-systems (such as the HU-Layer). For the sake of readability the figure reports just the task decomposition, and does not show the causal dependencies existing among the nodes. At the upper level the Supervision task is decomposed into the following nodes: *ExtractAction*, *CheckPreconditions*, *ExecuteAction*, *Recovery* and *SituationAwareness*.

The *CheckPreconditions* node is activated just after the next action to be performed has been extracted and consists of two sub-activities. The former gathers new pieces of information from the sensors and updates the current belief state. The latter, activated only when the previous sensing action returns ambiguous data, is the request for help toward the Human User. The interaction with the human is quite complex, in fact the node *QueryHU-checkPreconditions* is in turn subdivided into two nodes: *Send current status* which aims at providing the human with the pieces of knowledge relevant to describe the impasse, and *Query* which formulates a question to the human, captures the human's answer, and updates the current belief. Note that, the *Query* node takes also care of the management by consent and by exception we have previously described (i.e., a query node is not necessarily blocking).

The *CheckPrecondition* node activates the *ExecuteAction* node iff the precondition of the next action are satisfied in the belief state resulting after the sensing actions. The *ExecuteAction* node delegates the execution of the action to the EL that, in turns activates the services offered by the FL.

The *Recovery* sub-tree represents all the set of operations the RA performs when the action preconditions are not satisfied for sure. The *Recovery* node invokes first the Replanner algorithm previously discussed, and then the *QueryHU-selectPlan* node representing the help request toward the HU; in particular, such a node is subdivided into three other primitive nodes: the actions for sending the current belief state, the set of possible plans (these pieces of information contextualize the problem) and the action representing the real query waiting for the human response.

Finally, the *SituationAwareness* node is activated whenever relevant data have been collected and the user has to be informed about them.

7 Conclusions

The paper has addressed a control architecture for an autonomous robotic agent, which represents a first step towards the development of a more robust robotic agent. In fact, albeit the robotic agent is endowed with some form of autonomous behavior (which enables the agent to face some unexpected contingencies), the agent can request the help of the human user whenever the human, playing the role of *advisor*, can solve the impasses more effectively.

In this first attempt, we have shown a simple form of cooperation between the human user and the agent. On one hand, the computational power of the agent is used to present to the human a range of possible (conformant) solutions for an impasse. On the other hand, the human has the ability to assess them and to select the best one also by exploiting her personal experience without having to compute by himself/herself the conformant plans.

So far the interaction primitives supporting the human's situation awareness are system initiative based; however, we are currently extending them in order to develop more complex forms of human-agent cooperation including mixed initiative interactions (i.e., where the human could intervene without an explicit help request from the agent).

References

1. Alami, R., Chatila, R., Fleury, S., Ghallab, M., Ingrand, F.: An architecture for autonomy. *International Journal of Robotics Research* 17(4), 315–337 (1998)
2. Galindo, C., Fernandez-Madrigal, J., Gonzalez, J., Saffiotti, A.: Robot task planning using semantic maps. *Robotics and Autonomous Systems* 56, 955–966 (2008)
3. de Jonge, F., Roos, N., Witteveen, C.: Primary and secondary diagnosis of multi-agent plan execution. *Journal of Autonomous Agent and MAS* 18, 267–294 (2009)
4. Micalizio, R., Torasso, P.: Monitoring the execution of a multi-agent plan: Dealing with partial observability. In: *Proc. of ECAI 2008*, pp. 408–412 (2008)
5. Sellner, B., Simmons, R., Singh, S.: User modeling for principled sliding autonomy in human-robot teams. *Multi-Robot Systems. From Swarms to Intelligent Automata* 3, 197–208 (2005)
6. Fong, T., Thorpe, C., Baur, C.: Collaboration, dialogue, and human-robot interaction. In: *Proc. 10th International Symposium on Robotics Research* (2002)
7. Endsley, M.: Towards a theory of situation awareness in dynamic systems. *Human Factors* 37(1), 32–64
8. Bresina, J.L., Morris, P.H.: Mixed-initiative planning in space mission operations. *AI Magazine* 28(2), 75–88 (2007)
9. Micalizio, R., Torasso, P.: Recovery from plan failures in partially observable environments. In: *Research and Development in Intelligent Systems XXVII*, pp. 321–334 (2007)
10. Jensen, R.M., Veloso, M.M.: Obdd-based universal planning for synchronized agents in non-deterministic domains. *Journal of Artificial Intelligence Research* 13, 189–226 (2000)
11. Verma, V., Estlin, T., Jonsson, A., Pasareanu, C., Simmons, R., Tso, K.: Plan execution interchange language (plexil) for executable plan and command sequence. In: *i-SAIRAS 2005* (2005)

PART II
Knowledge Acquisition

Exer-Learning Games: Transferring Hopscotch from the Schoolyard to the Classroom

Martina Lucht, Steffi Domagk, and Martin Mohring

Fraunhofer Institute for Digital Media Technology IDMT,
Hirschlachufer 7, 99084 Erfurt, Germany
Martina.Lucht@idmt.fraunhofer.de

Abstract. In this paper we present the idea of Exer-learning games - integrating the element of exercise into serious games. The combination of motion, learning and playing is assumed to facilitate intrinsic motivation and learning achievements. For the application of exer-learning games, the concept of HOPSCOTCH is introduced that is inspired by the popular childrens game. Two demonstrators of this concept have been realized and evaluated: HOPSCOTCH*pad* and HOPSCOTCH*mobile*. First results show a positive feedback from scholars and teachers. Finally, future directions for our research on HOPSCOTCH are described that could model research on exer-learning games as well as their application in the classroom.

Keywords: digital game-based learning, serious games, exercise, learning, playing, exer-learning games, HOPSCOTCH, mobile learning.

1 Introduction

Learning and playing have long been tried to connect in order to use the intrinsic motivation of playing for persistent and joyful knowledge acquisition. In this paper we discuss the aspect of learning while playing, including a positive effect of integrating exercise in the learning process, as part of the play action. We therefore introduce a new genre of digital learning games that combine exercise, learning and playing: exer-learning games. The assumed positive effects of exer-learning games refer to the endurance of the willingness to learn and the quality of knowledge acquisition, especially the accessibility of knowledge. The popular children's game "Hopscotch" has been adapted to train English vocabulary while jumping on a sensor pad. This new concept for excer-learning games has been developed at the Fraunhofer Institute of Digital Media Technology (IDMT) and aims at the joyful and physically active acquisition of factual knowledge. First evaluations of two developed demonstrators "HOPSCOTCH*pad*" and "HOPSCOTCH*mobile*" are presented in this paper along with subsequent future directions for research and the integration of Excer-learning games in school lessons. We argue that Excer-learning games may facilitate knowledge acquisition and intrinsic motivation to learn as well as provide new approaches for teaching and practicing in classrooms.

2 Exercise Learning vs. Conventional Learning

Children’s play and learning are interconnected. “Children explore and acculturate the world through play, extend their skills and competencies, and experiment with possible selves. Only at a later point during elementary school do entertainment and learning start to drift apart. Older children may even associate play with being noneducational and learning with being anything but enjoyable” [23]. Formal learning seems mostly to be associated with sitting at a desk. This applies to school, in which the children spend most of the time sitting still, and to homework that is also mostly done in a sitting posture. Assuming a natural urge to move, especially among children of primary school age, it is liable that the necessary suppression of movement during learning can lead to a negative connotation.

2.1 Exercise

Unlike earlier, when students have physically worked after school or played outside, in the media age also recreation has become rather motionless. The German children and adolescents survey (KiGGS) notes that the “[...] chance on physical and sporting inactivity in the age between 11 to 17 years increases with each passing year by about 30% on average” [16]. Studies show that young people spend more and more free time with media reception instead of playing outside (e.g. [31][30]). This trend has enormous implications for the mental, physical and social development of children and adolescents, because exercise affects not only physical health. “In addition to positive effects on the organic and motoric development the importance of the psychosocial well-being, personal development and learning of social skills is to emphasize” [16]. Therefore the lack of exercise is not only problematic for the physical health of young people, it also affects psychological and social aspects.

2.2 Exergames

In the End of 2006 Nintendo presented the Wii in Europe and increased the popularity of Exergames. The so called Exercise Games (Exergames) are games where the input is done with body movements rather than with the finger. The body movements are targeted with special input devices that work with accelerometers. The players do not sit in front of a Computer but “bowl”, “box” or “hit” with the input device. These body movements are detected sensor based and are integrated in the game on the monitor. Sensor pads were used for dance games, where the player must step on a field with an arrow. The goal of the dance game is to hit the correct arrow direction at the correct point of time, given on the monitor. Most popular are “Dance Dance Revolution”, “In The Groove” and “StepMania”. In the field of games research, Exergames were studied mainly in terms of their health effects [34][19]. The health aspect of exergames is obvious, but could the combination of games, movement and learning content in digital games offer a new dimension of knowledge acquisition?

2.3 Digital Game-Based Learning

The concepts “serious games”, “game-based learning” and “digital educational games” - that are used widely synonymous - mark the initiative to use the potential of digital games to actively engage players for learning (e.g. [2,8,23,21]). They aim at the combination of joyful playing and learning that is defined as intentional acquisition of knowledge through practice and training. The primary goal of game-based learning however, is not the fun-aspect but the benefits for learning namely an increase of intrinsic motivation and knowledge acquisition [9,3]. Intrinsic motivation arises from activity-specific incentives that are inherent in games but not necessary in learning activities [35]. Learning activities are often performed to reach certain consequences (good marks, exams, etc.) and therefore have to be conceptualized as extrinsically motivated [22]. Game-based learning tries to create intrinsic motivation for learning through connected game activities that offer activity-specific incentives.

Recent game-based learning initiatives focus on deeper learning rather than the acquisition of factual knowledge and on broader educational issues outside the classroom [8,7,15,17,23]. “The current focus on the research and development of video games and simulations for learning is not translating into widespread practice in classrooms” [26]. However, first attempts have been made to integrate digital game-based learning into school lessons [27,11].

The combination of learning and gaming is a promising attempt to create new enjoyable ways of knowledge acquisition. The aspect of exercise however, has not yet been integrated into this discussion. We therefore propose a new genre of learning games: Excer-learning games.

2.4 Exer-Learning Games

The educational aspect of Exergames with regard to the acquisition of knowledge seems to be a desideratum. The advantage of a combination between learning and movement, however, can be concluded from different interdisciplinary scientific discoveries. Theories of exercise- and sports education for example, describe that development and differentiation are connected with psychomotor behaviour patterns, cognitive and personal-psychological elements [38,1]. Thus, exercise learning can be seen in context of a structuring process of the brain and also a process in which the learner’s self-confidence can be strengthened [36]. In addition, medical studies show that exercise leads to an increase in concentration [24].

Studies in pedagogics and psychology have shown that even minimal physical activities can support the learning process because actively performed tasks are memorized better than passively received information (e.g. [5,4,10,37]. In the current scientific discussion these outcomes are picked up. Tabbers [33] showed that an animation is better memorized when the learner can reproduce the movements of the animation on the screen with the mouse instead of just watching passively. These outcomes correspond with new findings in the neurosciences. These show that the human language- and exercise system in the brain are closely related. Physical activities may lead to additional links that support an

easier recall of the information learned (e.g. [12,28]). Critically it is to consider that the extensive involvement of various movements during the learning process, could increase cognitive load, thus exceed working memory capacity and even hinder learning. With increasing practice however, the task performance (performing the required movements) should become automated - not needing additional working memory resources [25]. Furthermore, cognitive load theory [32] considers learning in terms of efficiency rather than effectiveness: learning faster, without mental stress [13], whereas Excer-learning games consider learning in terms of effectiveness: investing more learning time due to a higher intrinsic motivation (having fun while learning) and challenge as one important incentive of game play.

We assume that Excer-learning games lead to better knowledge acquisition than traditional learning strategies. This assumption is justified by different theoretical approaches. The link between exercise and learning can fulfil a mnemonic function by the movement (e.g. dance steps) as a kind of crib [29]. In addition Piaget [20] states that the evolution of morphogenetic regulation is an important component of psychomotor and cognitive development.

3 HOPSCOTCH – A Concept for Excer-Learning Games

Hopscotch is a children’s game where numbered squares are to be hopped in the sequence of their numbers. The course is drawn with chalk on pavement. It is one of the world’s most popular outdoor games for children, often played in the schoolyard. The children’s game “Hopscotch” inspired the idea of HOPSCOTCH, presented in this chapter. The concept of HOPSCOTCH was developed at the Fraunhofer Institute for Digital Media Technology (IDMT), Germany. On the basis of software, scholars are motivated to solve various tasks by jumping on fields. The fields are reminiscent of the keyboard of a mobile phone, they are multidimensional: Tipping on a field once, twice or three times gives a different letter (e.g. A, B, C), comparable to writing an SMS. The tasks are shown on a monitor. The playing component of HOPSCOTCH is to move the body as quick as possible on the field, touching the correct fields in a given order as fast as possible. The learning component is the event that starts the game: What fields are to be touched in what order.

In serious games the learning aspects are often less enjoyable add-ons to the game [23]. In HOPSCOTCH the learning aspect is fully integrated into the game, since it is the event that starts the actual game process.

3.1 Two Applications of the HOPSCOTCH Concept

On the basis of the described concept, two applications were realized: HOPSCOTCH*pad* and HOPSCOTCH*mobile*.

Realisation on a Sensor Pad (HOPSCOTCH*pad*). For scientific research on Exer-learning games a first demonstrator of HOPSCOTCH was developed,




Fig. 1. Playing field of the childrens game “Hopscotch” (left), input device for HOPSCOTCH (right)

using a dance pad with nine sensor fields as an input device. As learning content an English vocabulary trainer was programmed on basis of “Unity” game engine. Via monitor the player is asked to translate a German word into English and enter the solution on the dance mat. The integrated high score takes the time that was used to “write” the correct answer. Incorrect letters are ignored, only when the correct letter was entered, the next can be jumped. The correct input is commented by a positive feedback with a short animation as it is usual in computer games (immediate feedback). HOPSCOTCH_{pad} can be played in single- and multipayer modus.

Realisation on a Mobile Phone (HOPSCOTCH_{mobile}). Studies confirm that mobile devices are very attractive to young people [30,31]. 95% of young people own a mobile phone and use it mostly for SMS and for playing. Additionally, the rapidly progressing technical development, also in the area of the mobile devices, opens a huge number of new possibilities for mobile learning, gaming and entertainment.

Because of this IDMT ported the gaming concept of HOPSCOTCH to mobile devices. After creating different interaction concepts one was implemented and evaluated. The focus lay on the fact that a mobile version of the game should run on many devices. During the implementation the concept was created on HTC Touch Diamond2 programmed on “C-Sharp”. This Version was ported on Sony Ericsson W995i based on “Java ME”.

The aspect of exercise was realized by using an accelerometer (g-sensor). This technology allows recognizing the exercises (in this case the jumps) of the player. The rotation around the axes (the pitch) of the mobile phone is used to determine the jumping-direction. For example: if the player rotates the mobile device to the back and jumps up, the players’ character moves one step back. Similar to the sensor pad version, wrong letters are ignored while only right letters are entered.

Because of the small display of mobile phones additional feedback is very important. Different sounds for any field and the use of vibration generates haptical and acoustical feedback. Correct spellings are rewarded with short animations.



Fig. 2. HOPSCOTCHpad (left) and HOPSCOTCHmobile (right)

3.2 Evaluation of the HOPSCOTCH Applications

The two prototypes first were evaluated in terms of acceptance and fun. This was to investigate whether the theoretical assumed positive effects were actually met by the integration of exercise into the learning process.

The sensor pad version of HOPSCOTCH was evaluated on a children’s trade show in Erfurt, a survey of scholars ($N = 57$) aged between eight and twenty years ($M = 11.5$, $SD = 1.8$). The sample consisted of young visitors to the fair, who self-chosen played HOPSCOTCH*pad* and school classes who were invited to hold English lessons with HOPSCOTCH*pad*. Children and adolescents who played at least 10 minutes were invited to answer a questionnaire, the exact play time was not measured. 78.9% ($n = 45$) of the respondents were female and 21.1% male ($n = 12$). Additionally the accompanying teachers were asked to fill out an extra questionnaire ($n = 11$).

The mobile version of HOPSCOTCH was evaluated in a child and youth centre in Gotha, evaluating twenty three children aged between eight and twenty one years ($M = 13.13$, $SD = 3.84$, $N = 23$). 43.5% ($n = 10$) of the probands were female and 56.5% ($n = 13$) were male. After a short instruction how to play, the children played at least five minutes followed by a questionnaire.

The Opinion of the Pupils. On a scale from 1 (“very good”) to 5 (“poor”) the probands were asked to rate their first impression of the game. 96.5% stated as their first impression as “good” (21.1%) or “very good” (75.4%) for HOPSCOTCH*pad*. When the scholars had to decide whether it is more of a sports or an educational game, the vast majority (77.2%) associated HOPSCOTCH*pad* as a learning game. However the element of exercise seems to be important: 48.2% of the players want to exercise with HOPSCOTCH*pad*. 62.5% of the respondents would use it for learning. The majority prefers to play together with friends (80.4%) instead of playing alone (1.8%). In general there was a very positive feedback to HOPSCOTCH*pad*; children came again and again to play. The fact that 56.1% of the interviewed would renounce on another present for getting HOPSCOTCH*pad* underlines this impression.

The first impression of HOPSCOTCH*mobile* was “very good” (65.2%) or “good” (34.8%) on a scale from 1 (“very good”) to 5 (“poor”). A second indicator for the attractiveness of HOPSCOTCH*mobile* gives the factor “play fun” that was also rated as very good to well ($M = 1.41$; $SD = 0.60$) on scale from 1 (“very good”) to 5 (“poor”). The questionnaire also measured the attractiveness of the new game control. Although the combination of pitching and jumping seems to be a little bit tricky (because the players have to direct the phone in a correct position while jumping) the probands rated the game control as very good ($M = 1.41$; $SD = 0.47$).

The Opinion of the Teachers. The teachers ($n = 11$) also rated HOPSCOTCH*pad* very positive. Their first impression of HOPSCOTCH*pad* was for 72.7% “very good” and for 27.3% “good”, rated again on a scale from 1 (“very good”) to 5 (“poor”). From their point of view, the game has “high” (81.8%) or “very high” (18.2%) potential for learning. Only 18.2% see the potential at an average, none of them see a low potential. The teachers could very well imagine using HOPSCOTCH for their classes: 72.7% see a “very high” or “high” potential for HOPSCOTCH as a teaching aid.

The mobile version of HOPSCOTCH was not yet presented to the teachers.

3.3 Interim Conclusion

The pedagogic concept of HOPSCOTCH was applied in two demonstrators: HOPSCOTCH*pad* and HOPSCOTCH*mobile*. Both got positive feedback from children. They love playing HOPSCOTCH even though they classify it as a learning game. The teachers also see high potential in this concept and could very well imagine using it as teaching aid in school.

4 Next Steps

Initial tests show positive feedback from students and teachers for HOPSCOTCH as a learning concept. The research interest now comprises two aspects (1) conducting empirical studies on the effectiveness of Excer-learning games based on the HOPSCOTCH concept and (2) bringing HOPSCOTCH into the classroom as a teaching aid.

4.1 Specifications for HOPSCOTCH as a Teaching Aid

To what extent is it possible to integrate HOPSCOTCH into school lessons and how could this happen? To answer this question a qualitative survey with teachers will be realized. It aims at specifications for a school version of HOPSCOTCH. As Simpson and Stansberry [26] state, it is important for the teachers to anchor innovations to what they are familiar with. Therefore initial semi-structured interviews will be organised, in which the conditions of school life, the learning standards, the lesson planning, etc. will be assessed, to find the possibilities of using HOPSCOTCH in classes and the opportunities and risks that occur.

While this first survey with interviews is focusing on the conscious schemes and opinions of the teachers, a second survey will also evaluate the unconscious opinions. By using the repertory grid method [14,6,18] the unknowingly constructs of HOPSCOTCH in school lessons will be made available. The major goal of this second survey is to evaluate the inner acceptance of the teachers to a refinement of their classes in general and to a refinement in terms of exercise in special.

4.2 Empirical Studies on the Effectiveness of Exer-Learning Games Based on HOPSCOTCH

Do children actually gain more factual knowledge when playing HOPSCOTCH than with traditional learning strategies? The assumed positive effect of Excer-learning games on intrinsic motivation and learning will be investigated in experimental studies. Therefore an Excer-learning game for English vocabulary will be implemented on HOPSCOTCH*pad* and compared to traditional learning strategies for the same set of English vocabularies.

4.3 Conclusions

In this paper we introduced the idea of Excer-learning games. As outlined it seems to be promising to integrate basic elements of serious games like competition and highscores and combine them with the element of exercise that goes beyond input devices like mouse or keyboard. Two applications of the HOPSCOTCH concept were realized, one with a sensor pad technology and one on a mobile phone. First results with these demonstrators show positive feedback from children and teachers. Exercise could be a catalyst that integrates the intrinsic motivation of playing into the learning process in a better way than traditional serious games. Further projects focus on the integration of this concept into school lessons. In sum, HOPSCOTCH could well be a new inspiration to the learning process at home and in school classes.

Acknowledgements

The present research has been partially funded by the Federal Ministry of Education and Research within the project “MEMORI” under contract 01GI0902 and the Thuringian Ministry of Education, Sciences and Culture within the project “iCycle” under contract PE 004-2-1.

References

1. Bietz, J., Laging, R., Roscher, M.: *Bildungstheoretische Grundlagen der Bewegungs- und Sportpädagogik*. Schneider Verlag, Baltmannsweiler (2005)
2. Brannigan, C., Owen, A.: Game based learning or learning based games? a case study. In: Burmester, M. (ed.) *Digital Game Based Learning - Proceedings of the 4th International Symposium for Information Design 2nd of June 2005 at Stuttgart Media University*, pp. 179–196. Universitaetsverlag, Karlsruhe (2006)

3. Dumbleton, T., Kirriemuir, J.: Understanding digital games. Sage Publ., London (2006)
4. Engelkamp, J.: Memory for actions. Psychology Press/Taylor & Francis, England (1998)
5. Engelkamp, J., Zimmer, H.D.: Human memory. Hogrefe & Huber, Seattle (1994)
6. Fromm, M.: Repertory Grid Methodik. Deutscher Studien-Verl., Weinheim (1995)
7. Gee, J.P.: Good video games and good learning: Collected essays on video games, learning, and literacy. No. v. 27 in New literacies and digital epistemologies, Lang, New York [u.a.] (2007)
8. Gee, J.P.: What video games have to teach us about learning and literacy. Palgrave Macmillan, Basingstoke (2007), Rev. and updated ed. edn.
9. Goldstein, J.H., Buckingham, D., Brougre, G.: Introduction: Toys, games, and media. In: Goldstein, J.H., Buckingham, D., Brougre, G. (eds.) Toys, games, and media, pp. 1–10. Erlbaum, Mahwah (2004)
10. Grabowski: Generative learning: Past, present, and future. In: Handbook of Research on Educational Communications and Technology, pp. 897–918. Macmillian, New York (2004)
11. Hawlitschek, A.: Spielend lernen aber wie? zum didaktischen design von serious games für die schule. In: Schwill, A., Apostolopoulos, N. (eds.) Lernen im digitalen Zeitalter
12. Hurley, S.: The shared circuits model (scm): how control, mirroring, and simulation can enable imitation, deliberation, and mindreading. *Behav. Brain Sci.* 31(1), 1–22 (2008), discussion 22–58 (February 2008),
<http://dx.doi.org/10.1017/S0140525X07003123>
13. Kalyuga, M.: Enhancing instructional efficiency of interactive e-learning environments: A cognitive load perspective. *Educational Psychology Review* 19, 387–399 (2007)
14. Kelly, G.A.: The psychology of personal constructs. Norton, New York (1955)
15. Kline, S.: Learners, spectators, or gamers? an investigation of the impact of digital media in the media-saturated household. In: Goldstein, J.H., Buckingham, D., Brougre, G. (eds.) Toys, Games, and Media, pp. 131–156. Erlbaum, Mahwah (2004)
16. Lampert, T., Mensink, G.B., Romahn, N., Woll, A.: Physical activity among children and adolescents in germany. results of the german health interview and examination survey for children and adolescents (kiggs). *Bundesgesundheitsblatt Gesundheitsforschung Gesundheitsschutz* 50(5-6), 634–642 (2007)
17. Linderoth, J., Lindström, B., Alexandersson, M.: Learning with computer games. In: Goldstein, J.H., Buckingham, D., Brougre, G. (eds.) Toys, Games, and Media, pp. 157–176. Erlbaum, Mahwah (2004)
18. Lucht, M.: Erfüllung der Informations- und Meinungsbildungsfunktion im Fernsehen: Eine experimentelle Studie zum Vergleich von herkömmlicher Dokumentation und Multiperspektivendokumentation. Müller, Saarbrücken (2007)
19. Newsvote: Computer games 'burn up calories' (2007),
<http://newsvote.bbc.co.uk/2/hi/health/6376637.stm>
20. Piaget, J.: Nachahmung, Spiel und Traum, 1st aufl edn., Klett, Stuttgart (1975)
21. Prensky, M.: Digital Game-Based Learning. Paragon House Publishers (2007)
22. Rheinberg, F., Vollmeyer, R., Burns, B.D.: Motivation and self-regulated learning. In: Motivational Psychology of Human Development, pp. 81–108. Elsevier, Amsterdam (2000)
23. Ritterfeld, U.: Introduction. In: Serious games: Mechanisms and Effects, pp. 3–9. Routledge, New York (2009)

24. Ritteser, M.J.: Bewegung und Lernen – Auswirkungen von Bewegung in der Schule auf Konzentration, Merkfähigkeit und Befindlichkeit, Grin, München (2007)
25. Schnotz, W., Kuerschner, C.: A reconsideration of cognitive load theory. *Educational Psychology Review* 19(4), 469–508 (2007),
<http://springerlink.com/content/nt077668737230p6/>
26. Simpson, E., Stansberry, S.: Video games and teacher development: Bridging the gap in the classroom. In: *Games: Purpose and Potential in Education*, pp. 1–22. Springer, US (2009),
<http://www.springerlink.com/content/p4448631p1800276/>
27. Squire, K.D.: *Replaying history: learning world history through playing “civilization iii”*. Ph.D. thesis, Indianapolis, IN, USA, chair-Barab, Sasha A (2004)
28. Stevanoni, E., Salmon, K.: Giving memory a hand: Instructing children to gesture enhances their event recall. *Journal of Nonverbal Behavior* 29(4), 217–233 (2005),
<http://springerlink.com/content/pn3g00501v1j8087/>
29. Stork, A.: *Vokabellernen. Giessener Beiträge zur Fremdsprachendidaktik*, Narr, Tübingen (2003)
30. Südwest, M.F.: *Jim-studie 2008: Jugend, information (multi-) media. basisuntersuchung zum medienumgang 12- bis 19-jähriger* (2008),
http://www.mpfs.de/fileadmin/JIM-pdf08/JIM-Studie_2008.pdf
31. Südwest, M.F.: *Kim-studie 2008: Kinder und medien, computer und internet: Basisuntersuchung zum medienumgang 6- bis 13-jähriger in deutschland* (2008),
<http://www.mpfs.de/fileadmin/KIM-pdf08/KIM2008.pdf>
32. Sweller, J., Chandler, P.: Evidence for Cognitive Load Theory. In: *Cognition and Instruction*, vol. 8, pp. 351–362. Taylor and Francis, London (1991)
33. Tabbers, H.: Enhancing learning from dynamic visualizations: An embodied perspective. Paper Presented at the Biennial Conference of the European Association for Research on Learning and Instruction (EARLI), Amsterdam, Netherlands, August 25-29 (2009)
34. Unnithan, V.B., Houser, W., Fernhall, B.: Evaluation of the energy cost of playing a dance simulation video game in overweight and non-overweight children and adolescents. *Int. J. Sports Med.* 27(10), 804–809 (2006),
<http://dx.doi.org/10.1055/s-2005-872964>
35. Vorderer, P.: Motivation. In: *Psychology of Entertainment*, pp. 3–18. Erlbaum Assoc., Mahwah (2006)
36. Werning, B.R., Palmowski, W., Sassenroth, M.: Der Bereich der Wahrnehmung und Bewegung als sonderpädagogischer Förderbedarf. In: *Sonderpädagogik. Lernen, Verhalten, Sprache, Bewegung und Wahrnehmung*, pp. 284–318. Oldenbourg, München (2002)
37. Wittrock, M.C.: Generative learning processes of the brain. *Educational Psychologist* 27(4), 531–541 (1992)
38. Zimmer, R. (ed.): *Wahrnehmen, bewegen, lernen - Kindheit in Bewegung*. Hofmann, Schorndorf (2004)

Social Relationships as a Means for Identifying an Individual in Large Information Spaces

Katarína Kostková, Michal Barla, and Mária Bieliková

Slovak University of Technology
Faculty of Informatics and Information Technologies
Ilkovičova 3, 842 16 Bratislava, Slovakia
kostkova.katka@gmail.com, {barla,bielik}@fiit.stuba.sk

Abstract. In this paper we describe a method for identification of a particular user in large information spaces such as the Web is. Our main goal is to be able to decide whether the particular information found on the Web is relevant to the person being looked-up or not by taking into account additional background knowledge about that person: his or her social network. Our method combines semantically as well as syntactically based metrics to compare different social networks acquired during the identification process. We describe evaluation of the proposed method and its comparison to the related works.

Keywords: disambiguation, individual identification, social network, social networks comparison, background knowledge.

1 Introduction

Searching for information about a particular person on the Web is one of the most popular search types. However, due to the vast information space of the Web it became a non-trivial task. The problem is not only insufficient personal information on the analyzed web page, but mainly ambiguities in person search. Ambiguity exists mainly due to the fact that many persons may have the same name (multi-referent ambiguity) or many names are used in various different written forms (multi-morphic ambiguity) [13].

Our goal is to support this kind of search by automatically deciding whether a particular information found on the Web (thus a web page), is relevant to the person we are interested in or not. Our approach is not limited to be used on the Web only, but can be applied to any large information space, which suffers from name variants and name disambiguation problems. An example, of such an information space apart from the Web is DBLP¹ database, where the way of adding and representing information about new publications may cause that one author is represented by several variants of his or her name. The second, name disambiguation, problem is also present in DBLP when linking publications to their authors, i.e., how to correctly assign a publication, if there are several authors with the same name?

¹ Digital Bibliography and Library Project.

Employing additional, background knowledge of the domain can solve both of the aforementioned problems. In our method the additional knowledge is represented by social networks, which identify the person by her relationships with other individuals. In reality, there is only marginal probability that two different people with the same or similar names would share the same social network. Thus social networks can be used to assign information found on the Web to a right, real world person.

In order to identify a person, we perform a comparison of two social networks, coming from different sources using both semantic (i.e., based on relations) as well as syntactic (based on text similarity) comparison metrics. The first social network is constructed automatically and is based on the data extracted from the Web. The extraction itself is driven from the starting point – a web page for which we want to decide whether it is relevant to the person we are interested in or not. The second social network comes as the person’s background knowledge, e.g., from social web-based applications such as Facebook. The result of the comparison is a probability of that the input web page is related to the person of our interest.

The rest of this paper is organized as follows. In the next section, we provide an overview of related works in the field of person identification. In Section 3, we describe our method for identification of a particular user. Section 4 describes the experiments we conducted along with obtained results. Finally, we summarize the main contributions of the paper and discuss possible extensions of proposed method.

2 Related Works

The name disambiguation and name variants problems in large information spaces are in focus of many research projects. Most of approaches identify an individual in order to deal with a problem of searching for personally related information on the Web using ordinary search engines such as Google, where results are filtered, reordered or clustered appropriately [13].

During the identification process, almost every approach uses some kind of background knowledge. Some of the authors use personal information, like birth date, wife’s or husband’s name etc. [7], although the nature of this information imposes that it is not always possible to get it. Another option is to use keywords selected from analyzed web page [13] or estimated person’s interests extracted from the web page content [10] or professional category knowledge [4]. In [2] the authors prefer to use the whole *context* of a person. By the term *context* they mean all terms extracted from the entire inspected document. The background information is used to cluster web sites found by a search engine. By assuming that the name is the only feature in common to any two namesakes, groups of pages where each group is related to one namesake are returned.

Another approach uses a social network as background knowledge [6]. The social network is created by analyzing the e-mail traffic and the content. This approach is based on an assumption that when people are connected via e-mail messages, their web pages would be connected as well. The authors issue a web search query for every name occurring in the social network, take first K results from each query and connect them in a graph-like structure. Pages in the most connected part of the graph are declared as pages related to a person they were interested in. The fact, that this

solution is built on searching several persons at the time substantially alleviate the identification problem. However, even if the approach uses the social network, it takes only the *names* portion of it and ignores connections between people.

We believe that not only names of other people, but also relations between them are important to the identification problem and thus should be considered, which is exactly what our method does.

Most of approaches devoted to solving the name variants problems are based on dictionaries of nicknames containing pairs: name – name variant. In [3] the authors created dictionary of nicknames, based on morphological rules transforming a full name to various name variants. Another possible way is to extract nicknames from the Web. In [3] the authors also used sentence pattern “My name is <name>, friends call me <nickname>” for web-based name variants extraction. This approach allowed them to get also nicknames, which are not related to any of full name variants.

An approach to the name variants solving based on social networks can be seen in [9], where comparison of two candidates (paper authors) is performed by comparing the names of their co-authors. If an overlap of co-authors was significant, a name variant was probably found. Similar idea was elaborated in [11], where they were analyzing common publications of two distinct co-authors. If such publications could not be found (considering time and topic factors), those two co-authors are probably representing one real world person.

Most of existing works in the name disambiguation domain use some information about a person. However, the satisfactory solution of the problem of identifying a particular user was not achieved yet. We propose a method based on an analysis of social networks, which uniquely describe a person. During the processing, we do not use only the names of people from the social networks, but consider also relations between them.

3 Method of Identification of an Individual on the Web

We specify the problem as follows: “For any web page, we want to determine the probability that the web page is related to a given person”. Our additional information (background knowledge) about that person is his or her social network. We propose a method, which compares two persons, represented by their names, which are either namesakes or two variants of a name of the same person. The process of the person identification consists of the following steps:

1. *Social network extraction*, where we extract social network from the person’s web page, to have a second input to compare with the given background knowledge
2. *Comparison*, where the given social network is compared with information extracted in previous step.

As an input, we get a candidate web page of a person we are looking for and background knowledge (social network) related to this person. After the process of identification, we get probability of that candidate web page is related to the person we have been looking for.

3.1 Social Network Extraction

The extraction of social network is based on an analysis of interconnected web pages. If a relation exists between two web pages, we add a relation between people (names), which are stated on those two web pages.

The process, consisting of the following steps, is being performed recursively until the *depth of recursion* set as a parameter is not achieved:

1. An input web page URL is added to the list of URLs to be processed, the *actual depth* is set to 0.
2. If the list of URLs is empty, the process ends. Otherwise, first URL from the list is picked-up.
3. Source text of the web page is obtained for the given URL.
4. URL extraction – if the *actual depth* is lower than the given *depth of recursion*, URLs linking to other web pages are extracted from the page source and added to list of URLs to be processed. This step is omitted if the required *depth of recursion* was reached.
5. Name extraction – person names are extracted from the page source. Our approach to name extraction is based on dictionary of English given names. We search for these names on the web page and then we check whether there is a surname, middle name, or middle name initial around the found given name, all using regular expressions.
6. Social network creation – names and relations between web pages are added to social network.
7. Increment the *actual depth* and continue with step 2.

In step 4 we extract references to other websites. By following links between web pages the relations between persons are obtained.

3.2 Social Network Comparison

In second step of our method we compare two social networks, the one given as background knowledge and the one extracted from the Web, by examining their nodes and edges. Each node represents one person. We do not merge the two networks, but rather perform a mapping of the nodes based on syntactical similarity (in our case, we decided to use Levenshtein distance) of the node names. The mapping produces one-to-one as well as one-to-many connections between the two networks (Fig. 1). These connections are subsequently used to compare the networks.

However, by employing only syntactic comparison, we achieve unsatisfactory results, with many name variants still unresolved. Therefore we use syntactically based metric only as a pre-filter to a semantic-based comparison (which is thus performed only for nodes, which have the syntactic similarity above the pre-defined threshold). The threshold differs according to the nature of the social networks to be compared. For example, if we have a social network of authors from DBLP, proper similarity of names can be higher than in a social network from some chat portal.




Fig. 1. Syntactic mapping of two social networks

We employed existing semantic metric based on relations *Connected Triple* [12] for comparison. When social network is defined as a graph G with vertices V representing people and edges E representing relations between them, $G(V, E)$, then a connected triple is a graph $G'(V_{CT}, E_{CT})$ of three vertices with $V_{CT} = \{X, Y, Z\}$, $V_{CT} \subset V$ and two edges $E_{CT} = \{E_{XY}, E_{YZ}\}$, $E_{CT} \subset E$, where $E_{XZ} \notin E$. On Figure 2, we present an example of a connected triple – the two vertices (persons) being compared must not be connected by an edge. This is obvious, as if there had been an edge between two compared vertices, these two persons would have knew each other or would have worked together, which means they are *really* two different persons. If the edge between them is missing, then there is a possibility that these two vertices represent the *same person*.

The process resulting in similarities of pairs of vertices from two social networks consists of following steps:

1. Select a pair of vertices, each coming from different social network.
2. Compute similarity of vertices using Levenshtein distance
 - a. If the similarity is above the threshold
 - i. compare the vertices using enhanced Connected Triples
 - b. If the similarity is below the threshold
 - i. continue with another pair of vertices
 - ii. finish if there are no more unprocessed pairs

The basic original formula for calculating similarity between two persons based on connected triples was defined in [12]:

$$similarity(i, j) = \frac{|C_{ij}|}{\max_{k=0..m, l=0..n} |C_{kl}|} \quad (1)$$

where i a j are compared persons, C_{ij} is a count of connected triples of compared persons and C_{kl} is a count of connected triples where at least one the compared persons is involved.




Fig. 2. A connected triple. Note, that there is no relation between Jozef and Jozeph.

The similarity is computed as a number of connected triples between compared vertices, divided by a maximum of connected triples in graph, between any two vertices. Notice, that this formula considers all relations in the graph, even those that are not related to the people being compared.




Fig. 3. An example of social network

Let us consider a social network depicted in Figure 3, where we are comparing two persons, Ivan and Ivo, and we know that it is the same person. If we use formula 1, the probability that these two vertices represent the same person in reality is $3/5$, because number of connected triples where both members are the nodes being compared is 3, but maximum number of connected triples between any two nodes is 5, between *Gustav* and *Zdeno*.

We modified the formula 1 considering only relations and connected triples, where at least one of its members is one of the compared persons. In other words, we count similarity of two people as a number of connected triples, where members are both of them, divided by a maximum number of connected triples, where at least one of them is a member. The new formula with the same meaning of variables as in formula 1:

$$\text{similarity}(i, j) = \frac{|C_{ij}|}{\max_{k=i \vee j, l=0..n} |C_{kl}|} \quad (2)$$

Taking the same example of social network from Figure 3, the probability that Ivan and Ivo is the same person will be $3/3$, because the number of connected triple where

both members are compared nodes is 3 and also maximum number of connected triple, where at least one member is one of the compared nodes is 3.

The connected triples are detected according to the following algorithm

1. take the first relationship from the first social network in the form "from node *A* to node *B*"
2. consider node *A* as an intermediate node
 - a. find a mapping of node *A* in the second social network (node 3 from the example in Fig. 4)
 - b. find all adjacent of mapped node in the network (nodes 8,9,1)
 - c. add corresponding triples (B,A,x) into connected triples, where x is the found adjacent node from the second network
3. consider node *B* as an intermediate node and perform the same procedure as in step 2




Fig. 4. Identification of connected triples between two social networks

After having identified the connected triples, we can proceed to the comparison itself based on the aforementioned formula 2.

4 Evaluation

We developed a software tool, which implements the aforementioned method in order to evaluate its properties. We evaluated name extraction, used during the social network construction based on linked web pages as was described in section 3.1 and our method for social networks comparison based on syntactic and semantic properties of the networks.

4.1 Name Extraction

We evaluated our name extractor on real web pages, from which we extracted names and then manually checked the results in order to determine precision and recall of our approach. We took pages of American universities and homepages of professors from these universities, usually containing lists of their students and/or publications. We did not consider any structural features of the pages (HTML markup) and were extracting names in different combinations of first name, middle name, middle name initial and surname from the full texts. Table 1 shows our results.

We achieved a satisfactory precision, but low recall. The reason of lower recall values is that we used a dictionary of English names, but our dataset contained several pages with many foreign names, like Bernd, Xuerui or Sameer which were not included in the dictionary. Using extended dictionary of first names can solve this problem.

Table 1. Name extraction results

	Mean	Standard Deviation	Min	Max
Precision	95,8%	2,6	94,7%	98,4%
Recall	51,3%	12,2	31,6%	59,8%

4.2 Connected Triples

We evaluated our approach to social network comparison based on the modified connected triple method. We obtained a social network from Slovak Companies Register on the Internet (<http://orsr.sk>), which has about 300 000 vertices and more than 460 000 edges, representing companies and individuals, which are somehow connected to those companies. For prototyping purposes, we decided to take only a part of this social network.

We have selected all people with surname “Havran” and all relations of these people. We did the same for surname “Novak” as these two surnames are very common in Slovakia. In order to evaluate precision and recall of our approach, we identified duplicities in that smaller social networks manually as well as by using a domain-specific heuristics (baseline) employing specific methods to identify duplicates in the whole social network. The duplicates, in this case, were persons with very similar names and addresses.

The results are shown in Table 2. We achieved good results in precision and recall and we also found more duplicates than the baseline method. However, we did not find all duplicates, which the baseline did, which can be caused by the fact that we took only a part of the whole social network whereas the baseline operated on the whole graph. We also compared our modified connected triple algorithm with an original connected triple algorithm. Results are shown in Table 3.

Table 2. Connected triple algorithm results

	Havrans	Novaks
Vertices	329	256
Edges	610	526
Levenshtein 85% and connected triples (our method)	28	23
Levenshtein 85% and connected triples and same address	23	15
Baseline	14	19
Our method \cap baseline	9	12
Precision	82,2%	65,2%
Recall	82,2%	60,0%

Table 3. Modified connected triple algorithm vs. original connected triple algorithm

	Havrans	Novaks
Identified duplicates with similarity > 50% (by our modified method)	9	20
Identified duplicates with similarity > 50% (by the original method)	3	0
Precision of modified	- 23,3%	+ 56,5%
Recall of modified	+ 14,3%	+ 52,5%

When we defined the required similarity of duplicates to be more than 50 %, we got better recall values with our modified algorithm and even if precision achieved on Havrans dataset is lower, we identified more duplicates than the original connected triple algorithm. We can thus conclude that our modification was useful.

5 Conclusions

In this paper we presented a novel method for identification of a particular user in the large information space such as the Web. We based our solution on social networks, which act as a background knowledge about the person. We combine syntactically and semantically based metrics in the process of network comparison, which determines whether the information we found on the Web is related to the person we are looking for or not.

We evaluated our approach to gain verification of our changes in the original connected triple metric as well as to verify our assumptions about overall usability of our method against specialized, domain dependent heuristics. The experiments showed that our *domain independent approach* based on a modified connected triple metric is performing well, compared to either domain dependent heuristics or original connected triple metric. More, we were able to discover *different* duplicities than the domain dependent heuristics, which promises that their combination would allow for the achievement of even better results as they eliminate the weak points of each other.

We should point out that our method is built on a broader concept of social networks comparison, which has a great usage potential especially in, but not limited to, the domain of user modeling [1]. We can determine relevance of data found on the Web, compare two social networks and find corresponding nodes within them. We see a nice application of our method in social portals. When a new user registers into a social portal, we can ask for his homepage (or any other page stating information which is relevant to that user). We then extract his social network from the “web environment” of this page and take it as our background knowledge about this new user. Subsequently, we can use it to find his friends, which are already registered within our social portal. Then we can recommend these people to the newly registered user as potential friends or we could offer him services, which his friends are also interested in.

Our method of identification of an individual can also serve to other approaches such as query expansion based on user context [8], where it can help to disambiguate entities present in the current user context. Apart from search-related tasks, it can also

assist in automatic content annotation [5], where it helps to associate a correct instance to a retrieved information.

We already see several possible extensions of our work. Probably the most interesting is to consider different attributes of people or relations in the algorithm. For instance, different types of relationships may have different weights of contribution to the final result, depending on their importance. Promising seems to be also a combination of our approach with various clustering methods and techniques.

Acknowledgement. This work was partially supported by VEGA 1/0508/09, KEGA 028-025STU-4/2010 and it is the partial result of the Research & Development Operational Programme for the project Support of Center of Excellence for Smart Technologies, Systems and Services, ITMS 26240120029, co-funded by the ERDF.

References

1. Barla, M., et al.: Rule-based User Characteristics Acquisition from Logs with Semantics for Personalized Web-based Systems. *Computing & Informatics* 28(4), 399–427 (2009)
2. Bollegala, D., et al.: Disambiguating Personal Names on the Web Using Automatically Extracted Key Phrases. In: *ECAI 2006*, pp. 553–557. IOS Press, Amsterdam (2006)
3. Driscoll, P., Yarowsky, D.: Disambiguation of Standardized Personal Name Variants (2007), <http://www-lipn.univ-paris13.fr/~poibeau/mmies/mmies04.pdf> (17.05.2010)
4. Han, X., Zhao, J.: Web personal name disambiguation based on reference entity tables mined from the web. In: *Proc. of the 11th Int. Workshop on Web information and Data Management, WIDM 2009, Hong Kong, China*, pp. 75–82. ACM, New York (2009)
5. Laclavík, M., et al.: Ontea: Platform for Pattern Based Automated Semantic Annotation. *Computing and Informatics* 28(4), 555–579 (2009)
6. Malin, B.: Unsupervised Name Disambiguation via Social Network Similarity. In: *Workshop on Link Analysis, Counterterrorism, and Security* (2005)
7. Mann, G.S., Yarowsky, D.: Unsupervised Personal Name Disambiguation. In: *Natural Language Learning Workshop at HLT-NAACL 2003, Association for Computational Linguistics*, pp. 33–40 (2003)
8. Návrat, P., Taraba, T., Bou Ezzeddine, A., Chudá, D.: Context Search Enhanced by Readability Index. In: *Artificial Intelligence in Theory and Practice. IFIP WCC. IFIP Series*, vol. 276, pp. 373–382. Springer Science+Business Media, LLC, New York (2008)
9. On, B., Lee, D., Kang, J., Mitra, P.: Comparative study of name disambiguation problem using a scalable blocking-based framework. In: *Proc. of the 5th ACM/IEEE-CS Joint Conference on Digital Libraries, JCDL 2005*, pp. 344–353. ACM, New York (2005)
10. Omelina, L.: Extracting Information from Web Pages Based on Graph Models. In: Bieliková, M. (ed.) *IIT.SRC 2009: Proc. of Student Research Conference, STU Bratislava*, pp. 105–112 (2009)
11. Reuther, P.: Personal Name Matching: New Test Collections and a Social Network based Approach. University of Trier, Mathematics/Computer Science: Tech. Report 06-01 (2006)
12. Reuther, P., et al.: Managing the Quality of Person Names in DBLP. In: Gonzalo, J., Thanos, C., Verdejo, M.F., Carrasco, R.C. (eds.) *ECDL 2006. LNCS*, vol. 4172, pp. 508–511. Springer, Heidelberg (2006)
13. Wan, X., Gao, J., Li, M., Ding, B.: Person resolution in person search results: WebHawk. In: *Proc. of the 14th ACM Int. Conf. on Information and Knowledge Management, CIKM 2005, Bremen, Germany*, pp. 163–170. ACM, New York (2005)

PART III
Data Mining and Machine
Learning

J-PMCRI: A Methodology for Inducing Pre-pruned Modular Classification Rules

Frederic Stahl, Max Bramer, and Mo Adda

University of Portsmouth, School of Computing,
Buckingham Building, Lion Terrace, Portsmouth PO1 3HE, United Kingdom
{Frederic.Stahl,Max.Bramer,Mo.Adda}@port.ac.uk

<http://www.port.ac.uk>

Abstract. Inducing rules from very large datasets is one of the most challenging areas in data mining. Several approaches exist to scaling up classification rule induction to large datasets, namely data reduction and the parallelisation of classification rule induction algorithms. In the area of parallelisation of classification rule induction algorithms most of the work has been concentrated on the Top Down Induction of Decision Trees (TDIDT), also known as the ‘divide and conquer’ approach. However powerful alternative algorithms exist that induce modular rules. Most of these alternative algorithms follow the ‘separate and conquer’ approach of inducing rules, but very little work has been done to make the ‘separate and conquer’ approach scale better on large training data. This paper examines the potential of the recently developed blackboard based J-PMCRI methodology for parallelising modular classification rule induction algorithms that follow the ‘separate and conquer’ approach. A concrete implementation of the methodology is evaluated empirically on very large datasets.

Keywords: Parallel Rule Induction, Blackboard Systems, Modular Rule Induction.

1 Introduction

Parallel Data Mining is motivated by the need to scale up to massive datasets. The runtime complexity of data mining is typically linear or worse with respect to both the number of instances and the number of attributes in a dataset, thus massive datasets can be computationally expensive to analyse. There are many examples of applications that need to deal with the problem of mining massive datasets. For example, in bioinformatics and chemistry there are large datasets which are generated in different kinds of Molecular Dynamics (MD) simulations and there is a considerable desire to find ways to manage, store and find complex relationships in this generated data [2,14]. Other examples of massive datasets are those in the area of astronomy for example generated by the NASA system of earth orbiting satellites and other space-borne probes launched in 1991 and still ongoing [17]. A further example in astronomy is the Sloan survey [16].

In the area of classification rule induction most of the approaches to scaling up the algorithms by parallelisation have concentrated on the Top Down Induction of Decision Trees (TDIDT), also known as the ‘divide and conquer’ approach. Notable parallel ‘divide and conquer’ algorithms and systems are SLIQ [9], SPRINT [11] and their successor, the ScalParC [8] algorithm. For building split points of attributes, SLIQ uses a sorted attribute list for each attribute of the form $\langle \text{record id}, \text{attribute value} \rangle$, and a class list of the form $\langle \text{class value}, \text{tree node} \rangle$. Each processor is assigned to a sublist of each attribute list plus the class list. Different to SLIQ, SPRINT and ScalParC build a sorted attribute list for each attribute of the form $\langle \text{record id}, \text{attribute value}, \text{class value} \rangle$ and no class list. Again each processor is assigned to a sublist of each attribute list and exchanges the statistics needed in order to determine the best split point. Each processor builds the tree simultaneously and updates the attribute lists accordingly.

On the other hand there have been virtually no attempts to parallelise algorithms that induce modular rules rather than decision trees. Modular rules are rules that do not necessarily fit into a decision tree. Most modular classification rule induction algorithms follow the ‘separate and conquer’ approach. The Parallel Modular Classification Rule Induction (PMCRI) methodology has recently been developed in order to provide a methodology that allows the parallelisation of a whole subset of modular classification rule induction algorithms. In this work we will outline the PMCRI methodology and highlight an extension of it, that integrates pre-pruning, the J-PMCRI methodology. J-PMCRI is designed to run in a local area Network (LAN) in order to provide an inexpensive data mining solution for large datasets for modest sized organisations that cannot afford supercomputers. This work provides a concrete implementation of the J-PMCRI methodology that incorporates pre-pruning and is presented and evaluated empirically.

2 J-PMCRI

The PMCRI/J-PMCRI methodology is designed to parallelise any member of the Prism family of algorithms. Section 2.1 will give a brief introduction to the Prism family and Section 2.2 will describe J-PMCRI.

2.1 Prism Family of Algorithms

Cendrowska’s criticism [7] of a tree representation of classification rules is that they do not directly allow the induction of modular rules such as:

$$\begin{aligned} & \text{IF } a = 1 \text{ and } b = 1 \text{ then class} = A \\ & \text{IF } c = 1 \text{ and } d = 1 \text{ then class} = B \end{aligned}$$

Such rules do not necessarily have attributes in common in their rule terms unlike for the representation in tree format. Inducing such rules will lead to the replicated subtree problem, first described by [7]. Prism has been examined

empirically and has shown to be comparable to decision trees and in some cases even shows a better predictive performance [3]. Subsequent research developed further variations of Prism algorithms notably the PrismTCS algorithm and PrismTC [4,5]. The theoretical worst case complexity of Prism algorithms is $O(N^2 \cdot M)$ where N is the number of data instances and M is the number of attributes. However an empirical examination revealed an average complexity of $O(N \cdot M)$ [15].

Our implementation of the basic Prism approach for continuous data only is summarised in the following pseudo code: where A_x is a possible attribute value and D is the training dataset.

```

For each class i do {
  Step 1: Calculate for each  $A_x$  probabilities  $p(\text{class} = i | A < A_x)$  and  $p(\text{class} = i | A \geq A_x)$ 
  Step 2: Select the condition with the maximum probability as rule term
          and create a subset  $D'$  of  $D$  that comprises all instances
          that match the selected condition.
  Step 3: Repeat 1 to 2 for  $D'$  until  $D'$  only contains instances
          of classification  $i$ . The induced rule is then a
          conjunction of all the selected conditions and  $i$ .
  Step 4: Create a new  $D'$  that comprises all instances of  $D$  except
          those that are covered by the rules induced for class  $i$  so far.
  Step 5: Repeat steps 1 to 4 until  $D'$  does not contain any
          instances of classification  $i$ .
}

```

J-pruning, a pre-pruning method, developed by Bramer [4], demonstrated good performance with respect to the predictive accuracy of the classifier [4,6]. J-pruning is based on the J-measure from Smyth and Goodman [13] and can be applied to TDIDT and any member of the Prism family. Also it has been observed in [15] that J-pruning lowers the number of rules and rule terms induced and thus the number of iterations of the algorithm, which in turn lowers the runtime.

According to Smyth and Goodman [13] the average information content of a rule of the form $IF Y = y THEN X = x$ can be quantified by the following equation:

$$J(X; Y = y) = p(y) \cdot j(X; Y = y) \quad (1)$$

The J-measure is a product of two terms. The first term $p(y)$ is the probability that the antecedent of the rule will occur. It is a measure of the hypothesis simplicity. The second term $j(X; Y=y)$ is the j-measure or cross entropy. It is a measure of the goodness-of-fit of a rule and is defined by:

$$j(X; Y = y) = p(x | y) \cdot \log\left(\frac{p(x|y)}{p(x)}\right) + (1 - p(x | y)) \cdot \log\left(\frac{1-p(x|y)}{1-p(x)}\right) \quad (2)$$

If a rule has a high J-value then it tends to have a high predictive accuracy as well. The J-value is used to identify when a further specialisation of the rule is likely to result in a lower predictive accuracy due to overfitting. The basic idea is to induce a rule term and if the rule term would increase the J-value of the current rule then the rule term is appended. If not then the rule term is discarded and the rule is finished.

2.2 The J-PMCRI Methodology

Both the PMCRI and J-PMCRI methodology distribute the workload of inducing rule terms over a network of workstations in a LAN by distributing the training data. The basic rule induction procedure of the J-PMCRI can be divided into three steps:

1. the training data is distributed over n workstations
2. learning algorithms on each workstation cooperate to induce a rule and communicate in order to get a global view of the state of the classifier
3. combine the local parts of the rule to a final classifier.

With regards to *step 1*, a workload balance is achieved by building attribute lists out of each attribute in the training data similar to those in the SPRINT [11] algorithm. Attribute lists are of the structure $\langle \text{record id, attribute value, class value} \rangle$. These attribute lists are then distributed evenly over n workstations. Unlike SPRINT, which achieves a workload balance by splitting each attribute list into n sublists and each workstation gets assigned a part of each attribute list. We distribute entire attribute lists evenly over all workstations. Distributing parts of the attribute lists may achieve a better workload balance at the very beginning, however it is likely that it will result in a considerable workload imbalance later on in the algorithm's execution as part attribute lists may not evenly decrease in size [12]. Distributing entire attribute lists may only impose a slight workload imbalance at the beginning of the algorithm in J-PMCRI, however the relative workload on each workstation will approximately stay the same. With regards to *step 2* every workstation holds a subset of the attribute lists and thus a subset of the feature space of the training data in memory. Each workstation can induce the conditional probabilities for candidate rule terms for their attribute lists independently and thus can derive a candidate rule term that is locally the best for the attribute lists in the workstation's memory. However in order to find out which is the globally best rule term the workstation needs to exchange information about the locally induced rule terms with the other workstations. For this purpose we use a distributed blackboard architecture like the one in [10]. A blackboard architecture can be seen as a physical blackboard, that is observed and used by several experts with different knowledge domains that have a common problem to solve. Each expert will use its own knowledge plus information written on the blackboard by other experts in order to derive new information and write it on the blackboard. As a software model this principle can be represented by a client server architecture. In the blackboard literature these experts are also called Knowledge Sources (KS). The basic architecture of J-PMCRI is shown in figure 1. The attribute lists are distributed over k KS machines in the network. The blackboard system is partitioned into two logical panels or partitions that have a different meaning to the KS machines, one for information about local rule terms, on the 'local rule term information' (LI) partition and one for global information (GI).

Every KS is hosted on a separate machine in the network and is able to induce the rule term that is locally the best one for the attribute lists it holds

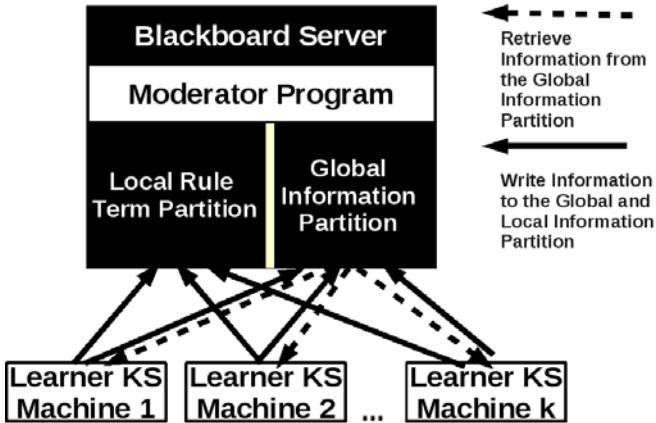


Fig. 1. PMCRI's communication pattern using a distributed blackboard architecture

in memory. It then writes information about the induced rule term on the LI partition and next it awaits the global information it needs in order to induce the next rule term being advertised on the GI partition. The information submitted about the rule term is the probability with which the induced rule term covers the target class on the local attribute list collection. The moderator program which is embedded in the blackboard server collects the probabilities from the LI partition, compares them and advertises the name of the KS that induced the globally best rule on the GI partition. The KSs that have only induced a locally best rule term will delete the term from their memory. The KS with the globally best rule term will keep the rule term in memory and then communicate the ids of the instances that are uncovered by this rule term to the other waiting KS, using the GI partition. Now the next rule term can be induced in the same way.

The information needed to calculate the J-value as outlined in formulae 1 and 2 in Section 2.1, is the count of how many data instances (list instances) the rule term covers, the count of how many instances covered by the rule term are assigned with the target class, the total number of instances and the total number of instances covering the target class. This information is solely contained in the attribute lists the locally best rule term was induced from. The basic idea is to not only write the probability with which the rule term was induced but also the J-value of the rule, if the rule term would be appended, on the LI partition. The moderator first compares the probabilities and thus identifies the KS that has induced the globally best rule term. Next the moderator compares the new J-value with the previous one. If the new J-value is larger than the winning KS, it is advertised on the GI partition otherwise the moderator writes on the 'global information partition' that the rule needs to be pruned. The KS will act accordingly. They will discard the last locally induced rule term and start inducing the first candidate rule term for the next rule. The moderator pseudo code is shown below:

```

bestJ=0; bestProb=0; ExpertInfo;
for each submitted rule term do{
  IF(t.p>bestProb){
    if(t.j>bestJ){
      ExpertInfo = Best term induced on t.ExpertName;}
    else{
      ExpertInfo = prune rule;}
  }
}

```

The following pseudo code describes how a KS induces rules based on the Prism algorithm outlined in Section 2.1. However it can easily be adapted to any other member of the Prism family.

```

Rule_Set{};
For each class i do{
  Step 1: Initialise empty rule r;
  Step 2: Calculate for each attribute list Ax p(class = i| A < Ax)
    and p(class = i| A >= Ax) and the corresponding J-value and coverage;
  Step 3: Select the condition with the maximum probability as locally best
    rule term t;
  Step 4: Write KS name plus p, J-value and coverage on the LI partition
    and observe the GI partition for information I:
    IF(I equals own name){
      add t to r;
      generate from t uncovered Ids and communicate them to the remaining KS;
      delete attribute list instances that match the uncovered Ids;
    }ELSE IF(I = "prune rule"){
      add r to Rule_Set;
    }ELSE (I equals name of different KS){
      retrieve Ids that are uncovered from the globally best rule term;
      delete attribute list instances that match the uncovered Ids;
    }
  Step 4: Restore all attribute lists to their initial size;
    delete all instances from all attribute lists that are covered by Rule_Set{};
} While still instances left in the training data

```

At the end of J-PMCRI's execution the KS will only have the rule terms in their memory that they induced locally and that were globally the best ones. They simply communicate the rule terms plus information about which rule and which class they were induced for to the blackboard server. There the globally best rule terms are assembled to the final rules.

The Prism algorithm parallelised J-PMCRI would induce exactly the same rules as the serial version would using J-pruning, only they are parallel and faster to generate and thus can computationally cope with much larger data volumes.

3 Evaluation of the J-PMCRI

Previous evaluations of the methodology have focussed mainly on the size up behaviour, which measures the runtimes of the system with respect to the number of training instances or data records. For both the number of training instances and data records in previous work a slightly better than linear size up behaviour has been observed, meaning that if the amount of data doubles then the runtimes nearly double as well, but not quite. These experiments have been performed with several smaller datasets from the UCI repository that have been appended

to themselves in order to increase the number of instances and attributes and thus the workload. The reason for this was to keep the number of rules and rule terms induced constant, so that they will not influence the runtime severely, as a different number of rules and rule terms results in a different number of iterations of the algorithm and thus in a different runtime.

In this work we evaluate the J-PMCRI methodology with respect to the number of processors (workstations) used assuming that there is only one KS hosted on each workstation. Also we will evaluate this on two real large datasets from the infobiotics repository [1], which comprises very large datasets for benchmarking purposes. The first dataset is called (here) the infobio2 dataset. It comprises 60 attributes, 4 classifications and more than 2.3 million continuous training instances. The second dataset used is called (here) infobio3 and comprises 220 attributes, 4 classifications and 2.5 million continuous data instances. A standard metric to evaluate a parallel algorithm or a parallel architecture is the speedup factor. With the speedup factors one can calculate by how much the parallel version of an algorithm is faster using p processors compared with 1 processor.

$$S_p = \frac{R_1}{R_p} \quad (3)$$

Formula 3 represents the speedup factors S_p . R_1 is the runtime of the algorithm on a single machine and R_p is the runtime on p machines. In the ideal case, the speedup factors are the same as the number of processors used. In reality, the speedup factor will be below the number of processors for various reasons, such as communication and synchronisation overheads imposed by each processor.

Figure 2 shows speedup factors plotted versus the number of processors used on fixed numbers of training instances of the infobio2 dataset. The speedup factors increase with an increasing number of processors then decrease again. This can be explained by the fact that using more processors will impose a

Speedup Factors of J-PMCRI on infobio2




Fig. 2. Speedup factors obtained for J-PMCRI with the PrismTCS algorithm on the infobio2 datasets

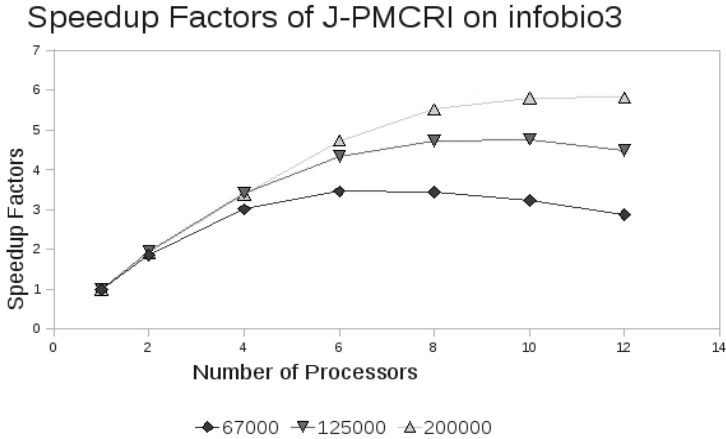


Fig. 3. Speedup factors obtained for J-PMCRI with the PrismTCS algorithm on the infobio3 datasets

larger communication overhead as well as managing overhead. However, what can also be observed is that the best speedup is reached for a larger number of processors if the number of training instances is large as well. Thus, loosely speaking, the larger the number of training instances the more processors are beneficial. Please note that the experiments illustrated in figures 2 and 3 do not include an experiment with all data instances. The reason for this is that it was not possible to perform this experiment on a 1 processor configuration due to memory constraints.

Figure 3 shows similar experiments as in figure 2 only this time on the infobio3 dataset in order to show that the results are reproducible.

Figure 4 shows the speedup factors obtained on a much larger portion of the training data compared with the experiments in figures 2 and 3. The experiments on infobio2 comprised a total of 2,338,121 (and 60 attributes) instances and on infobio 3 a total of 700,336 (and 220 attributes). The reason for these experiments is to show that J-PMCRI would perform well with many processors if the data were large enough. For the calculation of the speedup factors the system needs to be based on a configuration with fewer processors than the observed configurations. In the ideal case of one processor, however, the training data is too large to fit in the memory of a single workstation. We based the speedup factors on a 4 processor configuration. The results are shown in figure 4. What can be seen in figure 4 is that the speedup factors are still increasing with a total of 12 learner KS machines. This is different compared with the experiments outlined in figures 2 and 3, where 12 learner KS machines were not beneficial anymore. The fact that the speedup factors are still growing for 12 processors highlights even more the observed behaviour that using more processors is more beneficial the larger the amount of training data.




Fig. 4. Speedups obtained using J-PMCRI on very large training data that cannot be handled by a single machine

4 Conclusions

This paper has addressed the need for better scaling classification rule induction algorithms. The argument established in the introduction was that there are many approaches to speeding up the ‘divide and conquer’ approach which induces decision trees. But there are virtually no attempts to scale up the ‘separate and conquer’ approach for inducing modular rules. This work’s aim is to establish a methodology for parallelising and thus scaling up algorithms of the Prism family that induce modular rules. The methodology is outlined together with a pre-pruning facility and pseudo code is given in order to illustrate the basic concept. The system is based on a blackboard approach, can be run in a LAN and has been evaluated with respect to the number of processors (workstations) used. Using J-PMCRI has the very desirable property that the optimum number of processors increases as the amount of training data becomes larger.

Ongoing and future work comprises the implementation of further algorithms of the Prism family using the J-PMCRI methodology. Also the speedup factors illustrated in Section 3 may be improved by using more than just one blackboard. For instance the connection in the LAN (assuming a switched network) that connects the Blackboard server with the LAN will be used by all KS. However if further blackboard servers are introduced each blackboard only needs to serve a subset of the KS machines. These Blackboard servers will need to be synchronised as well, but this could easily be established using a separate smaller network for synchronising the blackboards.

References

1. Bacardit, J., Krasnogor, N.: The infobiotics PSP benchmarks repository. Technical report (2008)
2. Berrar, D., Stahl, F., Silva, C.S.G., Rodrigues, J.R., Brito, R.M.M.: Towards data warehousing and mining of protein unfolding simulation data. *Journal of Clinical Monitoring and Computing* 19, 307–317 (2005)

3. Bramer, M.: Automatic induction of classification rules from examples using N-Prism. In: *Research and Development in Intelligent Systems XVI*, Cambridge, pp. 99–121. Springer, Heidelberg (2000)
4. Bramer, M.: An information-theoretic approach to the pre-pruning of classification rules. In: Neumann, B., Musen, M., Studer, R. (eds.) *Intelligent Information Processing*, pp. 201–212. Kluwer, Dordrecht (2002)
5. Bramer, M.: Inducer: a public domain workbench for data mining. *International Journal of Systems Science* 36(14), 909–919 (2005)
6. Bramer, M.: *Principles of Data Mining*. Springer, Heidelberg (2007)
7. Cendrowska, J.: PRISM: an algorithm for inducing modular rules. *International Journal of Man-Machine Studies* 27(4), 349–370 (1987)
8. Joshi, M.V., Karypis, G., Kumar, V.: Scalparc: a new scalable and efficient parallel classification algorithm for mining large datasets. In: *Proceedings of the First Merged International... and Symposium on Parallel and Distributed Processing 1998, Parallel Processing Symposium, IPPS/SPDP 1998*, pp. 573–579 (1998)
9. Metha, M., Agrawal, R., Rissanen, J.: SLIQ: a fast scalable classifier for data mining. In: Apers, P.M.G., Bouzeghoub, M., Gardarin, G. (eds.) *EDBT 1996*. LNCS, vol. 1057, pp. 18–32. Springer, Heidelberg (1996)
10. Nolle, L., Wong, K.C.P., Hopgood, A.: DARBS: a distributed blackboard system. In: *Proceedings of the Twentyfirst SGES International Conference on Knowledge Based Systems and Applied Artificial Intelligence*, Cambridge. Springer, Heidelberg (2001)
11. Shafer, J., Agrawal, R., Metha, M.: SPRINT: a scalable parallel classifier for data mining. In: *Proc. of the 22nd Int'l Conference on Very Large Databases*, pp. 544–555. Morgan Kaufmann, San Francisco (1996)
12. Sirvastava, A., Han, E., Kumar, V., Singh, V.: Parallel formulations of Decision-Tree classification algorithms. In: *Data Mining and Knowledge Discovery*, pp. 237–261 (1998)
13. Smyth, P., Goodman, R.M.: An information theoretic approach to rule induction from databases. *Transactions on Knowledge and Data Engineering* 4(4), 301–316 (1992)
14. Stahl, F., Berrar, D., Silva, C.S.G., Rodrigues, J.R., Brito, R.M.M.s.: Grid warehousing of molecular dynamics protein unfolding data. In: *Proceedings of the Fifth IEEE/ACM Int'l Symposium on Cluster Computing and the Grid*, Cardiff, pp. 496–503. IEEE/ACM (2005)
15. Stahl, F., Bramer, M., Adda, M.: Parallel rule induction with information theoretic pre-pruning. In: *Research and Development in Intelligent Systems XXV*, Cambridge, pp. 151–164. Springer, Heidelberg (2010)
16. Szalay, A.: *The Evolving Universe*. ASSL 231 (1998)
17. Way, J., Smith, E.A.: The evolution of synthetic aperture radar systems and their progression to the EOS SAR. *IEEE Transactions on Geoscience and Remote Sensing* 29(6), 962–985 (1991)

Enhancement of Infrequent Purchased Product Recommendation Using Data Mining Techniques

Noraswaliza Abdullah, Yue Xu, Shlomo Geva, and Mark Looi

Discipline of Computer Science
Faculty of Science and Technology
Queensland University of Technology, 2 George Street
Brisbane QLD 4000
noraswaliza.abdullah@student.qut.edu.au,
{yue.xu, shlomo.geva, m.looi}@qut.edu.au

Abstract. Recommender Systems (RS) have emerged to help users make good decisions about which products to choose from the vast range of products available on the Internet. Many of the existing recommender systems are developed for simple and frequently purchased products using a collaborative filtering (CF) approach. This approach is not applicable for recommending infrequently purchased products, as no user ratings data or previous user purchase history is available. This paper proposes a new recommender system approach that uses knowledge extracted from user online reviews for recommending infrequently purchased products. Opinion mining and rough set association rule mining are applied to extract knowledge from user online reviews. The extracted knowledge is then used to expand a user's query to retrieve the products that most likely match the user's preferences. The result of the experiment shows that the proposed approach, the Query Expansion Matching-based Search (QEMS), improves the performance of the existing Standard Matching-based Search (SMS) by recommending more products that satisfy the user's needs.

Keywords: Recommender system, opinion mining, association rule mining, user review.

1 Introduction

The large amount of information that is available on the Internet leads to an information overload problem [1]. Recommender systems (RS) have emerged to help users deal with this problem by providing product suggestions according to their needs and requirements. Nowadays, recommender systems have been widely applied by major e-commerce websites for recommending various products and serving millions of consumers [2]. However, many of the recommender systems are developed for recommending inexpensive and frequently purchased products like books, movies and music. Many of the systems that are currently available for searching infrequently purchased products like cars or houses only provide a standard matching-based search function, whereby the system retrieves products that match exactly with the user's query. This query is normally short and does not reflect the user requirements fully. In

addition, many users do not have much knowledge about the products, thus, they cannot provide detailed requirements of the product attributes or features. Therefore, a recommender system that can predict users' preferences from the initial input given by the users is needed for recommending infrequently purchased products.

Many of the current recommendation systems are developed using a collaborative filtering (CF) approach [2][3][4]. The collaborative filtering approach utilizes a large amount of ratings data or users' previous purchase data to make meaningful recommendations. This approach is not suitable for recommending infrequently purchased products because there is no previous users' purchase history or explicit ratings data about the available products, as the products are not often purchased by the users during their lifetime, and users are not able to provide ratings for products they never use. Fortunately, with the popularity of e-commerce applications for selling products on the web, users are given more opportunity to express their opinion on products they previously owned via the online merchant websites and, as a result, more and more users share reviews concerning their experience with the products. These reviews provide valuable information that can be used by recommender systems for recommending infrequently purchased products.

This paper proposes a recommender system approach that utilizes knowledge extracted from user reviews for recommending infrequently purchased products. Opinion mining and rough set association rule mining are applied to extract knowledge from the user review data to predict a user's preferences. The knowledge about user's preferences is used to expand a user's query to improve the recommendation result.

The following sections of this paper are organized as follows. First, the related work will be briefly reviewed in section 2. Then, the proposed approach will be discussed in section 3. The experimental results and evaluation will be discussed in section 4. Finally, the conclusion will be given in section 5.

2 Related Work

Recently, automatic review mining and summarization of extracting product features values from user reviews is becoming a popular research topic [5][6][7]. Review mining and summarization, also called opinion mining, aims at extracting product features on which the reviewers express their opinion and determining whether the opinions are positive or negative [7]. [5] proposed a model of feature-based opinion mining and summarization, which uses a lexicon-based method to determine whether the opinion expressed on a product feature is positive or negative. The opinion lexicon or the set of opinion words used in this method is obtained through a bootstrapping process using the WordNet. Then, [6] proposed a technique that performs better than the previous methods by using the holistic lexicon-based approach. This technique deals with context dependent opinion words and aggregating multiple opinion words in the same sentence, which are the two main problems of the existing techniques.

Despite the growth in the number of online reviews and the valuable information that they can provide, not much work has been done on utilizing online user reviews for creating recommendations [4]. [8] employed text mining techniques to extract useful information from review comments and then mapped the review comments into the ontology's information structure, which is used by the recommender system

to make recommendations. In their approach, users must input the features of the product that are most important to them and the recommendations are generated based on the features provided by the users. In contrast, our approach aims to predict users' preferences about the product features from the initial input given by them, and use the knowledge to recommend products to the users. The following section will discuss the proposed approach in detail.

3 Proposed Approach

User reviews contain written comments expressed by previous users about a particular product. Each comment contains a user's opinion or how the user feels about the product's features (e.g. good or bad). Opinion mining techniques are applied on user reviews to determine each user's sentimental orientation towards each feature, which indicates whether the user likes or dislikes the product in terms of this feature. The overall orientation of each review is also determined to summarize whether a user's opinion about the product is positive, negative or neutral. The user's opinions generated from the reviews reflect their viewpoint concerning the quality of the products. A review with a positive orientation indicates that the reviewer (i.e. the user) was satisfied with the product in some aspects. This means that at least some attributes of this product were attractive to the user. If we can identify these attractive attributes for each product, based on these attributes we can determine the products that will be of most interest to the user. Based on this idea, we propose to apply association rule mining techniques to generate patterns and association rules from users' positive reviews. By using the extracted patterns and association rules for a target user, we can predict the user's preferred product attributes and, thus, recommend products that best match the user's preferences.

The proposed recommender system approach contains three main processes: i) Opinion mining to extract a user's sentimental orientations to the product features from the user online reviews, summarizing and presenting the reviews in a structured format, ii) Rough set association rule mining to generate association rules between the product attribute values, and iii) Query expansion to expand a user's query by using association rules between product attribute values. The following sections will provide the definitions of the concepts and entities involved and the specific problems of this research. In addition, they will also explain each process in detail.

3.1 Definitions

This section first defines the important concepts and entities used in this paper and then highlights the specific problems that we aim to solve.

- **Product**

Products include any type of product or online service for which users can search for information or purchase. This paper focuses particularly on infrequently purchased products such as cars or houses. A product p can be represented by two-tuple (C, F) , $C = \{c_1, c_2, \dots, c_n\}$ is a set of attributes representing the technical characteristics of the

product defined by domain experts and $F = \{f_1, f_2, \dots, f_m\}$ is a set of usage features representing the usage performance of the product defined by domain experts or the users of the product. The usage features are usually the aspects commented upon by the users of the product. In this paper, we assume that both the product attributes and usage features have been specified. For example, for the online car search domain on which we conducted our experiments, the following car characteristics and usage aspects were chosen as the car attributes and usage features:

$C = \{Make, Model, Series, Year, Engine Size, Fuel System, Fuel Consumption, Tank Capacity, Power, Torque, Body Type, Seating Capacity, Standard Transmission, Drive, Turning Circle, Kerb Weight, Dimension, Wheelbase\}$

$F = \{Comfort Practicality, Price Equipment, Under Bonnet, How Drives, Safety Security, Quality Reliability, Servicing Running Costs, Aesthetics Styling\}$

• User Reviews

For a product, there is a set of written reviews about the product given by users. Each review consists of a set of sentences comprised of a sequence of words. In many e-commerce websites, the product features to be reviewed have been specified so that users can provide their comments and opinions on each particular feature. For reviews that are not classified according to any specific feature, opinion mining techniques can be used to identify the product features that are addressed by each sentence in a review [5]. In this paper, we assume that the sentences in each review have been divided into groups, each of which consists of the sentences that talk about one feature of the product. Let $R = \{R_1, R_2, \dots, R_m\}$ be a review given by a user to a product, R_i is a set of sentences that are comments concerning feature f_i . By applying opinion mining techniques, which will be discussed in the next section, we can generate the user's sentimental orientation concerning each feature, denoted as $O = \{o_1, o_2, \dots, o_m\}$ and an overall orientation of the review O_{all} , where $o_i, O_{all} \in \{positive, negative, neutral\}$.

• Structured review

A structured review is a 4-tuple consisting of the sentimental orientations to a product generated from a review and the product's attributes and features, denoted as $sr = (C, F, O, O_{all})$, where C and F are the attributes and features of the product, O and O_{all} are the sentimental orientations to the features and the overall orientation of the review, respectively. Let $SR = \{sr_1, sr_2, \dots, sr_{|SR|}\}$ be a set of all structured reviews.

• Information System

Information system, I contains 2-tuple of information, denoted as $I = (U, A)$, where U is a set of objects, and A is a set of attributes for each object. In this paper, U is a set of structured reviews and A consists of the product attributes, features, the sentimental orientations to the features and the overall orientation of the review, i.e. $A = \{c_1, \dots, c_n, f_1, \dots, f_m, o_1, \dots, o_m, O_{all}\}$.

The problems that we aim to solve are as follows:

- i) Given a user review R on a product p , the review has to be summarized and represented in a structured review sr_i . Then from a set of all structured reviews SR , an information system I has to be generated using only reviews sr_i that have a positive or neutral overall orientation $O_{all} \in \{positive, neutral\}$.
- ii) From the information model I , a set of association rules between product attribute values c_i has to be extracted using rough set association rule mining to represent users' preferences.
- iii) To develop a query expansion technique by utilizing association rules extracted from information model I , to retrieve products that best meet the users' preferences.

3.2 Opinion Mining

We adopted the approach proposed by [5] to perform the opinion mining process. The first task in this process is to identify the sentimental orientations concerning the features. A user expresses a positive, negative or neutral opinion o_i , on each feature f_i , in a review R_i using a set of opinion words $W = \{w_1, w_2, \dots, w_n\}$. To find out opinion words used by the user (e.g. good, amazing, poor etc.) that express his/her opinion on a product feature f_i , all adjectives used by the user in a review R_i are extracted. The orientation of each opinion word $ow_i \in \{negative, positive\}$ is then identified by utilizing the adjectives synonym set and antonym set in WordNet [9]. In WordNet, adjectives share the same orientation as their synonym and opposite orientations as their antonyms. To predict the orientation ow_i of a target adjective word w_i , a set of common adjectives with known orientation $S = \{s_1, s_2, \dots, s_n\}$ called as seed adjectives, and WordNet are searched to find the word's synonym or antonym with the known orientation. If the word's synonym is found, the word's orientation is set to the same orientation as its synonym and the seed list is updated. Otherwise, if the word's antonym is found, the word's orientation is set to the opposite of the antonym and is added to the seed list. The process is repeated for the target words with unknown orientation and the words' orientations are identified using the updated seed list. Finally, the sentimental orientation o_i of each feature f_i is identified by finding the dominant orientation of the opinion words in the sentence through counting the number of positive opinion words $ow_i \in \{positive\}$ and the negative opinion words $ow_i \in \{negative\}$, for a review R_i . If the number of positive opinion words is more than the negative opinion words, the orientation o_i of the feature f_i is positive $o_i \in \{positive\}$, otherwise negative $o_i \in \{negative\}$. If the number of positive opinion words equals the negative opinion words, the orientation o_i of the feature f_i is neutral $o_i \in \{neutral\}$.

Finally, opinion summarization is performed to determine the overall orientation of each review R and represent the review in a structured review $sr_i = (C, F, O, O_{all})$.

O_{all} is determined by calculating the number of positive features $o_i \in \{positive\}$, neutral features $o_i \in \{neutral\}$, and negative features $o_i \in \{negative\}$ for the review. If the number of positive features and neutral features is more than negative features, the overall orientation O_{all} for the review is positive $O_{all} \in \{positive\}$, otherwise negative $O_{all} \in \{negative\}$. If the number of positive features and neutral features is equal to the negative features, the overall orientation O_{all} for the review is neutral $O_{all} \in \{neutral\}$.

3.3 Rough Set Association Rule Mining

Standard online product search engines perform a match process to find products that satisfy a user's query, which usually consists of the product attributes or characteristics that the users are looking for. However, many users do not have sufficient knowledge about the product and may not know the exact product attributes. Therefore, the attributes in the query may not be the right attributes to query. Online user reviews are provided by users who have used the product and the opinions about the product reflects the users' viewpoints concerning the product based on their experience of using the product. The products that are positively reviewed must possess attractive attributes or characteristics that pleased their users. Based on this intuition, we propose to find the associations between the product attributes from the users' reviews that have a positive orientation. These associations can be used to predict users' preferences to product attributes. In this paper, we utilize the rough set association rule mining approach [10] to find hidden patterns in data and generate sets of association rules from the data. We chose the rough set association rule mining technique because it allows us to easily select the condition and decision attributes of the rule.

Rough set data analysis starts from a data set that is also called a decision table or an information system. In the table, each row represents an object, each column represents an attribute, and entries of the table are attribute values. An attribute can be a variable or an observation or a property, etc. As we have defined above, an information system is written as $I = (U, A)$, in this paper, U is a set of structured reviews and A consists of the product attributes, features, and the sentimental orientations, i.e. $A = \{c_1, \dots, c_n, f_1, \dots, f_m, o_1, \dots, o_m, O_{all}\}$.

In this paper, the information system $I = (U, A)$ is created from the structured reviews with positive/neutral orientation. Let $sr \in SR$ be a structured review, $sr(a)$ be the value of attribute $a \in A$, $U = \{sr \mid sr \in SR, sr(O_{all}) \in \{positive, neutral\}\}$ is the set of objects in the table. The information system contains attribute values for a set of products that have received good comments from the reviewers.

The next step in rough set association rule mining is to partition the information system into two disjointed classes of attributes, called condition C and decision D attributes. The information system is then called a decision table $S = (U, C, D)$, where C and D are disjointed sets of conditions and decision attributes, respectively. The condition and decision attributes are selected from product attributes C and features F in A in the information system I . The attributes chosen as the condition are the product attributes or features that are usually provided by a user as the initial input in a query and the decision contains other attributes and features of the products. For

example, for the online car search on which we conducted our experiments, the car make, model, price, etc are chosen as the condition. Then, association rules are generated from the decision table through determining the decision attributes values based on the condition attribute values. The association rules between the product attributes values show the relationship between the initial product attribute values given by a user with other product attributes in which the user may be interested. Thus, these association rules can be used to represent the user's preferences to retrieve products that will most likely fit the user's requirements.

3.4 Query Expansion

The query expansion process aims to improve the initial user's query in order to retrieve more products that might fit the user's requirements. A user's query Q is represented by a set of terms $Q = \{q_1, q_2, \dots, q_n\}$ that the user provides to the search engine.

In the product search, the terms in the query are attribute values of the product that the user is looking for. The query, generally, is very short and lacks sufficient terms to present the user's actual preferences or needs. Query expansion involves adding new attribute values $E = \{e_1, e_2, \dots, e_n\}$ to the existing search terms $Q = \{q_1, q_2, \dots, q_n\}$ to generate an expanded query $EQ = \{eq \mid eq \in (E \cup Q)\}$. The attribute values $eq_i \in EQ$ are used to retrieve products to recommend to the user. All products that have attribute values that match with any attribute values of the expanded query eq_i are selected as the candidate products $CP = \{cp_1, cp_2, \dots, cp_n\}$. The similarity between each product $cp_i \in CP$ and the expanded query EQ is calculated by matching each attribute value of the product cp_i with the value of the same product attribute eq_i in the expanded query. The similarity value v_i for the product cp_i with the expanded query EQ is calculated as the total number of attribute values of the product that match with the attribute values in the expanded query. Then, the products are ranked based on the similarity value. The top-N products are recommended to the user based on their ranking. The proposed system may retrieve products that exactly match the user's input, as well as other products in which the user may be interested by predicting the user's preferences from his/her initial input.

4 Experiment and Evaluation

4.1 Experiment Method

A case study was conducted for the cars domain. Data was collected from one of the car selling websites that contains reviews provided by users for cars previously owned by them. The dataset contains 5,504 reviews and 3,119 cars. Online reviews on cars from previous users were used to extract rules between attribute values. The opinion mining technique was first applied to generate structured reviews from the online reviews. Then ROSETTA [11], a rough set mining tool was used for extracting rules from the information system generated from the structured review. Four processes were involved in extracting rules i) Data pre-processing and attribute selection, which

included preparing a dataset and selecting important attributes to be included in the decision tables, ii) Data categorization, which involved transforming the attribute values into categorical values to reduce the dimensionality of the data and to reduce the number of rules generated, iii) Selection of decision and condition attributes and iv) Rule induction, which generated rules representing associations between query terms (e.g. car make and model) and other product attributes (e.g. series year, price new, engine size, fuel system etc.) from the decision table. The example of the rule is shown below:

CarMake(Toyota), CarModel(Camry) -> Year(>2000_and<=2005), Price(>30000_and<=50000), EngineSize(>1.6L_and<=3.0L), Seat(4_5), BodyType(SEDAN), Drive(FWD), FuelSystem(FUEL_INJECTED), FuelConsumption(>9.0L_and<=11.5L), TankCapacity(>51L_and<=70L), StandardTransmission(4A), Power(>82Kw_and<=146Kw), Torque(>150Nm_and<=284Nm), TurningCircle(>11.20m), Wheelbase(>2710mm), KerbWeight(>1126Kg_and<=1520Kg), Dimension(>4794)

Two search techniques were developed - the Standard Matching-based Search (SMS) and the Query Expansion Matching-based Search (QEMS). The SMS technique retrieves cars that match with a user's query terms exactly. In addition, the QEMS technique retrieves cars based on the expanded query of the user's preferences predicted from the rules generated from rough set association rule mining process.

For evaluating the proposed approach, previous users' navigation history from the existing system log data was used as a testing data. A sequence of cars viewed by each user was generated from the log data. The first car in each user's sequence was chosen as the input and some of the attributes of the car, such as car make and car model, were considered as the query for that user. The other cars in the sequence were considered as the cars that the user is interested in and they were used as the testing cars to test whether the two search engines recommend these cars to the user. For each query (i.e. the first car of a user), cars recommended by both systems for a different number of the top retrieved results (N=10, 20, 30, 40 and 50) were compared with the testing cars for that user. The recall and precision values for each user were calculated for both techniques by using the following formula:

$$recall = \frac{NM}{NT} \quad \text{and} \quad precision = \frac{NM}{NR}$$

Where NM is the number of cars retrieved that match with the testing cars, NT is the number of testing cars, and NR is the number of retrieved cars. Finally, the average recall and precision values for all users were calculated for both techniques.

4.2 Results

The graphs in Figure 1 and Figure 2 show the evaluation results of the proposed approach. The evaluation results show that the Query Expansion Matching-based Search (QEMS) outperformed the Standard Matching-based Search (SMS), in that this approach can retrieve more car models that the users are interested in without requiring much effort from them. The expanded query can improve the retrieval performance of the Standard Matching-based Search as it provides more keywords to represent a user's requirements or preferences.




Fig. 1. Precision for different number of top retrieved results of the SMS and QEMS




Fig. 2. Recall for different number of top retrieved results of the SMS and QEMS

5 Conclusion

We have proposed a recommender system approach for recommending infrequently purchased products by utilizing user reviews data. The evaluation result shows that our recommendation approach leads to recommendations novelty or serendipity, where more unexpected or different items that meet the users' interests will be recommended to the users. This approach is able to predict a user's preferences and may suggest more products that fit the user's requirements and, also, may help online vendors promote their products. In future work, we intend to utilize sentimental orientations of the features for improving the product recommendations.

Acknowledgments. "This paper was partially supported by the Smart Services CRC (Cooperative Research Centres, Australia)".

References

1. Schafer, J.B., Konstan, J., Riedl, J.: E-commerce Recommendation Applications. *Data Mining and Knowledge Discovery* 5(1-2), 115–153 (2001)
2. Leavitt, N.: Recommendation Technology: Will It Boost E-Commerce? *Computer Society* 39(5), 13–16 (2006)

3. Sarwar, B., Karypis, G., Konstan, J., Riedl, J.: Analysis of Recommendation Algorithms for E-commerce. In: 2nd ACM Conference on Electronic Commerce, pp. 158–167. ACM, New York (2000)
4. Aciar, S., Zhang, D., Simoff, S., Debenham, J.: Recommender System Based on Consumer Product Reviews. In: 2006 IEEE/WIC/ACM International Conference on Web Intelligence, pp. 719–723. IEEE Computer Society, Washington (2006)
5. Hu, M., Liu, B.: Mining and Summarizing User Reviews. In: Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 168–177. ACM, New York (2004)
6. Ding, X., Liu, B., Yu, P.S.: A Holistic Lexicon-based Approach to Opinion Mining. In: International Conference on Web Search and Web Data Mining, pp. 231–240. ACM, New York (2008)
7. Zhuang, L., Jing, F., Zhu, X.Y.: Movie Review Mining and Summarization. In: 15th ACM International Conference on Information and Knowledge Management, pp. 43–50. ACM, New York (2006)
8. Aciar, S., Zhang, D., Simoff, S., Debenham, J.: Informed Recommender: Basing Recommendations on Consumer Product Reviews. *IEEE Intelligent Systems* 22(3), 39–47 (2007)
9. Miller, G., Beckwith, R., Fellbaum, C., Gross, D., Miller, K.: Introduction to WordNet: An Online Lexical Database. *International Journal of Lexicography (Special Issue)* 3(4), 235–312 (1990)
10. Pawlak, Z.: Rough Sets and Intelligent Data Analysis. *Information Science* 147(1-4), 1–12 (2002)
11. Øhrn, A.: ROSETTA Technical Reference Manual. Department of Computer and Information Science, Norwegian University of Science and Technology (NTNU), Trondheim, Norway (2000)

A Machine Learning Approach to Predicting Winning Patterns in Track Cycling Omnium

Bahadorreza Ofoghi^{1,2}, John Zeleznikow¹, Clare MacMahon¹, and Dan Dwyer²

¹ Victoria University, Melbourne VIC 3000, Australia

² Victorian Institute of Sport, Melbourne VIC 8006, Australia

Abstract. This paper presents work on using Machine Learning approaches for predicting performance patterns of medalists in Track Cycling Omnium championships. The omnium is a newly introduced track cycling competition to be included in the London 2012 Olympic Games. It involves six individual events and, therefore, requires strategic planning for riders and coaches to achieve the best overall standing in terms of the ranking, speed, and time in each individual component. We carried out unsupervised, supervised, and statistical analyses on the men's and women's historical competition data in the World Championships since 2008 to find winning patterns for each gender in terms of the ranking of riders in each individual event. Our results demonstrate that both *sprint* and *endurance* capacities are required for both men and women to win a medal in the omnium. Sprint ability is shown to have slightly more influence in deciding the medalists of the omnium competitions.

1 Introduction

The *Omnium* is a track cycling competition that has newly been introduced as one of the events in the Olympic Games by the International Cycling Union (UCI) [1]. The main motivation behind inclusion of the omnium in the Olympic Games is to create a better balance between the number of cycling competitions for men and women in the Olympics. In the Beijing 2008 Olympics, there were seven cycling competitions for men (team sprint, sprint, keirin, 4000-meter team pursuit, 4000-meter individual pursuit, 50-kilometer Madison, and 40-kilometer points race) whereas only three events were held for women (sprint, 3000-meter individual pursuit, 25-kilometer points race). According to the agreement made between UCI and the International Olympic Committee (IOC), in the London 2012 Olympics, the cycling competitions will involve five events per gender including the sprint, keirin, team sprint, team pursuit, and omnium.

The omnium was first included in men's cycling competitions at the UCI Track Cycling Championships in Spain in 2007. Primarily, the omnium included five individual events:

- Flying time trial, where cyclists compete in a short flying lap that has traditionally been used for qualification in other cycling events (e.g. individual pursuit).

¹ <http://www.uci.ch>

- Scratch race, in which all contestants start from a start point at the same time and need to complete a certain number of laps. If a rider gains an entire lap ahead of the other riders, she/he will have an advantage over the riders who have completed fewer laps.
- Individual pursuit, where two riders start a race from opposite sides of the track on the pursuit line at the bottom of the track. The riders start at the same time and both must complete the race distance to record a time for the ranking.
- Points race, which is a mass start event involving a large number of riders on the track at the same time. This is a race over a long distance (about 120 to 160 laps). Every 10 laps, a sprint is held and the top four finishing riders are awarded 5, 3, 2, and 1 points respectively. Any rider who can complete a lap is also awarded 20 points for each lap. In the end, the rider with the most points is the winner of the event.
- (Kilometer) Time trial, where riders compete against the clock to secure the fastest time. Riders are sent out individually in set time intervals. This event includes cycling in moderately longer distances than the flying time trial event.

Each rider in the omnium scores according to their rank in the individual events and the rider with the lowest score wins the competition.

In December 2009, UCI announced new changes to the omnium competition that will take place for the first time in the Olympic Games in London in 2012. The new changes were introduced to bring the balance back towards the endurance riders. According to UCI, the new omnium competition will now involve six events with the following format:

- Flying time trial: 250 meters for men and women
- Scratch race: 15 kilometers for men and 10 kilometers for women
- Individual pursuit: 4 kilometers for men and 3 kilometers for women
- Points race: 30 kilometers for men and 20 kilometers for women
- (Kilometer) Time trial: 1 kilometer for men and 500 meters for women
- Elimination race: 24 riders will compete in a race where every two laps, the last rider over the finish line will be eliminated until only a single rider remains and is decided as the winner.

The elimination race has never been part of the omnium, which creates an unknown factor in regards to predicting the overall winner. The omnium competition with its new format will be held over two days. The sequence of the events in each day has yet to be announced by UCI.

2 Related Work

Because of the nature of multiple-component competitions comprised individual events, it is difficult to determine what strategies need to be employed by coaches and athletes in order to maximize the chances of winning the whole competition.

In similar sports to the omnium such as the *Decathlon* (with 10 individual track and field events), overall success can be attributed to different performances for each individual component. Kenny et al. [1] argue that excellence in any individual event at the decathlon competitions may be achieved at the expense of finishing the whole competition with a poorer overall standing. Cox and Dunn [2], however, used statistical analysis on the data collected at the meetings of the International Amateur Athletic Federation (1991 – 1999) and concluded that the decathlon favors athletes who are proficient in the field events.

Zwols and Sierksma [3] studied training data from the decathlon. They made use of mathematical optimization methods to find the best combinations of training times/sessions for performing well in different types of championships/competitions. As they also make the point, it is, however, not straightforward to relate training performances to competition performances as there are other factors (such as mental/psychological preparedness) that can potentially interfere with performance outcomes.

Much of the previous work in the domain of sports data analysis is based on the statistical evidence. Although statistical analysis is a good approach to make better sense of sports data, it should not be taken as the only source of information that can be extracted from the data. Data mining and machine learning-based approaches have also been used widely in this domain. Examples include modeling and predicting competitive performances in swimming using Neural Networks [4], using Self Organizing Maps for classifying coordination patterns of participants in different types of basketball shots [5], cluster analysis of elite swimmers' race patterns [6], predicting player injuries in soccer using data mining [7], and knowledge discovery techniques implemented in *Advanced Scout* software used for finding series of consecutive events that may lead to scoring in NBA basketball games [8].

In this paper, we also employ machine learning-based approaches and try to confirm our results on the strategic planning for winning a medal in the omnium competitions by conducting a further complementary statistical analysis.

3 Research Problems

In terms of the omnium, especially with its new format, there is little to no previous research to the best of our knowledge. With this event to be included in the Olympics for the first time, it is, however, very useful for both coaches and athletes to gain a better understanding of the strategies to be employed and assist them towards winning the competition.

In particular, for both male and female cyclists, we would like to understand:

1. What overall ranking in each individual event is required to win a medal in the omnium?
2. What type of rider may have a greater advantage with regards to winning a medal in the omnium? Those with *sprint* or *endurance* capacity?
3. What individual events have more influence on the final standing in the omnium?

We used a number of unsupervised and supervised learning mechanisms as well as statistical analysis methods to answer these questions and assist coaches and riders with a better strategic plan for winning a medal in the omnium. We did not attempt at developing new artificial intelligence methods for this; instead, tried to see how existing methods can be used/adjusted in this context.

4 Data Collection

Since omnium competitions were only introduced to the World Championships after 2006, there is only presently a limited dataset available. The data that we have collected consist of the competition results for the following events:

- The 2009 Track Cycling World Championships (women, men)
- The 2009 Australian Track Championships (women, men)
- The 2009 Juniors Track Cycling World Championships (women, men)
- The 2008 Juniors Track Cycling World Championships (men)

The dataset includes 96 data records for men and 75 data records for women. Since there is inconsistency in the dataset in terms of the class of the riders (i.e., masters and juniors), it is not viable to conduct a study on non-relative attributes such as time, speed, and score for each individual event. Therefore, we only considered the ranking/placing of the riders in each individual event.

On the other hand, since the elimination race is a totally new individual event in the omnium, previous data/results do not involve any information regarding this event in the context of the omnium.

Therefore, our dataset only includes the ranking of the cyclists at the individual events excluding the elimination race and the overall final standing of the riders in omnium competitions.

5 Empirical Analysis

In order to answer the questions raised in section 3, we first pre-processed the dataset. Pre-processing included two steps:

1. Converting the data to the format that can be dealt with by the machine learning package that we use (the *weka* machine learning package²). This mainly consists of replacing missing values with a character (the question mark) and inserting the header information at the beginning of the data file.
2. Categorizing/generalizing the final standings in the omnium in to pre-defined classes:
 - 1: Any medal winners including the final standings 1st, 2nd, and 3rd,
 - 2: Non-medal winners ranked between 4 and 10 inclusive, and
 - 3: Non-medal winners ranked above 10.

After pre-processing was completed, we carried out different experiments to find answers to the research problems already mentioned.

² <http://www.cs.waikato.ac.nz/ml/weka/>

5.1 Experiment 1 – Unsupervised Learning of Rankings

The unsupervised learning experiment included a cluster analysis of the data records collected for the riders in different competitions. We used the *K-Means* clustering algorithm implemented in weka and set the number of clusters to 3. The three clusters were selected on the basis of the three categories of the final standings. We ran the K-Means algorithm on the men’s and women’s data separately. Table 1 and Table 2 summarize the results for each gender.

Table 1. Unsupervised (clustering) learning of rankings for each individual event on the *men’s* competition results. Numbers in parentheses represent the number of instances/records in each cluster. Within cluster sum of squared errors = 16.1598.

Attribute	Full data (96)	Cluster 1 (21)	Cluster 2 (15)	Cluster 3 (60)
PR_rank	7.8676	11.0476	3.8667	7.7549
TT_rank	8.6774	11.4808	5.5828	8.4699
IP_rank	8.0000	13.0000	3.1333	7.4667
SR_rank	9.0588	11.5266	7.2902	8.6373
FTT_rank	8.2877	13.0137	3.8667	7.7388
final_standing	2.0000	3.0000	1.0000	2.0000

PR= Points Race, TT= Time Trial, IP= Individual Pursuit, SR= Scratch Race, FTT= Flying Time Trial.

The column “Full data” represents the average values (means) of each attribute over the full set of data, whereas the other columns show the average values over the instances in each cluster.

From the results in both Table 1 and Table 2, it is not surprising that the average final standing under the column “Full data” is equal to 2 since most of the data records have the final standing of category 2 (finishing in position 4 to 10). It can also be seen that the clustering algorithm has correctly detected three clusters for men’s and women’s data where the final standings cover the three categories of final standings that we set. This is an important observation since it shows that there is a pattern of the rankings in individual events that correspond to each final standing category. In other words, the clustering algorithm has been able to bring together data points with similar final standings in certain groups.

For men, the results in Table 1 suggest that in order for a rider to win a medal (final_standing=1, Cluster 2), it is required to achieve the rankings of approximately 4, 5, 3, 7, and 4 in the points race, time trial, individual pursuit, scratch race, and flying time trial events respectively. In general, the flying time trial and individual pursuit components require strong sprint capacity, whereas the scratch race and time trial events are dealt with much better by riders with endurance capacity. The points race event requires a mixture of sprint and endurance abilities. Therefore, our results suggest that in order for a male rider to win a medal in the omnium it is required that he has a combination of sprint and endurance power without either being an outstanding sprinter or having excellent endurance ability. In fact, the endurance capacity can be neglected to

Table 2. Unsupervised (clustering) learning of rankings for each individual event on the *women’s* competition results. Numbers in parentheses represent the number of instances/records in each cluster. Within cluster sum of squared errors = 11.4309.

Attribute	Full data (75)	Cluster 1 (14)	Cluster 2 (46)	Cluster 3 (12)
PR_rank	7.5577	10.8255	7.5973	3.5833
TT_rank	8.0000	13.2857	6.9184	6.2500
IP_rank	7.9123	13.0714	7.4168	3.9167
SR_rank	8.7576	11.1082	8.6481	6.4621
FTT_rank	8.0847	13.8571	7.6603	3.0833
final_standing	2.0000	3.0000	2.0000	1.0000

some degree as the results show that the ranking of the rider in the scratch race can be as low as about 7 while winning a medal.

For women, the results in Table 2 show, however, in order to win a medal, more sprinting capacity is required with only moderate endurance power. This is a result of the average rankings under Cluster 3 with final_standing=1. A ranking between 3 and 4 in the points race, individual pursuit, and flying time trial events in addition to final rankings of about 6 in the time trial and scratch race components can potentially lead to winning a medal for a female cyclist.

5.2 Experiment 2 – Supervised Learning of Rankings

To see how accurate the predictions in Experiment 1 are, for both men and women, we conducted a supervised learning experiment. The supervised learning experiment consists of a classification task that classifies each instance/data record into one of the predefined categories of the final standing explained in section 5. For this, we used the *Naive Bayes* classifier implemented in weka and conducted a 10-fold cross validation process for testing the classification accuracies. In this experiment, we did not consider any inter-dependence between the different attributes (i.e. rankings in each individual event). Table 3 and Table 4 show the results of the classification task for men and Table 5 and Table 6 summarize the results of the same analyses for women.

For men, the results in Table 3 demonstrate that for instance if a rider finishes ranked 4, 3, 3, 6, and 4 in the points race, time trial, individual pursuit, scratch race, and flying time trial events respectively, then there is a chance of approximately 84% that he/she can win a medal in the omnium competition. It is promising to see that the trend of the results of the supervised learning analysis matches that of the results of the unsupervised learning analysis in terms of the average ranking in each individual event. To illustrate this, one might compare the average rankings under Cluster 2 in Table 1 with the mean values under final_standing=1 in Table 3.

The supervised analysis on the male athlete data again shows that both sprint and endurance capacities are required to win a medal with a slightly lesser endurance power than sprint ability. According to our analysis, finishing ranked

Table 3. Supervised (classification) learning of rankings for each individual event on the *men's* competition results. Numbers in parentheses represent the percentage of instances/records in each class.

Attribute		final_st.=1 (0.22)	final_st.=2 (0.47)	final_st.=3 (0.31)
<i>PR_rank</i>	mean	3.8667	7.6563	11.0476
	std. dev.	2.1868	4.5662	3.9094
<i>TT_rank</i>	mean	2.8750	7.7778	13.5833
	std. dev.	1.6154	2.9355	1.8465
<i>IP_rank</i>	mean	3.1333	6.7500	13.0000
	std. dev.	2.0613	3.0822	3.2514
<i>SR_rank</i>	mean	6.1111	7.4286	11.6500
	std. dev.	4.8177	4.0891	4.2694
<i>FTT_rank</i>	mean	3.8667	7.5625	13.2500
	std. dev.	2.4459	3.4635	3.6038
Correctly classified instances			58	84.0580%
Incorrectly classified instances			11	15.9420%

Table 4. Classification measures for supervised learning of rankings for each individual event on the *men's* competition results

Class	Precision	Recall	F-measure	ROC area
final-standing=1	0.833	0.667	0.741	0.927
final-standing=2	0.789	0.909	0.845	0.886
final-standing=3	0.947	0.857	0.900	0.959
<i>Weighted average</i>	<i>0.847</i>	<i>0.841</i>	<i>0.839</i>	<i>0.917</i>

about 6 in the scratch race is sufficient for winning the omnium given the rankings in the other events mentioned in the previous paragraph.

The results in Table 4 show the accuracy level of these arguments in more detail. The most accurate predictions (with an accuracy of approximately 95%) for where a competitor will place in the omnium, based on performance in the separate events, can be made for those who finish outside the top ten (final-standing=3). For the medal winners, the accuracy is slightly lower ($\sim 84\%$) and the lowest accuracy corresponds to the non-medal winners who may finish the competition with a final standing between 4 and 10 inclusive ($\sim 79\%$).

Overall, the average precision, recall, F-measure, and ROC area are all above 83% and this is a promising result emphasizing the correctness of concluding remarks drawn from the unsupervised analysis for men (see section 5.1).

For women, the results summarized in Table 5 suggest similar conclusions to those drawn from the men's dataset. Once again, the only major difference is the higher average rank (5.00) required for the time trial component compared to that for men (2.8750). As with the unsupervised analysis, this suggests that, for women, there is more bias towards sprinting capacity to win the omnium.

Table 5. Supervised (classification) learning of rankings for each individual event on the *women’s* competition results. Numbers in parentheses represent the percentage of instances/records in each class.

Attribute		final_st.=1 (0.24)	final_st.=2 (0.49)	final_st.=3 (0.27)
<i>PR_rank</i>	mean	3.5833	7.2308	11.0769
	std. dev.	3.0675	3.4453	3.9703
<i>TT_rank</i>	mean	5.0000	7.3000	13.6923
	std. dev.	2.5071	3.4799	2.0147
<i>IP_rank</i>	mean	3.9167	6.8846	13.0714
	std. dev.	2.5317	3.5008	3.3480
<i>SR_rank</i>	mean	4.1667	8.6429	11.5000
	std. dev.	2.9107	4.2191	4.2720
<i>FTT_rank</i>	mean	3.0833	7.7308	13.8571
	std. dev.	1.6562	3.2999	2.3561
Correctly classified instances			45	86.5385%
Incorrectly classified instances			7	13.4615%

Table 6. Classification measures for supervised learning of rankings for each individual event on the *women’s* competition results

Class	Precision	Recall	F-measure	ROC area
final-standing=1	0.818	0.750	0.783	0.896
final-standing=2	0.852	0.885	0.868	0.857
final-standing=3	0.929	0.929	0.929	0.980
<i>Weighted average</i>	<i>0.865</i>	<i>0.865</i>	<i>0.865</i>	<i>0.899</i>

The overall accuracy of the classification task is about 86%, slightly higher than that for men. The detailed classification measures shown in Table 6 demonstrate a slightly higher overall precision, recall, F-measure, and ROC area for women.

5.3 Experiment 3 – Statistical Analysis of Individual Events

In order to understand what individual events have more influence on the final standing in the omnium, we carried out a statistical analysis on the relationship between the values of each attribute (rankings) in our dataset and the values of the attribute that represents the final standing of the riders. For this, we used the *Correlation Coefficient* measure r . We back tracked in the pre-processing step (generalizing the final standings to classes 1, 2, and 3) and used the raw final standing values ranging from 1 to the number of riders in each competition.

The results of the statistical analysis for men and women are shown in Figure 1 and Figure 2. For both genders, the highest correlations are between the individual pursuit rank and flying time trial rank and the final standing. Especially in the case of women, this is very much consistent with the results of the unsupervised and supervised learning processes where sprinting capability has




Fig. 1. Correlation coefficient measures between the rankings in each individual event and the final standings of male riders




Fig. 2. Correlation coefficient measures between the rankings in each individual event and the final standings of female riders

been shown to be more important for a female rider to win omnium competitions. For men, this is slightly less consistent, however.

For women, the lowest correlation is between the scratch race rank and the final standing, whereas for men, the least correlation corresponds to the time trial rank and the final standing. In both cases, this emphasizes that endurance power, although required, can play a less important role in winning the omnium.

It is important to note that the results of our statistical analysis using the correlation coefficient r , however, do not show the *causation* relationship but only a correlation relationship that can be used in combination with the results of other experiments to provide strength to the interpretations of the results.

6 Conclusion and Future Work

We studied the newly introduced Olympic track cycling competition, the omnium, to understand what performance, in terms of ranking, in each individual event of this competition is required to win a medal, what type of riders may have the greatest chance of winning a medal (riders with sprint or endurance capacity), and what individual events may have more influence on the final standing in the omnium. For this, we used machine learning and statistical approaches to analyze the competition data available to us so far.

The results of our unsupervised clustering, supervised classification, and statistical correlation coefficient analysis suggest that both sprint and endurance capacities are required for both men and women to win a medal in the omnium. However, especially in the case of women, there can be less emphasis on endurance power and more on sprint ability. This is a result of our machine learning-based analyses and is supported by our statistical analysis that shows high correlations between the rankings of the riders in the sprint-power-requiring individual events (i.e., flying time trial and individual pursuit) and the final standings in the omnium. These results are based on a form of the omnium that

is about to change (to a six-event competition); therefore, our contributions may change when we repeat the analyses on competition results in the future.

Because of the inconsistent nature of our database in terms of speed, time, and scores of individual events, we are trying to collect more related data and repeat our experiments with a larger dataset that includes a greater number of data records for junior and senior men and women cyclists. This will give us the opportunity to carry out analyses on other more abstract attributes time, speed, and scores achieved by riders.

Other current work is analyzing the inter-dependence between different attributes (features of individual events) to construct a probabilistic model to tell us that, given certain ranks, speeds, times, and/or scores in certain individual events, what may be achieved by riders in the other individual events and potentially in the omnium competition.

References

1. Kenny, I.C., Sprevak, D., Sharp, C., Boreham, C.: Determinants of success in the Olympic decathlon: Some statistical evidence. *Journal of Quantitative Analysis in Sports* 1 (2005)
2. Cox, T.F., Dunn, R.T.: An analysis of decathlon data. *Journal of the Royal Statistical Society* 51, 179–187 (2002)
3. Zwols, Y., Sierksma, G.: Training optimization for the decathlon. *Journal of Operations Research* 57, 812–822 (2009)
4. Edelmann-nusser, J., Hohmann, A., Henneberg, B.: Modeling and prediction of competitive performance in swimming upon neural networks. *European Journal of Sport Science* 2, 1–10 (2002)
5. Choi, H.: Cluster analysis of Korean elite athletes' fitness using the self organised map. In: *Proceedings of International Association for Computer Science in Sport*, Australian Institute of Sport, Canberra, Australia (2009)
6. Chen, I., Homma, H., Jin, C., Hong Yan, H.: Identification of elite swimmers' race patterns using cluster analysis. *International Journal of Sports Science and Coaching* 3, 293–303 (2007)
7. Flinders, K.: Football injuries are rocket science. vnunet.com (2002), <http://www.vnunet.com/vnunet/news/2120386/football-injuries-rocketscience>
8. Bhandari, I., Colet, E., Parker, J., Pines, Z., Pratap, R., Ramanujam, K.: Advanced Scout: Data mining and knowledge discovery in NBA data. *Data Mining and Knowledge Discovery* 1, 121–125 (2004)

PART IV
Neural Nets and Fuzzy Logic

Learning Motor Control by Dancing YMCA

Rikke Amilde Løvliid

Department of Computer and Information Science
Norwegian University of Science and Technology
Trondheim, Norway

Abstract. To be able to generate desired movements a robot needs to learn which motor commands move the limb from one position to another. We argue that learning by imitation might be an efficient way to acquire such a function, and investigate favorable properties of the movement used during training in order to maximize the control system's generalization capabilities. Our control system was trained to imitate one particular movement and then tested to see if it can imitate other movements without further training.

1 Introduction

Humanoid robots assisting humans can become widespread only if they are easy to program. This might be achieved through learning by imitation, where a human movement is recorded and the robot is trained to reproduce it. However, to make learning by imitation efficient, good generalization capabilities are crucial. One simply cannot demonstrate every single movement that the robot is supposed to make.

How we want the agent to generalize depends on what we want the agent to do. When watching the demonstrator move, the robot can either learn to mimic the motion of the demonstrator or learn how the demonstrator acts in many situations, that is, extracting the intention of the movement. Mimicking the exact movement trajectory might be important when learning a dance movement, but this is less important when washing the dishes. Belardinell et al. taught a robot to extract salient features from a scene by imitating the gaze shifts of a human demonstrator [1]. Wood and Bryson used observations of an expert playing a computer game to make the agent learn what contexts are relevant to selecting appropriate actions, what sort of actions are likely to solve a particular problem, and which actions are appropriate in which contexts [2].

Learning by imitation is, in a sense, something in between pre programming the agent's control policy (i.e., the function that decides which action to choose in every situation), and letting the agent figure it out on its own through trial and error. According to a hypothesis in developmental psychology, learning to control one's own motor apparatus may be based on so called motor babbling, i.e., random exploration of joint angles [3,4]. Other findings suggest that children use more goal directed movements [5].

We argue that imitation can be used in an efficient way in learning to master the motor apparatus. In this paper we investigate the features of the training movement required to make the suggested control system generalize to new movements, and illustrate how imitation can be used to make the agent train on movements that are most valuable in terms of future generalization capabilities.

2 Feedforward Control

The goal of the work presented here is to make the agent capable of moving its limbs to the positions it desires, that is, we want the agent to learn *feedforward control*. In feedforward control the future desired state is monitored and a motor command is issued to drive the system towards this state. We could call this proactive motor control. The purely reactive alternative is *feedback control*, where the state of the system is compared with the desired state of the system and adjustive motor commands are issued accordingly. Often both feedforward- and feedback control is needed. In our experiments we have used a feedback controller to train the feedforward controller.

We consider feedforward control as a modular process where the control policy, i.e., the function that maps the current state and the future goal to a motor command, is decomposed into a *planning stage* and an *execution stage*. The planning stage generates a desired trajectory. This can be realized by generating the whole desired sequence in advance, or through a next state planner. In the presented work planning is done by the demonstrator, and our focus is on the execution stage.

2.1 Realization of Feedforward Control with Internal Models

There are several ways to realize the execution stage in feedforward control, but most research regarding voluntary motor control shows a tendency towards the use of *internal models*. An internal model is a system that mimics the behavior of a natural process. In control theory, two types of internal models are emphasized, *forward models* and *inverse models*. A forward model predicts the outcome of an action (i.e., motor command). An inverse model represents the opposite process of calculating an action that will result in a particular outcome, the desired next state. Existence of internal models in the brain is widely accepted and there are many theories of how they are used and where they are located [6,7,8].

Forward models, inverse models and feedback controllers can be combined in different ways to calculate the desired motor command [9,10,6,11]. The straightforward approach is to use only an inverse model. Since the input-output function of the inverse model is ideally the inverse of the body's forward dynamics, an accurate inverse model will perfectly produce the desired trajectory it receives as input. To acquire such an accurate inverse model through learning is, however, problematic. Kawato investigated different possibilities [9]. Among these, we have adopted the *feedback-error-learning* scheme, where a simple feedback controller is used together with the inverse model. The details are explained in section 3.2.

2.2 Implementation of the Inverse Model

In our control system, the inverse model was implemented as an *echo state network* (ESN) [12]. The basic idea with ESNs is to transform the low dimensional temporal input into a higher dimensional *echo state* by using a large, recurrent neural network (RNN), and then train the output connection weights to make the system output the desired information.

Because only the output weights are altered, training is typically quick and computationally efficient compared to training of other recurrent neural networks, and also simpler feedforward networks.

A typical task can be described by a set of input and desired output pairs, $[(i_1, o_1), (i_2, o_2), \dots, (i_T, o_T)]$ and the solution is a trained ESN whose output y_t approximates the teacher output o_t , when the ESN is driven by the training input i_t .

Initially, a random RNN with the Echo State property is generated. Using the initial weight matrixes, the network is driven by the provided input sequence, $[i_1, i_2, \dots, i_n]$, where n is the number of time steps. Teacher forcing is used, meaning o_t is used instead of y_t when computing the state of the network at $t + 1$. The state of each node at each time step is stored in a state collection matrix, \mathbf{M} . Assuming \tanh is used as output activation function, $\tanh^{-1}o_t$ is collected for each time step into a target collection matrix, \mathbf{T} .

If \mathbf{W}^{out} is the weights from all the nodes in the network to the output nodes, we want to solve the equation $\mathbf{M}\mathbf{W}^{out} = \mathbf{T}$. To solve for \mathbf{W}^{out} we use the Moore-Penrose pseudoinverse; $\mathbf{W}^{out} = \mathbf{M}^+\mathbf{T}$.

Note that when the desired output is known, the network will learn the input-output function after only one presentation of the training sequence.

The input of the inverse model is the current state together with the desired next state, and the desired output is the desired motor command. The desired motor command is only known indirectly through the desired position of the limbs. Generally, several motor commands may result in the same position of the limbs, and one does not want to bias the controller into choosing one specific solution. In sections 3.2 and 3.5 it is explained how an estimate of the desired motor command is used for teacher forcing and when generating the target collection matrix.

3 Learning the Inverse Model by Imitation

Our agent is implemented as a simple stick-man-simulator. After it has learned to imitate one movement, we want it to be able to imitate any movement presented by the demonstrator without further training. The input to the control system is always the current state and the desired next state (which is provided by the demonstrator). The goal is thus to learn the function mapping the current state and desired next state to the motor command, preferably with minimal effort.

3.1 The Movement Data

In the implementation of the experiments, we used a recording of the dance to the song YMCA by The Village People (se figure [1](#)). The movement data was gathered with a Pro Reflex 3D motion tracking system by Axel Tidemann [\[13\]](#).




Fig. 1. The Movement is described with four angles; the two shoulder abduction angles θ_1 and θ_2 and the elbow extension angles ϕ_1 and ϕ_2

The movement of each arm was described in two degrees of freedom, the angle between the arm and the body, i.e., the abduction angle θ , and the angle between the under- and upper arm, i.e., the extension in the elbow ϕ . Hence, the simulated robot was described by 4 degrees of freedom.

The YMCA movement was represented as a sequence of states, where each state t represents the four desired joint angles at time step t . The movement was manipulated in different ways to generate various training and testing sequences.

The goal of the control system is to produce motor commands that generate a movement where the state generated in each time step is as close as possible to the corresponding state in the desired sequence of states.

3.2 The Architecture

The control architecture that underlie the experiments is shown in figure [2](#). It consists mainly of an inverse model, but to achieve this model, a feedback controller is included during training.

At each time step t the control system receives as input the *desired next state* (i.e., the joint angles at time step $t + 1$ in the training sequence) and the *current state*, and outputs a *motor command*, u .

During the training phase the feedback controller translates analytically the difference between the *desired current state* and the *actual current state* (i.e., state error) to a motor command, u_{error} . This motor error, the error done by the control system in the previous time step, is used to adjust the motor command for the current time step. This works as an approximation to teacher forcing because the only connection from the output nodes back to the network is through the plant, providing the current state input at the next time step. How much the feedback controller is able to influence the motor command depends on the *feedback gain*, K , by letting $u_{feedback} = K * u_{error}$. Note that during testing $K = 0$. The motor command $u_{feedback}$ is added to the motor command produced by the inverse model, and the result is sent to the plant, i.e., the robot simulator.

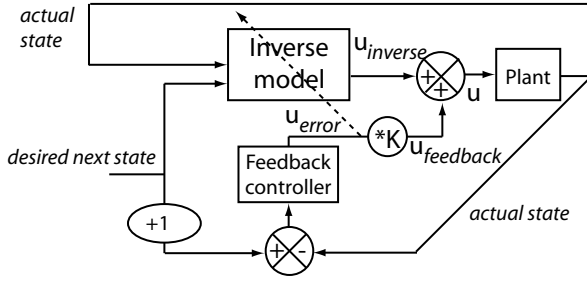


Fig. 2. The control architecture. Delays are shown with ellipses, i.e., the desired next state is delayed one time step, now representing the desired current state, before given as input to the feedback controller. An arrow across a model represents training, and u is a motor command.

3.3 Hypotheses

First, we need to verify that the control system presented is able to imitate novel movements when trained to perform one movement. This would imply that the system has learned at least parts of the function that computes the motor command needed to move the agent from its current position to the desired next position. Second, we investigate further what properties the training movement must possess in order to make the system generalize to any movement in the state space. Our aim can be conveyed through the following tree hypotheses:

Hypothesis 1: *When training on imitating one movement, the control system does not only learn to mimic that movement, but learns at least parts of the function mapping a current and a desired state to a motor command, which will make it able to imitate other movements without training.*

Hypothesis 2: *In order to learn to control one particular degree of freedom, it has to be included in the training movement.*

Hypothesis 3: *When trained on synchronous movements, i.e., the movement of the two arms are equivalent, mirroring each other, the control system is only able to imitate synchronous movements. Training on movements where the limbs follow different trajectories is necessary in order to make the control system generalize to all movements.*

3.4 Experiments

To generate different training and testing sequences the YMCA movement was manipulated in different ways. The movements were **YMCA** (the whole YMCA movement), **Y** (only the Y movement, moving back to start position by reversing the Y motion), **Y pure** (a manipulated version of the Y movement, where all movement in elbow angle is removed), **YM** (only the YM movement, moving back by reversing the YM sequence), **right arm mirror** (both arms does the

movement of the right arm in the YMCA movement, making the arms mirror each other) and **left arm mirror** (similar to *right arm*, but now both arms moves as the left arm in the YMCA movement).

3.5 Training

The training procedure was organized in epochs and cycles, where one cycle is one full temporal presentation of the training motion. In each epoch we ran seven cycles. First, we re-initialized the network by setting the internal states of the network to zero and run one cycle without updating the output weights. Subsequently, the training sequence was presented five times with enabled learning. The output connections were then adapted after each complete cycle. A final cycle was used to estimate the performance error on the training sequence while learning was disabled. The training was run for 150 epochs.

Multiple training epochs was needed because perfect teacher forcing could not be provided. To make the estimate of the desired motor command, u , as good as possible, the feedback controller should provide less influence as the inverse model gets more accurate. This was ensured by decreasing the feedback gain, K , by 10% each epoch.

The output from the feedback controller was also used when calculating the target collection matrix. Because we had batch learning, the motor command was stored and the target motor command u_{target} was calculated by using the u_{error} provided in the next time step, $u_{target}^t = u^t + u_{error}^{t+1}$ (because u_{error}^t reflects the error done at time step $t - 1$).

The inverse model had 8 input nodes, 1000 nodes in the internal layer and 4 output nodes. The ESN had spectral radius $\alpha = 0.1$ (determining the length of the memory with increasing $\alpha \in [0, 1]$) and noise level $v = 0.2$ (effectively adding 10% noise to the internal state of the network). The inputs and outputs where scaled to be in the range $[-1, 1]$. Normal-distributed noise with standard deviation 0.01 was added to the sensory signal from the plant. The system was implemented in MatLab, including the simple stick-man-simulator used as plant.

3.6 Testing

After the inverse model was trained to imitate one movement, we wanted to test whether it could imitate other movements without training. This was done by changing the desired sequence and run the network for one additional epoch with only two cycles, the initialization cycle and the evaluation cycle.

During training the feedback gain was decreased to ~ 0 , and the feedback controller was thus removed from the control system during testing.

4 Results

This section summarizes the results of the experiments. The results verifying hypotheses 1, 2 and 3 are described in section [4.1](#), [4.2](#) and [4.3](#) respectively. We

have illustrated the results by plotting the shoulder abduction angle, θ against the elbow extension angle, ϕ for each arm. The temporal dimension is not plotted because all discrepancies between actual and desired movement could be seen as spatial errors, the timing turned out to not be a problem in any of the experiments. Figure 3 shows a plot of the whole YMCA movement where the different parts of the movement is separated by the use of different markers.

4.1 Does the Control System Generalize?

The initial experiment was to train the control system to imitate the whole YMCA movement and see if it was able to produce the correct movements when tested on the other movements described in section 3.4. We also tested whether the control system would generalize to different speeds. The control system managed all these tests without significantly more error than when imitating the trained movement. We conclude that when trained with the whole YMCA movement, the control systems learned the complete mapping from current- and desired state to motor command in the state space.

4.2 Training All Degrees of Freedom

Does all degrees of freedom, used in the testing sequence, need to be included in the training sequence? To test this hypothesis the control system was trained on *Y pure* and tested on *YM*, see figure 4. The control system is clearly not able to utilize the joint it has not trained to use. To find out how small perturbations in the joint is sufficient for generalization, we tried training on *Y* and testing on *YMCA*. As figure 5 illustrates, small angular changes in the joint during training, makes it possible to generalize to larger changes during testing, but not large enough to perform *YMCA* without large errors.




Fig. 3. The whole YMCA movement. The shoulder abduction angle, θ is plotted against the elbow extension angle, ϕ for each arm. The left graph shows the movement of the right arm and the right graph the movement of the left arm. Each part of the movement is plotted with a different marker to distinguish them from each other.




Fig. 4. The figure shows the trajectory produced by the two arms when the control system is trained on *Y pure*, and tested on *YM*. In the training sequence, *Y pure*, there is no movement in the elbow joints, and thus the network is not able to utilize these during testing.

4.3 Synchronous and Asynchronous Movements

To test whether the system can generalize to asynchronous movements when trained with a pure synchronous movement, we used *right arm mirror* and *left arm mirror* as training sequences and tested on *YMCA* (see figure 6).

The system clearly does not generalize to asynchronous movements; the movement of the arms were more or less symmetric even though the test sequence is not. The movement of each arm was an average of the movements of the two arms in the desired sequence. As a consequence the results are practically identical when the control system was trained with *right arm mirror* compared to when trained with *left arm mirror*.




Fig. 5. The left figure illustrates the *Y* movement. Notice that there is hardly any motion in the elbow. The figure to the right shows the result when the control system trained on movement *Y* is tested on movement *YMCA*. The control system is able to generate more motion in the elbow joints than it learned during training. However, it is not able to produce the *YMCA* movement without large errors.

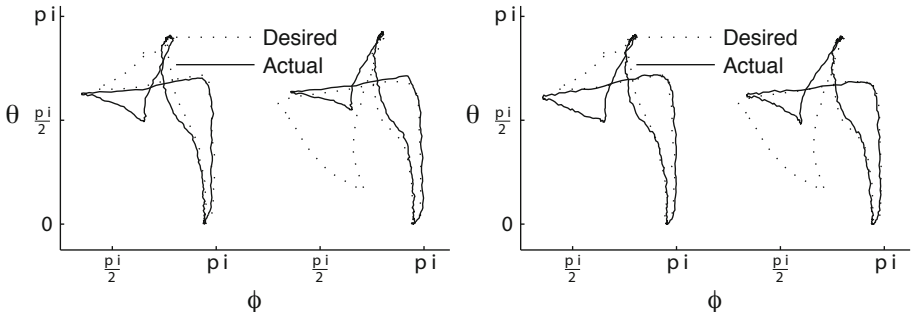


Fig. 6. The figures illustrates the results when trying to do an asynchronous movement when trained to do a synchronous one. The control system produces a synchronous movement that is the average of the desired movement of the two arms. In the figure to the left the system was trained on *left arm mirror* and in the right figure, *right arm mirror*. Both were tested on *YMCA*.

Remember that the opposite problem of imitating a synchronous movement when trained with an asynchronous movement does not pose any problem. When trained with the asynchronous movement *YMCA*, the system was able to generate all the movements without difficulty.

5 Discussion and Conclusion

Our aim was to use imitation to efficiently learn the inverse kinematics model of our robot simulator. We showed that when trained to imitate one movement, the control system has not only learned to imitate that particular movement, but is able to imitate novel movements without further training. This means that the system has learned at least parts of the desired inverse model, verifying hypothesis 1.

Our second hypothesis envisages that in order to learn to control one particular degree of freedom, it has to be included in the training movement. We showed this to be true. In addition, our results suggest that the control system does not have to train on the whole range of motion for each degree of freedom in order to generalize to all movements. This is important when we want to train the inverse model with minimal amount of effort.

Hypothesis 3 suggests that asynchronous movements are harder than synchronous movements, and that the control system will not be able to produce different motor commands for the two arms if it has not been trained to do so. For humans it is indeed true that it is easier to move the limbs synchronously. It is still very interesting that we get the same results for this control system, and interesting to see that a system trained to produce a synchronous movement and asked to generate an asynchronous movement provides the best solution it is able to, namely the average between the desired movement of the left and right arm.

Our findings suggests that imitation may be used as an efficient method to learn the inverse model, because one can choose the training sequence optimally, as opposed to exploration without guidance. This conclusion is supported by Rolf et. al. who suggests the use of goal directed exploration in contrast to motor babbling [14].

Further work should include more complex movements with larger degrees of freedoms where one target position can be reached through different motor commands. In addition more systematic evaluation of efficient training movements should be conducted.

References

1. Belardinelli, A., Pirri, F.: Bottom-up gaze shifts and fixations learning by imitation. *IEEE Trans. Syst. Man Cybern. B Cybern.* 37(2), 256–271 (2007)
2. Wood, M.A., Bryson, J.J.: Skill acquisition through program-level imitation in a real-time domain. *IEEE Trans. Syst. Man Cybern. B Cybern.* 37(2), 1083–1119 (2007)
3. Meltzoff, A.N., Moore, M.K.: Explaining facial imitation: A theoretical model. *Early Development and Parenting* 6, 179–192 (1997)
4. Demiris, Y., Dearden, A.: From motor babbling to hierarchical learning by imitation: a robot developmental pathway. In: *Proceedings of the Fifth International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, pp. 31–37 (2005)
5. von Hofsen, C.: An action perspective on motor development. *TRENDS in Cognitive Sciences* 8(6), 266–272 (2004)
6. Wolpert, D.M., Miall, R.C., Kawato, M.: Internal models in the cerebellum. *Trends in Cognitive Sciences* 2(9) (1998)
7. Shadmehr, R., Krakauer, J.W.: A computational neuroanatomy for motor control. *Experimental Brain Research* 185(3), 359–381 (2008)
8. Davidson, P.R., Wolpert, D.M.: Widespread access to predictive models in the motor system a short review. *Journal of Neural Engineering* 2(3), 313–319 (2005)
9. Kawato, M., Gomi, H.: The cerebellum and vor/okr learning models. *TINS* 15(11) (1992)
10. Mehta, B., Schaal, S.: Forward models in visuomotor control. *Journal of Neurophysiology* 88(2), 942–953 (2002)
11. van der Smagt, P., Hirzinger, G.: The cerebellum as computed torque model. In: Howlett, R., Jain, L. (eds.) *Fourth International Conference on Knowledge-Based Intelligent Engineering Systems & Applied Technologies* (2000)
12. Jaeger, H.: A tutorial on training recurrent neural networks, covering bppt, rtrl, and the echo state network approach. Techreport, Fraunhofer Institute for AIS (April 2008)
13. Tidemann, A., Öztürk, P.: Self-organizing multiple models for imitation: Teaching a robot to dance the YMCA. In: Okuno, H.G., Ali, M. (eds.) *IEA/AIE 2007. LNCS (LNAI)*, vol. 4570, pp. 291–302. Springer, Heidelberg (2007)
14. Rolf, M., Steil, J.J., Gienger, M.: Efficient exploration and learning of whole body kinematics. In: *IEEE 8th International Conference on Development and Learning* (2009)

Analysis and Comparison of Probability Transformations for Fusing Sensors with Uncertain Detection Performance

Edwin El-Mahassni, Samuel Davey, and Jonathan Legg

ISR Division

Defence Science and Technology Organisation

Edinburgh, SA, 5111, Australia

{Edwin.El-Mahassni, Samuel.Davey, Jonathan.Legg}@dsto.defence.gov.au

Abstract. In a recent paper by Davey, Legg and El-Mahassni a way of fusing sensors with uncertain performance was outlined using the Transferable Belief Model (TBM) theory. It showed that if the target prior was uncertain, then the resulting fused mass was also uncertain. That is, some belief mass was assigned to the case that the presence or absence of a target was unknown. Various methods have been proposed to transform an uncertain belief function into a probability mass. This paper analyses the relationship between an important subset of these methods and compares the resulting probability masses with those obtained via Bayesian methods using random priors.

Keywords: belief functions, probability transformations, sensor fusion.

1 Introduction

Sensor fusion is the combination of sensor data obtained from different sources to obtain more complete or accurate information than that available from any individual source. Two or more sensors with complementary capabilities may be used together to achieve a more reliable target state estimate. The scenario addressed by this paper is that of fusing two similarly performing sensors to enhance an operator's confidence in his/her decision making. The most popular methods for combining information from multiple sensors are Bayesian in nature, including networks and Kalman filters. Bayesian networks can make predictions based on a small number of observations [1]. They are models that represent variables and their probabilistic relationships.

A second method which has also attracted attention in the area of sensor fusion is that of Dempster-Shafer Theory (DST) [2,3]. DST may be thought of as a generalisation of probability theory; rather than assigning probability mass to an exhaustive collection of mutually-exclusive events, DST allows belief mass to be assigned to sets of events when the evidence or data is not able to distinguish between these events. Thus the DST is better suited for representing ambiguity and in particular can represent some measure of total ignorance by

assigning a non-zero mass to the universal set. An advantage of using DST is that it allows for random variables with uncertain probability mass, for example an event that occurs with probability at least p , or a prior between p_1 and p_2 .

In [4], the Transferable Belief Model (TBM) [6,7], a variant of DST, was used to fuse two sensors at the decision level (is there a target or not). It was found that uncertainty in the prior probability of a target being present lead to uncertainty in the fused belief mass. In [4] this was resolved by using the pignistic transform. However, alternative methods for transforming uncertain belief functions into probabilities have been proposed. Here, we extend the work of [4] and analyse several probability transformations that satisfy the minimal criterion of [5].

This paper also presents a Bayesian solution to the fusion problem when the sensor parameters are uncertain. This is achieved through treating the parameters as random variables with known priors. The Bayesian solution is compared with the TBM solution. The paper is divided in the following way: in Section 2 we give a short introduction to evidential reasoning and its application to sensor fusion when uncertainty exists. In Section 3 we describe and analyse some of the methods that transform belief masses into probabilities. In Section 4, we derive the Bayesian solution when target, detection and false alarm uncertainty exists. In Section 5, we provide a graphical example comparing the TBM probabilities and the Bayesian probabilities. Finally in Section 6, we provide some concluding remarks.

2 Applying Evidential Reasoning to Sensor Fusion

In DST, the *Basic Belief Assignment* (BBA) function (which is sometimes referred to as belief mass), m , defines a mapping of the power set to the unit interval. Note that $m(\emptyset) = 0$ is usually assumed for DST. For a given set $A \in \wp(X)$, where $\wp(X)$ is the powerset of a set X , the value $m(A)$ represents the proportion of all relevant and available evidence that supports the claim that a particular element of X belongs to the set A , but to no particular subset of A if A is not a singleton. We remark that conceptually, m can represent, but is not limited to, probability. The sum of all the masses is always 1.

Under the TBM variant of DST, mass is allowed to be assigned to the empty set, $m(\emptyset)$. One interpretation of the mass assigned to the empty set is that it is the degree of conflict between combined BBAs. The TBM combination rule does not contain the normalising term customary for DST and is given by

$$m_{1,2}(C) = \sum_{A \cap B = C} m_1(A) m_2(B). \quad (1)$$

The TBM was used in [4] to fuse detection level data from two sensors and derive the belief that a target is present or absent given the sensor outputs and uncertain knowledge of the sensor performance.

In [4], it was assumed that, for sensor i , the probability of detection, false alarm and prior probability of target being present are respectively given and bounded by

$$d_{\text{MIN}}^i \leq P_D^i \leq d_{\text{MAX}}^i, \quad f_{\text{MIN}}^i \leq P_{\text{FA}}^i \leq f_{\text{MAX}}^i \quad \text{and} \quad t_{\text{MIN}} \leq P_t \leq t_{\text{MAX}}.$$

Letting $D = d_{\text{MAX}}^1 d_{\text{MAX}}^2$ and $F = f_{\text{MAX}}^1 f_{\text{MAX}}^2$, [4] showed that the normalised belief that a target is present given detections from both sensors is given by

$$\begin{aligned} m(1) &= D t_{\text{MAX}} - D F (t_{\text{MAX}} - t_{\text{MIN}})/C, \\ m(0) &= F (1 - t_{\text{MIN}}) - D F (t_{\text{MAX}} - t_{\text{MIN}})/C, \\ m(x) &= D F (t_{\text{MAX}} - t_{\text{MIN}})/C, \end{aligned} \tag{2}$$

where $C = t_{\text{MIN}}(1 - D) + (1 - F) + (t_{\text{MAX}} - t_{\text{MIN}})(1 - D)(1 - F)$ is the mass assigned to the empty set and $m(x)$ reflects the uncertainty and is the mass assigned to the case that we do not know if there is a target.

When a decision is made using TBM, BBAs need to be transformed into probabilities. In order to do this the uncertain mass, $m(x)$, should be distributed between the target present and target absent cases. Note that if the target prior is known, $t_{\text{MIN}} = P_t = t_{\text{MAX}}$, and so $m(x) = 0$. A certain probability of target presence is obtained by simply normalising $m(0)$ and $m(1)$ even though the sensor performance is still uncertain. This solution is equivalent to the Bayesian solution for known sensor priors where these have been replaced with their upper bound values.

3 Probabilities from Belief Function Models

Different methods have been proposed to transform belief functions into probability masses, many of which were explored in [5]. Here we consider those which satisfy the *minimal criterion* of [5], which states that the result of applying the transform to the vacuous belief function should be a uniform probability. The vacuous belief function is one describing complete uncertainty. Three such transforms were identified in [5]: the Aggregate Uncertainty (AU); the pignistic transformation; and the Plausibility transformation. Each is described in detail in [5]. In addition, we will consider the generalized pignistic transformation [8,9] since it also satisfies the minimal criterion.

We will consider only the case of a binary frame of discernment, $X = \{A, B\}$, since this is appropriate for the decision fusion example. Assume that after using the TBM for sensor fusion, the masses are normalized to force the empty set mass to be zero. Thus there are three masses, $m(A)$, $m(B)$ and $m(A \cup B)$. Without loss of generality, we let $m(A) \geq m(B)$. The general requirement of the transform is to map these three masses into binary probabilities, $p(A)$ and $p(B)$.

3.1 Aggregate Uncertainty Approach

The Aggregate Uncertainty method distributes the mass $m(A \cup B)$ between $m(A)$ and $m(B)$ and then normalises the two to arrive at a probability mass.

The AU selects the proportion of $m(A \cup B)$ distributed to each in such a way as to maximise the entropy of the resulting probability mass [3]. For a binary frame, this process is straightforward: if $m(A) \geq 0.5$, the transform yields $P_1(A) = m(A)$ and $P_1(B) = m(B) + m(A \cup B)$, otherwise $P_1(A) = P_1(B) = 0.5$.

3.2 The Plausibility Transformation

Also known as the Bayesian approximation method, it has been argued that this transformation is the most consistent technique for translating Dempster's rule of combination to the Bayesian probability domain. For a binary frame, if $x \in X$, the probabilities given by this transformation become

$$P_2(x) = \frac{m(x) + m(A \cup B)}{1 + m(A \cup B)}.$$

3.3 The Pignistic Transformation

The pignistic transform moves the belief mass from the union elements of the power set and distributes it equally amongst the singleton members. For a binary frame, if $x \in X$, the transform gives

$$P_3(x) = m(x) + \frac{1}{2}m(A \cup B). \quad (3)$$

For a binary frame, it can be shown that the Baroni-Vicig transformation [10] is equivalent to the pignistic transformation.

3.4 The Generalised Pignistic Transformation

This transformation aims to maximise the Probabilistic Information Content (PIC). Here it suffices to say that maximization of the PIC is equivalent to minimization of Shannon's entropy [8,9,11]. For a binary frame, if $x \in X$, the transformation is given by

$$P_4(x) = \frac{m(x)(1 + 2\varepsilon) + \varepsilon m(A \cup B)}{m(A) + m(B) + 2\varepsilon},$$

for some small arbitrary $\varepsilon \geq 0$.

The entropy of $P_4(x)$ is minimised if $\varepsilon = 0$ [8,9]. However, if $m(B)$ is zero and $m(A \cup B)$ is not, then the transform gives $p(A) = 1$ despite the uncertainty, which is inappropriate.

3.5 Analysis of Probability Transformations

We now present a relationship between these probability transformations for a binary frame and $m(A) \geq m(B)$.

Theorem 1

If $X = \{A, B\}$ and $m(A) \geq m(B)$ then $P_1(A) \leq P_2(A) \leq P_3(A) \leq P_4(A)$.

Proof. For the case when $m(A) \geq 0.5$, then showing $P_1(A) \leq P_2(A)$ is equivalent to proving that

$$m(A) \leq \frac{1 - m(B)}{1 + m(A \cup B)}$$

$m(A) + m(A)m(A \cup B) + m(B) \leq 1$, so that $P_1(A) \leq P_2(A)$ when $m(A) \geq 0.5$. If $m(A) < 0.5$, $m(A) \geq m(B)$ then $P_1(A) = P_1(B) = 0.5$. In such a case we note that $0.5(1 + m(A \cup B)) \leq m(A) + m(A \cup B)$ so that $P_1(A) \leq P_2(A)$.

Similarly, after some manipulation we can also see that when $m(A) \geq m(B)$

$$m(A) + m(A \cup B) \leq \left[m(A) + \frac{1}{2}m(A \cup B) \right] [1 + m(A \cup B)].$$

This means that $P_2(A) \leq P_3(A)$. And lastly, we note that $P_3(A) \leq P_4(A)$ because it can easily be shown that

$$\left[m(A) + \frac{1}{2}m(A \cup B) \right] [m(A) + m(B) + 2\varepsilon] \leq m(A)[1 + 2\varepsilon] + \varepsilon m(A \cup B).$$

Finally, from above the following lemma can be established. It provides a lower bound on the probability values for all the transformations.

Lemma 1. *For all the probability transformations listed above we have $p(A) \geq m(A)$ and $p(B) \geq m(B)$.*

As mentioned in section 2, when the prior probability of a target is known, then there is no mass assigned to the uncertain event $A \cup B$ and so all of the transforms are equivalent.

4 Bayesian Solution

Section 3 establishes a relationship between a number of different transforms, but it does not provide any advice regarding which, if any, is preferred. One potential way to discriminate between them is to compare the transforms with the result of Bayesian analysis. Is it possible, for example, to craft a Bayesian prior to give the same result as the pignistic transform?

When the parameters of a prior distribution are unknown, the Bayesian approach is to treat these parameters as random variables. Such random parameters are referred to as hyperparameters and their distributions as hyperpriors [12]. Within the sensor fusion context, we will now treat the probability of detection, the probability of false alarm and the probability of target presence as hyperparameters with known hyperpriors. That is, P_D^i is a random variable with known distribution (hyperprior) $p(P_D^i)$.

Let t , s^1 and s^2 be binary indicator variables that denote the presence (or absence) of a target and sensor reports respectively.

The target is present with probability P_t . This probability is unknown, but it has a known hyperprior, $p(P_t)$. The mean target prior is $E\{P_t\}$ and the hyperprior is zero outside of the region $[t_{\text{MIN}}, t_{\text{MAX}}]$. Explicitly,

$$p(t|P_t) = \begin{cases} 1 - P_t & t = 0, \\ P_t & t = 1. \end{cases} \quad (4)$$

The sensor output is a function of the sensors' probability of detection, probability of false alarm, and the presence of a target, i.e.

$$p(s^i|t, P_D^i, P_{FA}^i) = \begin{cases} 1 - P_{FA}^i & t s^i = 00, \\ P_{FA}^i & t s^i = 01, \\ 1 - P_D^i & t s^i = 10, \\ P_D^i & t s^i = 11, \end{cases} \quad (5)$$

where again the probability of detection, P_D^i , and probability of false alarm, P_{FA}^i , are unknown but have known hyperpriors, $p(P_D^i)$ and $p(P_{FA}^i)$.

Theorem 2. *Let $p(t|s^1, s^2)$ be the conditional probability of a target being present ($t = 1$) or not ($t = 0$), given two sensor outputs. Then, the Bayesian solution when target, false alarm and detection priors are unknown is given by*

$$p(t|s^1, s^2) \propto p(t|E\{P_t\}) \prod_{i=1}^2 p(s^i|t, E\{P_D^i\}, E\{P_{FA}^i\}) \quad (6)$$

Proof. The fused output, $p(t|s^1, s^2)$, is derived through the use of Bayes' Rule, the law of total probability and conditional independence:

$$\begin{aligned} p(t|s^1, s^2) &\propto p(t, s^1, s^2) \\ &= \int_0^1 \int_0^1 \int_0^1 \int_0^1 \int_0^1 p(t, s^1, s^2, P_t, P_D^1, P_D^2, P_{FA}^1, P_{FA}^2) \\ &\hspace{20em} dP_t dP_D^1 dP_D^2 dP_{FA}^1 dP_{FA}^2, \\ &= \int_0^1 \int_0^1 \int_0^1 \int_0^1 \int_0^1 p(P_t) p(t|P_t) \\ &\hspace{10em} \times \prod_{j=1}^2 p(P_D^j) p(P_{FA}^j) p(s^j|t, P_D^j, P_{FA}^j) dP_t dP_D^1 dP_D^2 dP_{FA}^1 dP_{FA}^2, \\ &= \left\{ \int_0^1 p(P_t) p(t|P_t) dP_t \right\} \\ &\hspace{10em} \times \prod_{j=1}^2 \left\{ \int_0^1 \int_0^1 p(P_D^j) p(P_{FA}^j) p(s^j|t, P_D^j, P_{FA}^j) dP_D^j dP_{FA}^j \right\} \quad (7) \end{aligned}$$

The first term in (7) can be simplified by substituting (4)

$$\begin{aligned} \int_0^1 p(P_t) p(t|P_t) dP_t &= \begin{cases} \int_0^1 p(P_t) (1 - P_t) dP_t & t = 0, \\ \int_0^1 p(P_t) P_t dP_t & t = 1, \end{cases} \\ &= \begin{cases} 1 - E\{P_t\} & t = 0, \\ E\{P_t\} & t = 1. \end{cases} \\ &= p(t|E\{P_t\}). \end{aligned} \quad (8)$$

Similarly, the second term in (7) is simplified by substituting (5) to get

$$\begin{aligned} \int_0^1 \int_0^1 p(P_D^i) p(P_{FA}^i) p(s^i|t, P_D^i, P_{FA}^i) dP_D^i dP_{FA}^i &= \begin{cases} 1 - E\{P_{FA}^i\} & t s^i = 00, \\ E\{P_{FA}^i\} & t s^i = 01, \\ 1 - E\{P_D^i\} & t s^i = 10, \\ E\{P_D^i\} & t s^i = 11, \end{cases} \\ &= p(s^i|t, E\{P_D^i\}, E\{P_{FA}^i\}). \end{aligned} \quad (9)$$

It is very interesting, and perhaps unexpected, that the fused probability of a target given multi-sensor data depends only on the means of the hyperpriors. The implication is that we could solve this problem using the Bayesian method for an unspecified hyperprior constrained only by its mean.

It is intuitive that (6) extends directly to an arbitrary number of sensors simply by appropriately changing the domain of the product. For the special case of N identical sensors where $M \leq N$ sensors report a target, the result resembles a binomial distribution.

For the example of two sensors reporting a detection, the probability of a target being present is given by

$$p(t = 1|1, 1) = \frac{E\{P_t\}E\{P_D^1\}E\{P_D^2\}}{E\{P_t\}E\{P_D^1\}E\{P_D^2\} + (1 - E\{P_t\})E\{P_{FA}^1\}E\{P_{FA}^2\}}. \quad (10)$$

Let the Bayesian probability distribution be denoted by P_5 . An obvious lemma when all the marginals are known can now be stated.

Lemma 2. *If $t_{\text{MIN}} = t_{\text{MAX}}$, $d_{\text{MIN}}^i = d_{\text{MAX}}^i$, $i = 1, 2$, then*

$$P_1(j|(1, 1)) = P_2(j|(1, 1)) = P_3(j|(1, 1)) = P_4(j|(1, 1)) = P_5(j|(1, 1)),$$

that is, the TBM solution is the same as the Bayesian solution, independent of the belief to probability transformation method used.

Recall that when the target probability is known then all the probability transformations are equivalent. For this case, denote the TBM derived probability value of an event A by $P^*(A)$ and establish the following resulting lemma which indicates when the Bayesian solution will yield a higher or lower value than the probability transformations when the target prior probability is known.

Lemma 3. *When $t_{\text{MIN}} = t_{\text{MAX}}$, then*

$$\frac{D}{F} \geq \frac{E\{P_D^1\}E\{P_D^2\}}{E\{P_{FA}^1\}E\{P_{FA}^2\}} \iff P^*(t|(1,1)) \geq P_5(t|(1,1))$$

The dependence between the parameters and the probability values is more complicated when the target prior is uncertain and depends on the transformation method. For example, under the pignistic transform, Lemma 3 becomes:

Lemma 4.

$$\frac{t_{\text{MAX}}D - \frac{1}{2}(t_{\text{MAX}} - t_{\text{MIN}})DF}{(1 - t_{\text{MIN}})F - \frac{1}{2}(t_{\text{MAX}} - t_{\text{MIN}})DF} \geq \frac{E\{P_t\}E\{P_D^1\}E\{P_D^2\}}{(1 - E\{P_t\})E\{P_{FA}^1\}E\{P_{FA}^2\}}$$

$$\iff P_3(t|(1,1)) \geq P_5(t|(1,1))$$

While there is some relationship between the transforms and the Bayesian solution, it is too complicated to draw any intuitive conclusion.

5 Examples

In this section we provide some graphical examples which help to visualize the relationships between the different transformations and the Bayesian solution. In all cases, we assume that the false alarm and detection parameters of both sensors are the same.

First consider the case of a known target prior, $P_t = 0.5$. Assume that the hyperpriors for the probability of detection and the probability of false alarm are both uniform, so

$$p(P_D^i) = \begin{cases} \frac{1}{d_{\text{MAX}} - d_{\text{MIN}}} & d_{\text{MIN}} \leq P_D^i \leq d_{\text{MAX}}, \\ 0 & \text{otherwise,} \end{cases}$$

with $E\{P_D^i\} = \frac{1}{2}(d_{\text{MAX}} + d_{\text{MIN}})$, and similarly for $p(P_{FA}^i)$.

Figure 1 shows the difference between the TBM probability of a target and the Bayesian probability of a target as d_{MAX} and f_{MAX} are varied. Both d_{MIN} and f_{MIN} were also varied to maintain fixed values of $E\{P_D\} = 0.8$ and $E\{P_{FA}\} = 0.25$. A white line is plotted showing the values where the two probabilities are equal, as defined by Lemma 3. For this example, when the uncertainty in the probability of detection is greater than the uncertainty in the probability of false alarm, the TBM fused probability of a target is higher than the Bayesian fused probability and when the probability of false alarm uncertainty is greater, the TBM fused probability is lower.

Next we change the target prior to be random and uniformly distributed on $[0.3, 0.7]$. As before, the detection probability hyperprior and false alarm probability hyperprior were varied but constant means were maintained. Figure 2 shows the difference between the pignistic transform probabilities and the Bayesian probabilities. The general shape of the function is the same as for a fixed P_t , but it has been shifted to the right. This means that for a particular false alarm hyperprior, equivalence is achieved for a smaller d_{MAX} when the target prior is uncertain.




Fig. 1. TBM probability - Bayes Probability, P_t known




Fig. 2. TBM probability - Bayes Probability, P_t random

6 Concluding Remarks

In this paper, we investigated several Dempster-Shafer mass probability transformations and analysed how they related to one another in the context of decision fusion in a dual sensor example. We showed that if the target prior is fixed, then an intuitive relationship can be established between the probability transformations and the Bayesian solution. However, even in this simplified case, the two are only equivalent in special cases. There is no apparent relationship between the TBM and the Bayesian solutions or particular form of prior that makes the two equivalent. This is because the TBM belief depends on the boundary conditions of the parameter hyperpriors, whereas the Bayesian solution presents an overall average result. Potential work could concentrate on determining in which situations it would be preferable to use a particular transformation over another.

References

1. Niedermeyer, D.: An Introduction to Bayesian Networks and Their Contemporary Applications. In: Holmes, D., Jain, L. (eds.) *Innovations in Bayesian Networks*. SCI, vol. 156, pp. 117–130 (2008)
2. Dempster, A.P.: A Generalisation of Bayesian Inference. *J. of the Royal Statistical Society, Series B* 30, 205–247 (1968)
3. Shafer, G.: *A Mathematical Theory of Evidence*. Princeton University Press, Princeton (1976)
4. Davey, S., Legg, J., El-Mahassni: Fusing Sensors with Uncertain Detection Performance. In: *Proc. 5th. International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP)*, pp. 267–272 (2009)
5. Oxenham, M., Challa, S., Morelande, M.: Fusion of Disparate Identity Estimates for Shared Situation Awareness in a Network-Centric Environment. *Information Fusion* 7, 395–417 (2006)
6. Smets, P.: The Combination of Evidence in the Transferable Belief Model. *IEEE Trans. PAMI* 12, 447–458 (1990)
7. Smets, P.: The nature of the unnormalized beliefs encountered in the transferable belief model. In: Dubois, D., Wellmann, M.P., D’Ambrosio, B., Smets, P. (eds.) *Proc. 8th Conf. on Uncertainty in AI*, pp. 292–297. Morgan Kaufmann Publ., San Mateo (1992)
8. Dezert, J., Smarandache, F.: A New Probabilistic Transformation of Belief Mass Assignment. In: *11th. International Conference in Information Fusion* (2008)
9. Dezert, J., Smarandache, F.: Advances and Applications of DS_mT for Information Fusion. In: Smarandache, F., Dezert, J. (eds.) *Collected Works, Perfect Paperback*, vol. 3 (2009)
10. Baroni, P., Vicig, P.: Transformations from Imprecise to Precise Probabilities. In: Nielsen, T.D., Zhang, N.L. (eds.) *ECSQARU 2003. LNCS (LNAI)*, vol. 2711, pp. 37–49. Springer, Heidelberg (2003)
11. Sudano, J.: The System Probability Information Content (PIC) Relationship to Contributing Components, Combining Independent Multi-Source Beliefs, Hybrid and Pedigree Pignistic Probabilities. In: *Proc. 5th Conf. in Information Fusion*, vol. 2, pp. 1277–1283 (2002)
12. Bernardo, J.M., Smith, A.F.M.: *Bayesian Theory*. Wiley, Chichester (2000)

PART V
Case-Based Reasoning/Planning
and Scheduling

A Case-Based Approach to Business Process Monitoring

Stefania Montani and Giorgio Leonardi

Dipartimento di Informatica, Università del Piemonte Orientale, Alessandria, Italy

Abstract. The *agile workflow* technology deals with flexible workflow adaptation and overriding, in case of foreseen as well as unforeseen changes and problems in the operating business environment. One key issue that an agile workflow system should address is *Business Process (BP) monitoring*. This consists in properly highlighting and organizing non-compliances and adaptations with respect to the default process schema. Such an activity can be the starting point for other very critical tasks, such as quality assessment and process reengineering.

In this paper, we introduce an automated support to BP monitoring, which exploits the Case-based Reasoning (CBR) methodology. CBR is particularly well suited for managing exceptional situations, and has been proposed in the literature for process change reuse and workflow adaptation support. Our work extends these functionalities by retrieving traces of process execution similar to the current one, which can then be automatically clustered. Retrieval and clustering results can provide support both to end users, in the process instance execution phase, and to process engineers, in (formal) process quality evaluation and long term process schema redefinition. Our approach in practice is illustrated by means of a case study in the field of stroke management.

1 Introduction

Business Process (BP) Management is a set of activities aimed at defining, executing, monitoring and optimizing BP, with the objective of making the business of an enterprise as effective and efficient as possible, and of increasing its economic success. Such activities are highly automated, typically by means of the workflow technology [1,22].

BP optimization, in particular, may ask the enterprise to be able to flexibly change and adapt the predefined process schema, in response to expected situations (e.g. new laws, reengineering efforts) as well as to unanticipated exceptions and problems in the operating environment (e.g. emergencies) [6].

Agile workflow technology [26] is the technical solution which has been invoked to deal with such adaptation and overriding needs. It can support both ad-hoc changes of individual process instances [18,25], operated by end users, and modifications at the general process schema level, operated by process engineers - applicable even if the default schema is already in use by some running instances [18,5].

In order to provide an effective and quick workflow change support, many agile workflow systems share the idea of recalling and reusing concrete *examples of changes* adopted in the past. To this end, Case-based Reasoning (CBR) [2] has been proposed as a natural methodological solution. CBR is a reasoning paradigm that exploits the specific knowledge of previously experienced situations, called *cases*. It operates by *retrieving* and *reusing* - by possibly *revising* - them in order to solve the problem at hand. CBR is particularly well suited for managing exceptional situations, even when they cannot be foreseen or pre-planned. As a matter of fact, in the literature cases have often been resorted to in order to describe exceptions, in various domains (see e.g. [20]), and many examples of CBR-based process change reuse and workflow adaptation support have been proposed (see e.g. [12,10,25,14,15]).

A less explored (see however e.g. [8]), but very critical issue to be addressed in agile workflow systems is the one of BP *monitoring*. This activity consists in highlighting and organizing non-compliances with respect to the default process schema. It can be the starting point for a formal verification of process conformance to proper semantic constraints (see e.g. [11]), or for suggesting long term changes, in front of frequent, similar non-compliances. Providing BP monitoring functionality is non trivial. As a matter of fact, deviations from a default process schema generate process instances, typically stored as traces of actions, that are different from how they were supposed to be. Monitoring BP from traces is particularly hard when no contextual information, which could justify the reasons for deviation, is recorded in the traces themselves. A facility able to intelligently exploit traces of process executions, by retrieving similar ones, and by automatically organizing them, would then be an added value for an agile workflow tool. It could also provide support both to end users, in the process instance execution phase, and to process engineers, in (formal) quality evaluation and/or in long term process schema redefinition.

In this paper, we propose a CBR-based approach to BP monitoring, based on execution traces retrieval. In particular, in this work we focus on the definition of a proper similarity measure, and on its use for calculating trace distance. Our next step will be the implementation of a trace clustering technique, in order to support automatic non-compliances organization. Technical details of the approach are presented in section 2, and a case study on stroke management execution traces is described in section 3. Finally, section 4 addresses some comparisons, discussion and concluding remarks.

2 Case-Based BP Monitoring

Our framework is meant to support end users in process instance modification by retrieving traces of execution similar to the current one: suggestions on how to modify the default process schema in the current situation may be obtained by analyzing the most similar retrieved examples of change, recorded as traces that share the starting sequence of actions with the current query.

Interestingly, we could also automatically cluster the execution traces stored in the database on the basis of their similarity, and allow process engineers to

inspect the obtained clusters, in order to visualize the most frequent changes. Since changes can be an indicator of non-compliance, clustering can be seen as the first step in a process quality evaluation activity, which can be realized by means of formal (e.g. logic-based) verification approaches. Additionally, since changes can also be due to a weak or incomplete initial process schema definition, engineers could exploit clustering results to draw some suggestions on how to redefine process schemas, in order to incorporate the more frequent and significant changes once and for all.

From a technical viewpoint, both the decision support functionalities illustrated above require a preliminary design phase, in which:

1. the case structure is defined;
2. a proper similarity measure, to be used for retrieval and clustering, is identified.

Such issues are the main topic of the paper, and will be illustrated in this section.

Issue 1 is rather straightforward. We have defined a case as a trace of execution of a given process schema. In particular, every trace is a sequence of actions, each one stored with its execution time and duration information.

Issue 2 is more interesting. In the literature, a number of similarity measure definitions for agile workflows exist, that require further information in addition to the workflow structure, such as semantic annotations [21], or conversational knowledge [26,25]. The approaches are typically context-aware, that is, the contextual information is considered as a part of the similarity assessment of workflows.

Unfortunately, any contextual information, as well as conversational knowledge, is not always available, especially when instances of process execution are recorded as traces of actions. Starting from this observation, a rather simple graph edit distance measure [4] has been proposed and adapted for similarity assessment in workflow change reuse [14,8]. Our approach moves from the same graph edit distance definition, and properly adapts it in order to work on trace comparison.

Basically, in our approach similarity is modeled through a set of edit operations on traces. Each edit operation performs a modification of the following kinds: (i) *substitute* one action with a different one, (ii) *insert* a new action, or (iii) *delete* an action.

The cost of a *substitution* is not always set to 1, as in the classical edit distance. In fact, as in the weighted edit distance (see e.g. [9]), we define it as a value $\in [0, 1]$ which depends on what action appears in a trace as a substitution of the corresponding action in the other trace. In particular, we organize actions in a *taxonomy*, on the basis of domain knowledge. The more two actions are close in the taxonomy, the less penalty has to be introduced for substitution ([17]; see also [3,19,16]). In detail, in our work substitution penalty is set to the *taxonomical distance* dt between the two actions ([17], see Definition 2), i.e. to the normalized number of arcs on the path between the two actions in the taxonomy.

Insertions do not always cost 1 as well. In fact, an insertion may introduce a degree of indirection in a path otherwise connecting the same pair of actions in the two traces (see figure 1, third case). Our distance definition allows to capture this situation - which is very relevant for BP monitoring. The definition introduces a knowledge-based parametrized weight $\in [0, 1]$ for such insertions, depending on the action type. The final penalty of an insertion which generates an indirection is therefore equal to 1 multiplied by the weight. Naturally, the more insertions are performed, the more indirection is obtained in the path, and the more penalties are added. *Deletions* simply work dually with respect to insertions.




Fig. 1. Comparing the i -th action B in trace P with the j -th action B' in trace Q

The global distance $d(P, Q)$ between two traces P and Q is finally defined as the total cost of a sequence of edit operations which transform one trace into the other.

Formally, we provide the following definitions:

Definition 1: Trace edit distance. Let P and Q be two traces of actions, and let α and β be two actions. The *trace edit distance* between P and Q is defined as:

$$d(P, Q) = \sum_{i=1}^k c(e_i)$$

where (e_1, \dots, e_k) transforms P into Q , and:

- $c(e_i) = dt(\alpha, \beta)$, if e_i is the substitution of α (appearing in P) with β (appearing in Q), with $dt(\alpha, \beta)$ defined as in Definition 2 below;
- $c(e_i) = 1 * w_\alpha$, if e_i is the insertion (the deletion) of action α in P (from Q), with w_α defined as in Definition 3 below.

Definition 2: Taxonomical distance [17]. Let α and β be two actions in the taxonomy t , and let γ be the closest common ancestor of α and β . The *taxonomical distance* dt between α and β is defined as:

$$dt(\alpha, \beta) = \frac{N_1 + N_2}{N_1 + N_2 + 2 * N_3}$$

where N_1 is the number of arcs in the path from α and γ in t , N_2 is the number of arcs in the path from β and γ , and N_3 is the number of arcs in the path from the taxonomy root and γ [9].

Terminology. Two actions α and β are **comparable** if $dt(\alpha, \beta) \leq \tau$, where τ is a threshold to be set on the basis of domain knowledge.

Definition 3: Action weight. Let P and Q be two traces of actions and let α be an action (appearing in Q). The *action weight* w_α of α is defined as:

- $w_\alpha \in [0, 1[$ - to be set on the basis of domain knowledge - if α generates an indirection (see figure 1, third case) in Q with respect to P ;
- $w_\alpha = 1$ otherwise (e.g. if P is a substring of Q , and α is an action in the head or tail portion of Q , not matched to any action in P).

Operatively, at a generic iteration of our algorithm we compare the i -th action (say B) in trace P , with the j -th action B' in trace Q . If they are *comparable* (see figure 1, first case), the distance is not increased (or minimally increased, if the actions are not identical), and the next actions are considered. Otherwise, in order to discover insertions that generate indirections, we also compare the $(i - 1)$ -th action A in P with the $(j - 1)$ -th action A' in Q . If their are not comparable (see figure 1, second case), no indirection is being generated, and $dt(B, B')$ is added to the distance calculation. On the other hand, if A and A' are comparable, we need to compare B with the $(j + 1)$ -th action in Q : if they are comparable as well (see figure 1, third case), B' has been inserted to create an indirection on trace Q , and the cost $1 * w_{B'}$ must be added to the distance calculation; otherwise (see figure 1, fourth case), a cost equal to $dt(B, B')$ will be added.

In our approach similar traces retrieval is then performed by classical K-Nearest Neighbor techniques, where the most proper value of k has to be experimentally set according to the specific application domain needs.

As a next step, we will concentrate on the implementation of the clustering facility, which will be based on hierarchical clustering techniques.

3 The Framework in Practice: An Application to Stroke Management

Health-Care Organizations (HCO) place strong emphasis on efficiency and effectiveness, to control their health-care performance and expenditures. Therefore,

¹ Note that, if $dt(\alpha, \beta) > th$, being $th \in [0, 1]$ a proper, domain-dependent threshold, $dt(\alpha, \beta)$ can be forced to 1.

it is important to evaluate existing infrastructures and the services provided. To perform this operation, it is crucial to explore and process the data collected by the HCO systems, organizing them in form of process logs (i.e. traces of execution), which can be seen as the history of what happened in the HCO. Traces can be helpful to gain a clear picture of the actual care process, through the use of BP monitoring techniques [13], like the ones introduced in the previous section. As an example, our framework is currently being tested in the stroke management domain.

A stroke is the rapidly developing loss of brain function(s) due to disturbance in the blood supply to the brain. This can be due to ischemia (lack of glucose and oxygen supply) caused by thrombosis or embolism, or to a hemorrhage. As a result, the affected area of the brain is unable to function, leading to inability to move one or more limbs on one side of the body, inability to understand or formulate speech, or inability to see one side of the visual field. A stroke is a medical emergency and can cause permanent neurological damage, complications, and death. It is the leading cause of adult disability in the United States and Europe. It is the number two cause of death worldwide and may soon become the leading one. The best medical practice [7] requires that stroke patients are treated according to a management protocol, which is basically composed by four steps: (1) emergency management; (2) hospitalization; (3) dismissal; (4) follow up. Each step is in turn composed by a sequence of actions, which must respect some criteria, although inter-patients and inter-hospitals variations are admissible. In particular, in step (1), symptoms onset must be recognized, the patient must be taken to the hospital, and a brain computer-assisted tomography (CAT) must be executed. In step (2), diagnosis has to be finalized, by means of a neurological evaluation and of several additional diagnostic investigations, meant to confirm the stroke hypothesis. Diagnostic procedures may vary, but most patients undergo electrocardiogram (ECG) and chest X-ray. At the same time, administrative patient admission procedures must be fulfilled. Finally, a proper therapy has to be initiated: for instance, up to 90% patients are treated with antiaggregants. Rehabilitation also must be started as soon as possible during hospitalization.

Our system is not fully implemented; therefore it has not entered its validation phase yet. Nevertheless, we can describe some first experiments, to show how the tool works. In our experiments, we used traces collected on real patients, detailing the actions of steps (1) and (2). Our case base is currently composed by more than 300 traces, collected at one of the major Stroke Units in Lombardia Region, Italy.

As an example, we will show the retrieval results related to the query case no. 103101, which presents a rather atypical situation. As a matter of fact, patient no. 103101's CAT results, in step (1), allowed to immediately diagnose a stroke episode, without the need of additional diagnostic procedures. This fact enabled the responsible physician to start the antiaggregant therapy already in step (1). During step (2), the severity of the patient conditions was investigated by means of additional tests, and a further anticoagulant therapy was started. However,

Query case: 103101 → Case retrieved: 109984

	Events from case 103101:	Events from case 109984:	
Time ↓	stroke_onset	stroke_onset	
	arrival_emergency_ward	arrival_emergency_ward	
	CAT	CAT	
	antiaggregant		→ indirection
	admission	admission	
	neurological_evaluation	neurological_evaluation	
	ECG	ECG	
	coagulative_screening	coagulative_screening	
	anticoagulant	antiaggregant	→ substitution → indirection
	ric_invest_XR thorax		
	NMR brain with DWI	NMR brain with DWI	
	angio NMR	angio NMR	
	trans-thoracic ECG		→ indirection
	echo doppler SAT	echo doppler SAT	
	transcranial doppler	transcranial doppler	
	discharge	discharge	
follow_up	follow_up		

Fig. 2. The best matching retrieved case in our experiment

rehabilitation was not started in phase (2) - which is a very anomalous situation - probably due to the patient's very critical health status.

Our tool retrieved the 20 Nearest Neighbor cases with respect to case no. 103101, adopting the similarity measure described in the previous section - which took about 1.2 second on an Intel Core 2 Duo T9400, equipped with 4 Gb of DDR2 ram. In particular, the most similar retrieved case, reported in figure 2, is a more standard one with respect to the query case as regards step (1) actions, since diagnosis did not take place during emergency management, but was clarified only during hospitalization, as it usually happens. Therefore, the antiaggregant therapy was not started in step (1). Therapy start in case no. 103101 was then recognized as an indirection in an otherwise identical sequence of actions (see figure 1, third case). Analogously, a couple of additional diagnostic procedures took place in case no. 103101, step (2), determining indirections with respect to the retrieved case. On the other hand, antiaggregant therapy was started in the retrieved case after all the diagnostic tests were completed. This action is a substitution of the anticoagulant therapy in case no. 103101, and is also *comparable* to it, having set $\tau = 0.2$: actually, the two drugs have a very similar effect, and consequently the two therapeutic actions are very close in the domain taxonomy. Interestingly, rehabilitation was also missing in the retrieved case: this means that our tool was able to correctly retrieve one of the few cases in the case base matching the query case with respect to this very atypical feature.

The physicians working with us judged our first experimental results as very reasonable. Encouraged by these outcomes, in the next future, we plan to extensively test the performances of our tool on additional real cases, after having worked at the implementation and validation of the clustering procedure.

4 Discussion and Conclusions

In this work, we have described a CBR-based approach to BP monitoring. In particular, we have defined a proper case structure and a new similarity measure, that are exploited to retrieve traces of execution similar to the current one. In the next future, the similarity measure will also be applied to cluster the traces available in the database. Such functionalities will help end users who need to adapt a process instance to some unforeseen situation, by retrieving changes applied in the past to other instances of the same process. Moreover, process engineers will take advantage of the retrieval and clustering results for identifying the most frequent changes to the same process schema. Such changes can be an index of non conformance of process executions with respect to proper constraints, but can also be a suggestion for properly revising an incorrect or obsolete process schema definition.

From the technical viewpoint, as observed in section 2, the graph edit distance measure, that we have adapted in this work, was originally resorted to in [14]. However, with respect to that approach, by focusing just on traces of execution we did not need to consider extensions to the similarity measure able to deal with control flow elements (such as alternatives and iterations). As a matter of fact, traces are always linear, i.e. they just admit the sequence control flow element. For the same reason, we did not insert any penalty for pairs of arcs in different traces with a different input or output action, which is already accounted for by the comparison between pairs of actions themselves. On the other hand, when focusing on linear traces our approach is more general and flexible than the one in [14]. As a matter of fact, we resort to taxonomical knowledge for comparing pairs of actions, so that two different actions do not always have a zero similarity. Moreover, we are able to recognize an indirect path from two actions, and to properly weight the degree of indirection in a parametrized way.

Explicitly comparing pairs of arcs in different traces will be needed if considering temporal distances between actions and actions durations, which we will explore as a future work. In particular, in the similarity measure definition, we plan to calculate a penalty for arcs connecting the same pair of actions in both the traces at hand, but introducing a different temporal delay between them. The penalty will be proportional to the difference between the two delays. This calculation will also be adapted to the case in which the same pair of actions is directly connected in one trace, and indirectly in the other trace. In this case, the difference in the arc durations will also need to properly take into account the degree of indirection. Some effort in the direction of considering time in trace comparison has been proposed in [8]. The similarity function in [8], however, does not exploit action duration, and does not rely on taxonomical information about actions, as we do. The authors also do not distinguish between direct and indirect paths connecting the same actions, so that our approach, once extended to deal with temporal information, will potentially be more flexible.

As a final consideration, in the future we will also explore the possibility of incorporating our work as a plug-in in the ProM tool [24], which is an open source framework for process mining. Process mining [23] describes a family of

a posteriori analysis techniques exploiting the information recorded in traces of actions. This process information can be used to *discover* the underlying process schema, when no a priori model is available. Once the process schema is obtained, the incorporation of our similarity measure and of the clustering facility could support an analysis of deviations from the process schema itself, which can be the input for a (formal) compliance verification. Moreover, in cooperation with other performances analysis plug-ins already embedded in ProM, our work could support a principled reengineering activity.

References

1. Workflow management coalition, <http://www.wfmc.org/wfmc-publications.html> (last accessed on October 6, 2009)
2. Aamodt, A., Plaza, E.: Case-based reasoning: foundational issues, methodological variations and systems approaches. *AI Communications* 7, 39–59 (1994)
3. Bergmann, R., Stahl, A.: Similarity measures for object-oriented case representations. In: Smyth, B., Cunningham, P. (eds.) *EWCBR 1998*. LNCS (LNAI), vol. 1488, p. 25. Springer, Heidelberg (1998)
4. Bunke, H., Messmer, B.T.: Similarity measures for structured representations. In: Wess, S., Richter, M., Althoff, K.-D. (eds.) *EWCBR 1993*. LNCS, vol. 837, pp. 106–118. Springer, Heidelberg (1994)
5. Casati, F., Ceri, S., Pernici, B., Pozzi, G.: Workflow evolutions. *Data and Knowledge Engineering* 24, 211–238 (1998)
6. Heimann, P., Joeris, G., Krapp, C., Westfechtel, B.: Dynamite: dynamic task nets for software process management. In: *Proceedings International Conference of Software Engineering*, Berlin, pp. 331–341 (1996)
7. Inzitari, D., Carlucci, G.: Italian stroke guidelines (spread): evidence and clinical practice. *Neurological Sciences* 27, s225–s227 (2006)
8. Kapetanakis, S., Petridis, M., Ma, J., Bacon, L.: Workflow monitoring and diagnosis using case based reasoning on incomplete temporal log data. In: *Proc. Uncertainty, knowledge discovery and similarity in Case Based Reasoning Workshop, International Conference on Case Based Reasoning (ICCBR)*, pp. 171–180 (2009)
9. Kurtz, S.: Approximate string seraching under weighted edit distance. In: *Proc. 3rd South American Workshop on String Processing*, Recife. Carlton University Press (1996)
10. Luo, Z., Sheth, A., Kochut, K., Miller, J.: Exception handling in workflow systems. *Applied Intelligence* 13, 125–147 (2000)
11. Thao Ly, L., Rinderle, S., Dadam, P.: Integration and verification of semantic constraints in adaptive process management systems. *Data & Knowledge Engineering* 64(1), 3–23 (2008)
12. Madhusudan, T., Zhao, J.L., Marshall, B.: A case-based reasoning framework for workflow model management. *Data and Knowledge Engineering* 50, 87–115 (2004)
13. Mans, R., Schonenberg, H., Leonardi, G., Panzarasa, S., Cavallini, A., Quaglini, S., van der Aalst, W.: Process mining techniques: an application to stroke care. In: Andersen, S., Klein, G.O., Schulz, S., Aarts, J. (eds.) *Proc. MIE, Studies in Health Technology and Informatics*, vol. 136, pp. 573–578. IOS Press, Amsterdam (2008)
14. Minor, M., Tartakovski, A., Schmalen, D., Bergmann, R.: Agile workflow technology and case-based change reuse for long-term processes. *International Journal of Intelligent Information Technologies* 4(1), 80–98 (2008)

15. Montani, S.: Prototype-based management of business process exception cases. *Applied Intelligence* (published online in February 2009), doi:10.1007/s10489-009-0165-z
16. Page, R., Holmes, M.: *Molecular Evolution: A Phylogenetic Approach*. Wiley, Chichester (1998)
17. Palmer, M., Wu, Z.: Verb Semantics for English-Chinese Translation. *Machine Translation* 10, 59–92 (1995)
18. Reichtert, M., Rinderle, S., Dadam, P.: Adept workflow management system: Flexible support for enterprise-wide business processes. In: van der Aalst, W.M.P., ter Hofstede, A.H.M., Weske, M. (eds.) *BPM 2003*. LNCS, vol. 2678, pp. 370–379. Springer, Heidelberg (2003)
19. Resnik, P.: Using information content to evaluate semantic similarity in a taxonomy. In: *Proc. IJCAI*, pp. 448–453 (1995)
20. Surma, J., Vanhoof, K.: Integration rules and cases for the classification task. In: Aamodt, A., Veloso, M.M. (eds.) *ICCBR 1995*. LNCS, vol. 1010, pp. 325–334. Springer, Heidelberg (1995)
21. van Elst, L., Aschoff, F.R., Barbardi, A., Maus, H., Schwarz, S.: Weakly-structured workflows for knowledge-intensive tasks: An experimental evaluation. In: *Proc. of 12th IEEE International Workshops on Enabling Technologies (WETICE), Infrastructure for Collaborative Enterprises*, pp. 340–345. IEEE Computer Society, Los Alamitos (2003)
22. Van der Aalst, W., ter Hofstede, A., Weske, M.: Business process management: a survey. In: van der Aalst, W.M.P., ter Hofstede, A.H.M., Weske, M. (eds.) *BPM 2003*. LNCS, vol. 2678, pp. 1–12. Springer, Heidelberg (2003)
23. Van der Aalst, W., van Dongen, B., Herbst, J., Maruster, L., Schimm, G., Weijters, A.: Workflow mining: a survey of issues and approaches. *Data and Knowledge Engineering* 47, 237–267 (2003)
24. van Dongen, B., Alves De Medeiros, A., Verbeek, H., Weijters, A., Van der Aalst, W.: The proM framework: a new era in process mining tool support. In: Ciardo, G., Darondeau, P. (eds.) *Knowledge Management and its Integrative Elements*, pp. 444–454. Springer, Heidelberg (2005)
25. Weber, B., Reichert, M., Wild, W.: Case-based maintenance for CCBR-based process evolution. In: Roth-Berghofer, T.R., Göker, M.H., Güvenir, H.A. (eds.) *ECCBR 2006*. LNCS (LNAI), vol. 4106, pp. 106–120. Springer, Heidelberg (2006)
26. Weber, B., Wild, W.: Towards the agile management of business processes. In: Althoff, K.-D., Dengel, A.R., Bergmann, R., Nick, M., Roth-Berghofer, T.R. (eds.) *WM 2005*. LNCS (LNAI), vol. 3782, pp. 409–419. Springer, Heidelberg (2005)

A Survey on the Dynamic Scheduling Problem in Astronomical Observations

Matias Mora^{1,2} and Mauricio Solar¹

¹ Universidad Técnica Federico Santa María,
Computer Systems Research Group, Valparaíso, Chile
`mmora@csrg.inf.utfsm.cl`

² Atacama Large Millimeter/submillimeter Array,
Santiago, Chile

Abstract. The tasks execution scheduling is a common problem in computer science. The typical problem, as in industrial or computer processing applications, has some restrictions that are inapplicable for certain cases. For example, all available tasks have to be executed at some point, and ambient factors do not affect the execution order.

In the astronomical observations field, projects are scheduled as observation blocks, and their execution depends on parameters like science goals priority and target visibility, but is also restricted by external factors: atmospheric conditions, equipment failure, etc. A telescope scheduler is mainly in charge of handling projects, commanding the telescope's high level movement to targets, and starting data acquisition. With the growth of observatories' capacities and maintenance costs, it is now mandatory to optimize the observation time allocation. Currently, at professional observatories there is still strong human intervention dependency, with no fully automatic solution so far.

This paper aims to describe the dynamic scheduling problem in astronomical observations, and to provide a survey on existing solutions, opening some new application opportunities for computer science.

1 Introduction: The (Dynamic) Scheduling Problem

In general terms, the scheduling problem is an assignment problem of a number of tasks or jobs, which execution may be restricted by one or more constraints or performance criterion. A survey of the most important results in (multi-criterion) scheduling, and a summary of basic concepts and notations, have been presented in [8], starting with first researches in the mid 50s. The most classical scheduling problem is the *job-shop*, where N ideal jobs are assigned to M identical resources (machines, processors, etc.), while trying to minimize the total execution time.

Besides the static problem, the dynamic scheduling has been frequently researched in recent years. It presents the challenge of processing the schedule on-line, while new jobs are coming into the queue, or new restrictions appear. A complete survey on the dynamic problem for manufacturing systems has been presented in [15], concluding that Multi-Agent Systems are one of the most

promising techniques, although its application needs to be further developed in the future. In general terms, resolution techniques should match a balance between flexibility and robustness.

A simple example of dynamic scheduling can be found in [7]. It is proposed to divide the dynamic job-shop scheduling problem under fuzzy environment into static fuzzy sub-scheduling problems, and uses an improved Giffler & Thompson algorithm is used to solve them. A different approach using a Neuro-Fuzzy Network (NFN) is proposed in [20]. A fuzzy classification of jobs is made in real-time, and then scheduled with a neural network, trained off-line with a genetic algorithm. A simulation shows that this is effective and highly efficient compared to other algorithms (FIFO and Lagrangian Relaxation). As mentioned before, Multi-agent systems have also been used to solve dynamic scheduling decision systems. For example, [3] uses agents to explore more useful decisions in the variable environment, constantly updating the knowledge base. The result is a very flexible algorithm that will adapt itself to the environment. Detected problems are the slow convergence of the algorithm, and the relatively slow response to the environment.

A variation of the dynamic problem is scheduling with dynamic priorities, which has been widely used in CPU and cluster processes scheduling. As described in [11], the main idea is to optimize the processor utilization by dynamic priorities assigned according to deadlines, obligating pending jobs to be executed in a reasonable time. Priorities in these cases are adjusted by the scheduler itself, and not directly by external factors.

Finally, [8] mentions some interesting new development lines, outside the traditional scopes: “interfering job sets” (jobs that have to be run at the same time on the same machine), “scheduling with rejection” (allows the possibility to not run all jobs, by rejecting some of them) and “rescheduling for new orders” (introducing new conditions and jobs). As we will see, all these relatively new developments are related to the problem described in this paper.

2 Astronomical Observations Scheduling

Astronomical observations require specific conditions for their execution, such as the instrument to be used, visibility, sky brightness, etc. All this information, together with the scientific goals of the observation, are presented by an astronomer in a so called *proposal* to apply for observation time. Its format can vary from one institution to another, including the following fields: telescope and instrument (one telescope can work with more than one instrument), main investigator, program description and target(s) list.

In the case of Chile, there are three main institutions managing some of the world’s most important telescopes: European Southern Observatory, ESO (La Silla, Paranal, APEX); Association of Universities for Research in Astronomy, AURA (Tololo, Gemini, SOAR); and Observatories of the Carnegie Institute of Washington, OCIW (Las Campanas). Proposals have to be sent to the corresponding Telescope Assignment Committee, which evaluates all proposals, assigning a scientific priority (importance of execution), and approving or rejecting

the requested observing times. As most observatories are joint ventures between several organizations and/or countries, the list of approved projects has to comply with the percentage of total time assigned to each part. Normally, telescope time can be applied once per observations period. An observation can be executed in visitor or service mode. Visitor mode observations require the presence of the main investigator (or a collaborator) on site to take the data, while service mode observations are executed by the telescope operator and observatory's science staff members.

Approved proposals will be later scheduled as one or more observation blocks for execution on the corresponding telescope and instrument. The execution of each observation depends on external factors specified in the proposal, and the automatic scheduler and/or telescope operator will have to decide the best time for each case. This problem can consider various thousands of observations per period, although not all targets will be visible during the entire time. Normally, a long term plan will consider factors like visibility over the horizon and moon brightness, while a short term plan will consider specific observation blocks for the next night(s), based on more immediate factors. For a detailed description of a typical observation process at ESO observatories, see [21]. As it could be concluded from this *modus operandi*, it will certainly happen that not all approved observation blocks are executed in a given period. The execution probability of a project will depend on its science priority, but also on the fulfillment of needed conditions for its observation. In order to complete a project, all its observation blocks have to be successfully executed.

The scheduling of astronomical observations is a variation of the dynamic scheduling problem, which has been treated in various ways since several decades by the scientific community. This is a multi-objective problem, as normally various optimizations are required. The most important ones are the maximization of the executed scientific priorities (scientific throughput), and the maximization of observing time usage. This includes minimizing gaps between executions (including readout time, telescope movement to the new source, instrument change and/or calibrations, etc.), and carefully planning required maintenance. Also, the total exposure time of an observation may depend on atmospheric conditions, as it could be necessary to do larger exposures with poor visibility.

Although most of modern professional observatories use certain degree of scheduling automation, there is still a huge part of human intervention to build up the daily plan and to take last minute decisions. External parameters can vary at any time during observations, and therefore a dynamic re-scheduling is needed. If we consider a given execution priority for each block, depending on the quality of observation conditions and importance of scientific goals, external parameters can certainly cause priority changes. Some observations may not be anymore suitable to execute under new conditions. Moreover, as observation blocks depend on target's visibility on the sky, they might be only valid during certain day/night time, and/or certain periods of the year. Therefore, it might happen that initially high priority tasks have to be executed with less priority, or cannot be executed at all within one observations period. Particular

observation blocks may also depend on others to be executed. For example, it may be required to execute blocks sequentially or with a certain frequency.

3 Current Approaches

The common astronomy scheduling problem is NP-hard. A good description and basic mathematical model is presented in [2]. In the same publication, the author proposes long and short-term scheduling scopes, and tries to solve the problem using neighborhood search (Lin-Kernighan heuristic) and genetic algorithms. The result is that under certain circumstances (short size and good pre-ordered sample) the neighborhood search can perform better than the genetic algorithm. Nevertheless, the genetic algorithm is in general a better and faster alternative, and does not need a pre-knowledge of the main constraints. The scientific policy (committee) imposes some restrictions that are difficult to handle, depending strongly on the sample characteristics (proposal quality and duration). To take better advantage of automatic scheduling, it is important to have a small degree of over-subscription in the final allocated time, and also a large number of short exposure (duration) project proposals.

Figure 1 shows a typical observations scheduling process, with a list of approved, and *conditionally* approved (over-subscription) observation blocks. The scheduler will sort these blocks according to their execution priority into a queue, and then allocate them in the available time slots. According to varying conditions, this process has to be repeated and the schedule properly adapted.




Fig. 1. Example astronomical observations scheduling process (*source: self elaboration*)

The most referenced scheduling solution for astronomical observations is the SPIKE scheduler for the Hubble Space Telescope, developed by the Space Telescope Science Institute (STScI). SPIKE is largely cited as a reference scheduling system, and has also been adapted to other (ground based) telescopes. The current trend is to increase the observations automation, as astronomical projects are getting more complex, observation time more expensive, and decisions more difficult. Some scheduling approaches for current professional projects are discussed below.

3.1 Hubble Space Telescope (HST)

The HST is probably the most famous space telescope, launched in 1990, and best known for its exploration of the deep space from the Earth orbit. It is a collaboration between NASA and the European Space Agency. Space telescopes have the advantage of not depending on atmospheric interference, but are also much more complex to maintain and repair. The HST SPIKE scheduling system, described in [10] and [23], treats the scheduling as a constraint satisfaction problem (CSP), including a toolkit to handle this type of problems. Short and long term scheduling concepts are applied, and several schedule steps are considered: trial assignment heuristic (min-conflicts times), repair heuristic (neural network) and de-conflict (priority selection). Also, rescheduling of observations is possible through the CSP toolkit (task locking and conflict-cause analysis). Since its original implementation in 1987 SPIKE is entirely implemented in Common Lisp and Common Lisp Interface Manager for user interfaces. [13] presents a report about studies related to the generalization of constraint-based scheduling theories and techniques with application to space telescope observation scheduling. For this goal, the Heuristic Scheduling Testbed System (HSTS) planning and scheduling framework was developed. Planning and scheduling are treated as complimentary problems to produce good results. This was validated by producing observation schedules for the HST.

3.2 Very Large Telescope (VLT)

The VLT, operated by ESO, is one of the world's largest optical telescopes, with four 8.2 meter aperture telescopes, designed to work separately or as a single instrument: the VLT Interferometer. It is located in the northern Chilean dessert, on top of the Paranal mountain. The early observations plan for the VLT is discussed in [9], describing the SPIKE scheduling tools, initially developed for the HST, and the requirements to adapt it for VLT use and other ground based observatories. The automated scheduling is thought just as an assistant for human decisions on value judgments, mixing visitor and service mode observations. Nevertheless, it is an important step starting from the visitor only mode in La Silla. Also, first concepts of artificial intelligence are introduced, and envisioned as a good way to significantly reduce calculation times. Later, [21] describes the proposal selection, prioritizing and scheduling process at ESO observatories. Specially interesting is the mention of many factors that normally affect observations, such as time of the year, technical down-time and lunar phase (brightness). Observations are organized in periods of 6 months each.

3.3 Subaru Telescope

This 8.2 meters telescope is located on the top of the Mauna Kea mountain in Hawaii, together with many other large telescopes. [17] describes the (partially implemented) scheduling system for the Subaru Telescope, part of the Subaru Observation Software System (SOSS). It uses SPIKE as scheduling engine, with

various extensions to the VLT adaptation, for particular needs. The main objective is to provide “computer-aided planning of observation proposals”, and decision support for the Telescope Allocation Committee and the observers.

3.4 Gemini Observatory

The Gemini Observatory is made up of two almost identical 8.1 meters telescopes, located in both hemispheres: Chilean Pachón mountain and Hawaiian Mauna Kea. Simulation results for the Gemini project scheduler are presented in [16]. The general observations plan description is similar to [21], but with less institutional experience. This work focuses on how to distribute observing time in the future, based on ambient factors, but also the partners’ proportion. Observable nights are divided into 6-months periods, and A and B types, depending on sky quality.

3.5 Stratospheric Observatory for Infrared Astronomy (SOFIA)

The SOFIA telescope is a collaboration between NASA and the German Aerospace Center. The telescope is mounted on a Boeing 747 aircraft, and observes during flight time. A description of the various computational challenges of the SOFIA project is presented in [6]. In this very special case, scheduling includes flights planning (duration, position and speed, gas consumption, etc.), and also the observations set for each flight. On the previous project of this kind, KAO (Kuiper Airborne Observatory), the entire planning was done by hand and took about 8 hours for each flight. The idea for SOFIA is to assist this process by computational algorithms, rather than replacing human flight planners. In this context, similar experiences from the space are cited: HST, Earth Observing Satellites and Mars Exploring Rovers.

3.6 The Robert C. Byrd Green Bank Telescope (GBT)

The GBT is the world’s largest fully movable radio telescope antenna. It is located in West Virginia, and operated by the National Radio Astronomy Observatory (NRAO). A prototype automatic scheduling system has been recently introduced: the GBT Dynamic Scheduling System (DDS), described in [14]. Other than in other dynamic scheduling outlines, this telescope is very interactive and requires the presence of the observing astronomer. Therefore, “the Dynamic Scheduling System is scheduling people rather than scripts”, and predictions are an important need, mainly achieved through weather forecasts and weather data from the last four years. The candidates determination for certain periods based on this data is described in [1]. The scoring algorithm in the GBT DDS assigns a priority number to each proposal, taking specially into account the weather predictions, but also stringency and observing efficiency. The actual scheduling algorithms are described by [19], in a 2-phase approach: Sudoku solver for fixed and windowed sessions, and Knapsack algorithm for optimal schedules of remaining time intervals. In conclusion, by dynamic scheduling it is possible to substantially improve telescope time usage efficiency.

4 The ALMA Scheduling Problem

The Atacama Large Millimeter/submillimeter Array (ALMA) is a major collaboration effort between European, North American and East Asian countries, under construction on the Chilean Chajnantor plateau, at 5.000 meters altitude. When completed in 2013 it will be the largest radio-telescope on earth, with more than 60 antennas of 12 and 7 meters diameter, distributed over a wide extension, with up to 16 kilometers of baseline separation. The ALMA interferometer will provide the possibility to be used as a single array, or as up to six (due to limited centralized equipment) minor independent arrays or groups of antennas. As each array is equivalent to one instrument, this can be seen as a multi-telescope problem. Also, the antennas will be changing their position during the year, as different distributions will be used to exploit various kinds of observations. As ALMA is a radio-telescope, observations are not limited to nighttime, and a 24/7 operation with as small downtime as possible is expected.

ALMA will operate exclusively in service mode. Therefore, the Scheduling Subsystem is supposed to provide a fully automatic quasi-real-time dynamic scheduling platform, mostly with the only human participation of supervision. This subsystem is still under development, and the needed algorithm(s) not defined yet. The main problem is the dynamic priorities scheduling, which differs widely from the traditional dynamic job-shop. This particular problem is very similar for all ground based observatories, and ALMA is one real example, that needs this kind of scheduling to accomplish its operations requirements.

A huge part of the telescope operations will be handled through the ALMA software, which is divided into various subsystems, such as Control, Correlator, Pipeline, Archive, etc. (for a complete overview of the ALMA software architecture, see [18]). The Scheduling Subsystem is the one in charge of managing antenna arrays and executing observation blocks dynamically. Its design was first discussed in [22], where a prioritized list of observations is scheduled in real-time. After a critical retrospective of automatic scheduling it is recommended that ALMA should still use human intervention to schedule observations, but assisted by a dynamic scheduler which should be able of adjusting parameters according to weather conditions. A prototype time allocation program called *taco* was implemented and tested. This was not successful, as it wasn't able to adjust parameters (static behavior) and the dynamic decision queue was not suitable for actual observations, as it didn't consider astronomers' availability (which has been later discarded as exclusively service mode will be used). More recently, [5] presented the architecture for the ALMA dynamic scheduler (sections 5 and 7), with emphasis on the scheduler simulation, and some sort of detail about the algorithm implementation (section 6), and current simulated results. This paper is also a good introduction to some basic dynamic scheduling concepts (sections 1 to 4), citing mainly [23]. Furthermore, [12] is an extension to the previous paper, concentrating on scheduling policies and policy factors, like the observations success probability formula.

The Scheduling Subsystem requirements and design are detailed in the "Scheduling Subsystem Design Document" [4]. The current design considers four

main parts: Master Scheduler, Project Manager, Scheduler and Simulator (see Figure 2). Observation proposals are represented as projects, divided into Scheduling Blocks (SB), which are handled as a queue by the Master Scheduler. The queue is created and maintained by the Project Manager, which will also start the reduction pipeline once an observation is completed. SBs from the master queue are assigned to different Schedulers, which will execute them on an associated antennas array (one of these Schedulers is assigned to one array). A Scheduler will be able to operate in interactive or dynamic (automatic) mode. Additionally, the Simulator will be used to simulate the overall scheduling behavior, for planning purposes and performance tuning. The dynamic mode Schedulers are the main problem, which goal consist in providing a full automation of SB execution, with manual override capabilities. The general requirements of the ALMA scheduler include to be able to handle about 18.000 SBs per period (~ 12 months), and to do a rescheduling process in much less than 1/2 of a typical scheduling block length (~ 30 minutes). The Scheduling Subsystem is currently under development, and the defined architecture could still change in the near future.

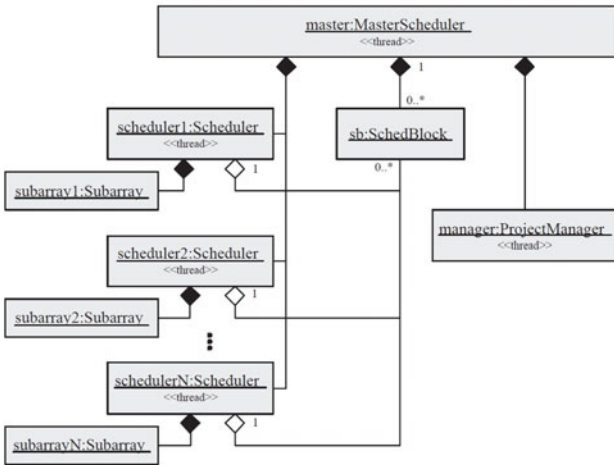


Fig. 2. ALMA scheduling system (source: [5])

As each array can be formed by any combination of antennas (different combinations make actually a difference depending on the physical positions), the number of possible junctions for a 1-array problem is in theory $O(n!)$, with n the number of available antennas. Nevertheless, it is expected that most of the major observations will use all available antennas to exploit their full potential. A simplification of this kind would be still a very useful approach.

Different alternatives need to be analyzed for the resolution of the ALMA problem, first considering the automation requirement, but also the distributed nature of the project, mainly regarding the varying antenna locations. Although

this problem differs from the manufacturing job-shop, we can see some correlation in the recent development of both areas, in terms of changing environment and even the distribution of jobs and processing machines. It would be therefore recommendable to start with techniques that have been successful for distributed job-shops, like Multi-Agent Systems (see [15]).

5 Conclusions

The astronomical observations scheduling problem has been presented, together with some of its most extreme variants, and a current example case description for the ALMA project. As it has been identified, this new problem considers some non-traditional elements, like interfering jobs, rejection possibility and rescheduling for new order.

None of the existing professional astronomical scheduling solutions provides a full automatic scheduler under changing conditions and for several telescopes, and none of the existing algorithms has solved the dynamic priorities problem for this case. As new astronomical projects will depend on scheduling optimization to make the best use of limited and expensive observing time, this problem is an interesting research topic to be analyzed from many different views. The new conditions for the scheduling problem are the major engineering challenges for this case.

This is just a first overview, to motivate further work on this topic. A first approach of specialization in this area, for current and future projects, is being done by the Computer Systems Research Group (CSRG) at Universidad Técnica Federico Santa María in Chile, in collaboration with the ALMA Scheduling Subsystem development team and the UTFSM Informatics Department.

Acknowledgements

The authors would like to thank the revisions and contributions by UTFSM and ALMA colleagues, and the ALMA-Conicyt Fund #31080031.

References

- [1] Balser, S., et al.: GBT Dynamic Scheduling System: Algorithms, Metrics and Simulations. In: *Astronomical Society of the Pacific Conference Series. Astronomical Data Analysis Software and Systems XVIII*, vol. 411, p. 330 (2009)
- [2] de Castro, A.I.G., Yanez, J.: Optimization of telescope scheduling: Algorithmic research and scientific policy. *Astronomy and Astrophysics* 403, 357–367 (2003)
- [3] Chao, W., Zhen-qiang, B., Chang-yi, L., Wen-yu, B.: A Self-adaptive Knowledge Based Dynamic Scheduling Decision Mechanism. In: *International Conference on Management Science and Engineering*, pp. 463–467 (October 2006)
- [4] Clarke, D., Lucero, S., Farris, A.: Scheduling Subsystem Design Document. Tech. Rep. COMP-70.55.00.00-003-E-DSN, National Radio Astronomy Observatory, Socorro, NM (May 2009)

- [5] Farris, A.: Scheduling simulation in ALMA. In: Proceedings of SPIE 2004. Optimizing Scientific Return for Astronomy through Information Technologies, vol. 5493, pp. 302–312 (June 2004)
- [6] Frank, J.: SOFIA’s challenge: automated scheduling of airborne astronomy observations. In: Second IEEE International Conference on Space Mission Challenges for Information Technology, SMC-IT (2006)
- [7] Fu-ming, L., Yun-long, Z.: Study on the Job Shop Dynamic Scheduling under Fuzzy Environments. In: International Conference on Computational Intelligence and Security, vol. 1, pp. 90–96 (November 2006)
- [8] Hoogeveen, H.: Multicriteria scheduling. *European Journal of Operational Research* 167, 592–623 (2005)
- [9] Johnston, M.D.: Automated observation scheduling for the VLT. Tech. rep., Space Telescope Science Inst., Baltimore (1988)
- [10] Johnston, M.D.: Spike: AI scheduling for NASA’s Hubble Space Telescope. In: Sixth Conference on Artificial Intelligence Applications, vol. 1, pp. 184–190 (1990)
- [11] Liu, C.L., Layland, J.W.: Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environment. *Journal of the ACM* 20, 46–61 (1973)
- [12] Lucero, S., Farris, A.: Dynamic Scheduling in ALMA. In: Astronomical Society of the Pacific Conference Series. Astronomical Data Analysis Software and Systems XVI, vol. 376, p. 673 (2007)
- [13] Muscettola, N., et al.: Investigations into Generalization of Constraint-Based Scheduling Theories with Applications to Space Telescope Observation Scheduling. Tech. rep., The Robotics Inst., Carnegie-Mellon University, Pittsburgh (1996)
- [14] O’Neil, K., et al.: The GBT Dynamic Scheduling System: A New Scheduling Paradigm. In: Astronomical Society of the Pacific Conference Series. Astronomical Data Analysis Software and Systems XVIII, vol. 411, p. 147 (2009)
- [15] Ouelhadj, D., Petrovic, S.: A survey of dynamic scheduling in manufacturing systems. *Journal of Scheduling* 12, 417–431 (2009)
- [16] Puxley, P.J.: Execution of queue-scheduled observations with the Gemini 8-m telescopes. In: Proceedings of SPIE 1997. Telescope Control Systems II, vol. 3112, pp. 234–245 (July 1997)
- [17] Sasaki, T., et al.: Observation scheduling scheme for the Subaru telescope. In: Proceedings of SPIE 2003. Advanced Telescope and Instrumentation Control Software, vol. 4009, pp. 350–354 (June 2003)
- [18] Schwarz, J., Farris, A., Sommer, H.: The ALMA software architecture. In: Proceedings of SPIE 2004. Advanced Software, Control, and Communication Systems for Astronomy, vol. 5496, pp. 190–204 (June 2004)
- [19] Sessoms, E., et al.: The GBT Dynamic Scheduling System: Scheduling Applications of the Knapsack Problem and Sudoku. In: Astronomical Society of the Pacific Conference Series. Astronomical Data Analysis Software and Systems XVIII, vol. 411, p. 351 (2009)
- [20] Shugang, L., et al.: Job shop scheduling in real-time cases. *The International Journal of Advanced Manufacturing Technology* 26, 870–875 (2005)
- [21] Silva, D.R.: Service mode scheduling at the ESO Very Large Telescope. In: Proceedings of SPIE 2002. Observatory Operations to Optimize Scientific Return III, vol. 4844, pp. 94–103 (June 2002)
- [22] Wright, M.C.H.: Dynamic Scheduling. Tech. rep., Radio Astronomy Laboratory, University of California, Berkeley (December 1999)
- [23] Zweben, M., Fox, M.S.: SPIKE: Intelligent Scheduling of HST Observations, ch. 14, pp. 391–422. Morgan Kaufmann Publishers, San Francisco (1994)

PART VI
Intelligent Systems Engineering

Combining Determinism and Intuition through Univariate Decision Strategies for Target Detection from Multi-sensors

Edwin El-Mahassni

ISR Division
Defence Science and Technology Organisation
Edinburgh, SA, 5111, Australia
Edwin.El-Mahassni@dsto.defence.gov.au

Abstract. In many surveillance systems, an operator will have to make a decision on whether a target is present or not from the outputs of a sensor. Here, we assume that such an operator has at his or her disposal two sensors to give him or her more confidence in his decision-making process. In addition, we propose several univariate decision strategies which combine sensor characteristics, target probabilities and reward schemes together with deterministic and intuitive operator decision-making parameters. Further, using Chebyshev's inequality we develop a method for selecting the most robust strategy that mitigates poor performance. We illustrate through examples how different strategies will be more advantageous over others depending on the reward scheme and sensor parameters.

Keywords: Decision-making, Target detection, Intuition.

1 Introduction

The ability to make decisions is ubiquitous in many defence settings. Sometimes, an operator has to make a deterministic decision based on the outputs of a sensor or sensors. On the other hand, an operator's experience should not be discounted and his intuition in events should also be used [1]. To add to the complexity of decision-making, an operator has to be aware of the rewards and costs associated with making the right or wrong decision [2,3].

In general, there are two distinct philosophical systems toward decision making. These are the rational (or logical or analytical) vs. the naturalistic (or intuitive) [4] techniques. The latter are based upon an expert's background in dealing with similar situations in the past [1] and assume that experience is vital in making the right decision. The solutions are based on finding a satisficing outcome, rather than the optimal. In fact, it was found that over 95% of naval commanders and army officers use their intuition. We also remark that intuitive decision-making is being used to replace the standard military decision-making process [1].

More mathematical approaches can be found in [5], where the reward of making the right decision diminishes the longer it takes for the decision to be made. Other approaches which also claim to be rational-based can be found in [6,7], where criteria and weights are used; though the fact that these are set by humans themselves makes their claim to objectivity weak [5].

Here, we derive and examine some decision strategies that combine determinism (where the same decision is always made given the same conditions) and intuition (greater flexibility with no fixed rules) for an operator to make a judgement from two sensors. For target detection we recall the well-known formula for the expected value of a decision strategy found in [2]. This formula is not only dependant on detection, false alarm and target probabilities, but also considers the rewards and costs of correct and incorrect decision-making. The decision strategies listed here will all consist of one only variable which we define as x ; and it will be used to quantify an operator's probability of deciding when a target is present for some scenarios. We remark that at all times, the distribution of x is known and could be predetermined through testing or simulations. Further, any knowledge, such as sensor parameters, target probability, reward schemes, etc. that the operator might have is assumed to be accommodated within this distribution.

Therefore, if an operator has to make the same choice for a particular multi-sensor reading we deem this to be a deterministic decision. However, as indicated in [1], the benefits of intuitive decision-making should not be discounted; with the variable x being used to model this intuition. On the other hand, because we are also relying on personal judgement, we should not rule out the effects that fatigue or his emotions can have in his decision-making ability [8]. In these cases, poor performances are modelled by all the expected values of a decision strategy $E(x)$ which are themselves given by varying x from its optimal value of 0 or 1. Alternatively, even within experienced operators, there could be different biases and life experiences which could result in different decisions being made for the same scenario [9]. Thus, we are interested in analysing decision strategies which can accomodate some intuition, and then select the most robust strategy when poor performance ensues. We note, however, that this given freedom to make decisions will be restricted to only some sensor readings, thereby combining the intuition or the experience of an operator with the certainty of determinism. The following two assumptions are made: 1) when an operator has poor or underperforming intuitive skills, then the expected value of his decision strategy is given by the mentioned formula in [2] and 2) when intuition aids an operator, then we assume this expected value is greater than that given by the same formula.

The paper will be divided as follows. In section 2, we outline 5 decision strategies and then derive their respective expected value formulae. Section 3 is devoted to analysis and description of the algorithm; in section 4 we provide some examples, and finally in section 5 we provide some concluding remarks.

2 Univariate Decision Strategies

In this section, we recall the original expected value of a decision strategy function and then proceed to define some univariate decision strategies based on the output of 2 sensors.

2.1 The Decision Value Function

We note that originally the expected value function was framed in terms of signal detection theory [2] and whether a target was present or not. That is, first, we let h_0 and h_1 denote the hypotheses indicating the absence or presence of a target respectively. And, we let H_0 and H_1 be the response of the decision maker accepting the h_0 or h_1 hypothesis respectively. And, we also define the following:

- Let V_{00} be the reward value associated with a correct choice of H_0 , which occurs with probability $p(H_0|h_0)$. This is also known as a correct rejection.
- Let V_{01} be the cost value associated with an incorrect choice of H_1 (when, in fact, H_0 is the correct alternative); that is, the person loses V_{01} when this type of incorrect choice is made which occurs with probability $p(H_1|h_0)$. This is also known as a false alarm.
- Let V_{11} be the reward value associated with a correct choice of H_1 , which occurs with probability $p(H_1|h_1)$. This is also known as a hit.
- Let V_{10} be the cost value associated with an incorrect choice of H_0 (when, in fact, H_1 is the correct alternative); that is, the person loses V_{10} when this type of incorrect choice is made, which occurs with probability $p(H_0|h_1)$. This is also known as a miss.

Thus, the expected value of the decision strategy is given by

$$E = V_{00}P(h_0)P(H_0|h_0) + V_{11}P(h_1)P(H_1|h_1) - V_{10}P(h_1)P(H_0|h_1) - V_{01}P(h_0)P(H_1|h_0) \quad (1)$$

noting that of course $\sum_{i=0,1} P(H_i|h_j) = 1$ for $j = 0, 1$.

2.2 Decision Strategies

Now, in formulating decision strategies from sensor outputs, we need to make a link between $p(H_i|h_j)$, $i, j \in \{0, 1\}$ and $p(H_i|(s^1, s^2))$, where $s^1, s^2 \in \{0, 1\}$. The values 1 and 0 for (s^1, s^2) indicate the state where a sensor reports that a target is present or not respectively. But, we also note that the sensors' reports are dependant on whether there is an actual target present or not. Thus, target probability will also influence sensor reporting. These reports then form the basis of any judgement by an analyst or operator. This relationship can be expressed below as:

$$p(H_i|h_j) = \sum_{s^1, s^2 \in \{0, 1\}} p(H_i|(s^1, s^2))p((s^1, s^2)|h_j). \quad (2)$$

If we now define d^i and f^i to be the probabilities of detection and false alarm for sensor i respectively, then we have that

$$\begin{aligned} p((1, 1)|h_0) &= f^1 f^2 & , p((1, 1)|h_1) &= d^1 d^2 \\ p((0, 1)|h_0) &= (1 - f^1) f^2 & , p((0, 1)|h_1) &= (1 - d^1) d^2 \\ p((1, 0)|h_0) &= f^1 (1 - f^2) & , p((1, 0)|h_1) &= d^1 (1 - d^2) \\ p((0, 0)|h_0) &= (1 - f^1)(1 - f^2) & , p((0, 0)|h_1) &= (1 - d^1)(1 - d^2). \end{aligned}$$

Next, we outline 5 decision strategies based on variations of $p(H_i|(s^1, s^2))$ through the use of one variable. But, first, for brevity, we denote by $\tilde{s} = (s^1, s^2)$ as the instantaneous outputs of 2 sensors. Lastly, for all instances we let the variable $0 \leq x \leq 1$ be the probability that an operator will make a decision that a target is present for some \tilde{s} . The 5 decision strategies, remembering that $p(H_1|\tilde{s}) + p(H_0|\tilde{s}) = 1$, now follow, and they are used to model the mean decision value when an operator performs poorly; in the sense that he or she is not using intuition. The first strategies assume that $p(H_1|\tilde{s}) = 0$ if $\tilde{s} = (0, 0)$. This means that we assume that in this instance no target can possibly be present when neither sensor reports. The last two strategies, on the other hand, imply that $p(H_1|\tilde{s}) = 1$ if $\tilde{s} = (1, 1)$ always indicating the presence of a target. The third strategy assigns a deterministic choice when the sensors agree, and leaves an operator to make non-deterministic decisions when only 1 sensor reports. We denote for each strategy the value $x_i, i = 1, \dots, 5$ to stress the fact x is not necessarily the same value for all decision strategies; although some dependencies within different strategies could also be accommodated. Further, the $x_i, i = 1, \dots, 5$ variables might lie between different upper and lower bounds also.

1. The first strategy is given by

$$p(H_1|\tilde{s}) = \begin{cases} x_1, & \tilde{s} = (1, 1) \\ 0, & \tilde{s} \in \{(0, 0), (0, 1), (1, 0)\} \end{cases} \quad (3)$$

Here, we label this case as very pessimistic. This is because we consider the sensors to be so unreliable that we do not even consider a target being present unless both sensors report. The operator is told to ignore all the cases where both sensors do not report and only use his discretion when both indicate the presence of a target. Combining (1) and (2), we have that the expected value is:

$$\begin{aligned} E(x_1) &= x_1(p(h_1)d^1 d^2(V_{11} + V_{10}) - p(h_0)f^1 f^2(V_{00} + V_{01})) \\ &\quad + p(h_0)V_{00} - p(h_1)V_{10}. \end{aligned}$$

2. The second strategy is given by

$$p(H_1|\tilde{s}) = \begin{cases} x_2, & \tilde{s} \in \{(1, 1), (0, 1), (1, 0)\} \\ 0, & \tilde{s} = (0, 0) \end{cases} \quad (4)$$

Here, we label this case as pessimistic. This strategy stipulates that even if only 1 sensor reports, the operator, to be on the safe side should treat this scenario on the same level as if they both reported. However, just like in the previous strategy, when both sensors do not report, it is assumed that a target is not present. Combining (1) and (2), we have that the expected value is:

$$E(x_2) = x_2 (p(h_0)((1 - f^1)(1 - f^2) - 1)(V_{01} + V_{00}) + p(h_1)(1 - (1 - d^1)(1 - d^2))(V_{10} + V_{11})) + p(h_0)V_{00} + p(h_1)V_{10}$$

3. The third strategy is given by

$$p(H_1|\tilde{s}) = \begin{cases} 1, & \tilde{s} = (1, 1) \\ x_3, & \tilde{s} \in \{(0, 1), (1, 0)\} \\ 0, & \tilde{s} = (0, 0) \end{cases} \quad (5)$$

Here, we label this case as the unbiased case. The operator makes fully deterministic judgements when both sensors agree and is asked to use his discretion only when they disagree in their outputs. Combining (1) and (2), we have that the expected value is:

$$E(x_3) = x_3 (p(h_0)(V_{00} + V_{01})(2f^1f^2 - f^1 - f^2) + p(h_1)(V_{11} + V_{10})(d^1 + d^2 - 2d^1d^2) + p(h_0)(V_{00} - (V_{00} + V_{01})f^1f^2) + p(h_1)(V_{10} + (V_{11} + V_{10})d^1d^2))$$

4. The fourth strategy is given by

$$p(H_1|\tilde{s}) = \begin{cases} 1, & \tilde{s} = (1, 1) \\ x_4, & \tilde{s} \in \{(0, 0), (0, 1), (1, 0)\} \end{cases} \quad (6)$$

Here, we label this case as the optimistic case. The operator always assumes that a target is present when both sensors report; all other \tilde{s} are treated equally and left up to the operator's discretion to decide whether there is a target present or not. Combining (1) and (2), we have that the expected value is:

$$E(x_4) = x_4(p(h_0)(V_{00} + V_{01})(f^1f^2 - 1) + p(h_1)(V_{11} + V_{10})(1 - d^1d^2)) + p(h_0)(V_{00}(1 - f^1f^2) - V_{01}f^1f^2) + p(h_1)(V_{10}(d^1d^2 - 1) + V_{11}d^1d^2)$$

5. The fifth strategy is given by

$$p(H_1|\tilde{s}) = \begin{cases} 1, & \tilde{s} \in \{(1, 1), (0, 1), (1, 0)\} \\ x_5, & \tilde{s} = (0, 0) \end{cases} \quad (7)$$

Here, we label this case as the very optimistic case. Apart from $\tilde{s} = (0, 0)$, the operator in all other instances determines that there is a target present.

When both sensors do not report, then it is left up to the operator to decide whether they could have both missed a target or not. Combining (11) and (12), we have that the expected value is:

$$\begin{aligned} E(x_5) = & x_5 (p(h_0)(V_{00} + V_{01})(f^1 - 1)(1 - f^2) \\ & + p(h_1)(V_{11} + V_{10})(1 - d^1)(1 - d^2)) \\ & + p(h_0)((V_{00} + V_{01})(1 - f^1)(1 - f^2) - V_{01}) \\ & + p(h_1)((V_{11} + V_{10})(d^1 - 1)(1 - d^2) + V_{11}) \end{aligned}$$

We note of course that for all decision strategies $E(x)$ models poor performances as x varies. Of course, we remark that it is obvious that when all the values, target, detection and false alarm probabilities are fixed then the value of x that maximises $E(x)$ will be either 0 or 1. However, we recall that from the beginning the operator was given some discretion to use his intuition within some constraints to make some decisions. If we next take the worst case scenario and assume that the operator has performed poorly, we then wish to mitigate his poor decision-making. And because we only know his choice of x will be within a given range and with a certain distribution, we have to consider all the possible variable values (or at least a large percentage of them) of x and all the corresponding $E(x)$. Lastly, we note that we assumed that when an operator effectively uses his intuition, then for his choice of x and a given decision strategy, the expected decision value of his decision-making process will be greater than $E(x)$.

3 Analysis and Selection

Throughout the rest of the paper, we have assumed the distribution of x to be uniform. However, the focus of this paper is to develop a robust intelligent decision support system, with the techniques developed here applicable to any distribution of x . Now, for the uniform distribution, we can obtain a closed form solution for the mean and standard deviation using simple linear regression theory [10]. That is, the *centre of mass* (average value) is simply given by $\bar{E} = E(x = 0.5)$ and the standard deviation of $E(x)$ is given by ms , where m is the slope of the line and s is the standard deviation of the uniform distribution (note that the correlation coefficient in this case is 1). Thus, in this case, the variance of the expected values is given by $m^2/12$. Thus, we can select a strategy by making use of Chebyshev's inequality [10] which we list next.

Theorem 1. *If the random variable X has mean μ and variance σ^2 , then for every $k \geq 1$,*

$$P(|X - \mu| \geq k\sigma) \leq \frac{1}{k^2}$$

This theorem, which applies to all distributions, in effect states that at least $(1 - 1/k^2) \times 100\%$ of all the variable values lie within k standard deviations of the mean. A weaker implication of this is that at least $(1 - 1/k^2) \times 100\%$ of the variable values are greater than $\mu - \sigma k$. In our case, for the uniform distribution, this is given approximately by $E(x = 0.5) - 0.29mk$. For other distributions, if a closed form solution is not available, then a Montecarlo simulation could be used to determine a set of x values, from which the expected values of each decision strategy can be calculated along with their means and variances.

This methodology ensures that the expected value and variance of the chosen decision strategy is not only analytically tractable, but is also best overall when an operator's performance deteriorates. Note that in making his decision as to whether there is a target or not, the operator is not directly using any information from other sources, and at most only drawing on his previous experience of similar situations or "gut feeling". We end this section by briefly outlining the algorithm.

3.1 Algorithm

We now propose to outline a possible way this technique might be employed.

Step 1. Obtain sensors' performance parameters, target probability and reward and cost values of object(s) of interest

Step 2. From a range of available decision strategies, plot the graphs of expected value formulae.

Step 3. Select the strategy which is optimal by using Chebyshev's inequality. That is, suppose we wish to determine the largest lower bound for at least $q\%$ of the values in the decision strategy. Then $k = \sqrt{100/(100 - q)}$ and so for each decision strategy, this value, which we call the decision figure, is given by $E(0.5) - 0.29m\sqrt{100/(100 - q)}$

An obvious application of the decision strategies listed here would be in offline processing of imagery with reports from sensors. That is, suppose we have an image with some sensor reports, indicating they are either from 1 or both sensors simultaneously. The operator (using for instance strategy VP) is then asked to make judgement on whether there is an actual target there or not, but only when both sensors report. Clearly, we do not know what value of x he will choose with the only information regarding x is that it is uniformly distributed. If we take a risk mitigation approach, then we can model the operator's performance in terms of $E(x)$; a model of poor operator performance as defined here. We extend this process to all the other decision strategies and then apply the methodology described with the aim of choosing the most robust decision strategy when performance deteriorates. That is, we wish to find the decision strategy with the largest lower bound as determined by (II) that corresponds to at least $q\%$ of the

possible x values. Such a strategy is then deemed to be optimal for the purposes of this paper.

4 Examples

We now present some examples that illustrate how the reward and cost values can influence which decision strategy to choose, given some fixed sensor parameters and target probabilities. For the purposes of this example, we have assumed that x is uniformly distributed, however, the same methodology can be used for any distribution. For instance, for a normal distribution, we simply perform some Montecarlo sampling to get some values of x and the corresponding values of $E(x)$. Then, we calculate the mean and variance of the expected value samples, thus allowing Chebyshev’s inequality to be employed. In all cases we have let $f^1 = f^2 = 0.35$, $d^1 = d^2 = 0.75$ and the probability of a target is 0.175. We present 4 examples with different reward and cost values. For Table 1, we let VP, P, U, O, VO represent the very pessimistic, pessimistic, unbiased, optimistic and very optimistic decision strategies as defined above. Note that the reward and cost values labelled neutral (N), airport (A) and gain (G) were obtained from [3]. In N , both types of errors were given the same value as were both types of correct decision. The second scheme was originally called “airport”. Here, this could be the reward scheme of a bomb hidden in a suitcase. The cost of missing such bomb would be prohibitively expensive compared to false alarms that resulted in a relatively quick manual inspection. The last scheme, denoted by D for deterrent, is indicative of what would happen if a security guard was employed at the exit of a shop to inspect bags. Here, the cost of employing such security guard would be perhaps much higher than any potential gains from catching someone who is trying to steal an item. However, anyone who is falsely stopped and accused of stealing an item could be upset enough to never shop in that store again thereby losing a customer and potential business.

For instance, for example N , if we want to select the strategy which gives the largest decision figure for at least 75% of the variable values ($k = 2$), then, through Chebyshev’s inequality, the pessimistic strategy should be chosen. But, if instead we want to select the strategy which gives the largest decision figure

Table 1. Means (E) and Variances (V) of examples with the same sensor parameters but varying reward and cost values

Ex.	Rewards/Costs Values				Decision Strategy									
					VP		P		U		O		VO	
	V_{00}	V_{01}	V_{10}	V_{11}	E	V	E	V	E	V	E	V	E	V
N	1	50	50	1	-7.99	0.0015	1.61	21.2	1.54	21.0	-24.6	90.8	-32.5	24.7
A	1	50	900	100	-110.0	725.1	228.2	1630	274.8	180.0	-43.6	131.0	-20.3	3.90
G	1	50	50	950	38.7	725.1	79.5	1630	126.1	180.0	105.2	131.0	128.4	3.90
D	900	225	50	50	681.8	898.8	491.5	22500	439.5	1440	226.5	54200	18.6	12700




Fig. 1. Example with $f^1 = f^2 = 0.35$, $d^1 = d^2 = 0.75$ and target probability of 0.175 with varying reward/cost values

for at least 97% of the variable values ($k = 6$), then the neutral strategy should be selected.

5 Concluding Remarks

In this paper, we have derived decision strategies from multisensor outputs which combine an operator's intuition for some sensors' outputs with deterministic decisions in other scenarios. The strategies listed here are by no means the only ones and combinations using 2 or more of the options listed above are possible. More importantly, we have outlined a method for selecting the most robust decision strategy (which combine determinism and intuition) when intuition fails and the operator's decision-making process thus deteriorates.

For more than 2 sensors, these schemes can be easily adapted. However, the number of possible sensors' outputs grows from 4 to 2^N if N sensors are employed. Through the use of examples, we have shown how different reward and cost value schemes can influence the choice for the most desirable strategy. If the variable which represents the operator's intuition is uniformly distributed then this is simply calculated by discretizing the probability range and then selecting the most robust strategy as explained above. But for other distributions of x , a Montecarlo approach might be required to generate the distribution of expected values of the decision strategies. The most robust decision strategy can then be chosen after direct calculation of the appropriate mean and variance measures,

and employing Chebyshev's inequality as given above. Of course, this inequality does not provide the tightest bounds for all distributions, but its advantage lies in its universality and ease of implementation. Of course, we note that although our examples were theoretical, they were, however, realistic in terms of the likely reward and cost values for the scenarios mentioned. Other tests such as ANOVA could also be employed to see if two or more distributions, corresponding to different decision strategies, significantly differ from one another. Lastly, the types of decision strategies and algorithm listed here could also be implemented in other critical environments such as hospital and business theatres. Potential future work could lie in examining other types of distributions.

References

1. Klein, G.: *Intuition at Work*. Doubleday (2003)
2. Green, D.G., Swets, J.A.: *Signal Detection Theory and Psychophysics*. Peninsula Publishing (1988)
3. Navalpakkam, V., Koch, C., Perona, P.: *Homo Economicus in Visual Research*. *J. Vision* 9(1), 36, 1–13 (2009)
4. Azuma, R., Daily, M., Furmanski, C.: *A Review of Time Critical Decision Making Models and Human Cognitive Processes*. In: *Proc. 2006 IEEE Aero. Conf.* (2006)
5. Horvitz, E., Barry, M.: *Display of Information for Time-Critical Decision Making*. In: *Proc. 11th Conf. on Uncertainty in Artificial Intelligence*, pp. 296–305 (1995)
6. Keeney, R.L., Raiffa, H.: *Decisions with Multiple Objectives*. John Wiley & Sons, Inc., Chichester (1976)
7. Anderson, J.: *Cognitive Architectures in a Rational Analysis*. In: VanLehn, K. (ed.) *Architectures for Intelligence*, pp. 1–24. Lawrence Erlbaum Associates, Mahwah (1991)
8. Straker, D.: *Changing Minds*. Syque Press (2003)
9. Friedman, S.: *Learning to Make More Effective Decisions: Changing Beliefs as a Prelude to Action*. *Qualitative Market Research: The Learning Organisation* 11(2), 110–128 (2004)
10. Christensen, R.: *Analysis of Variance, Design and Regression: Applied Statistical Methods*. Chapman and Hall, Boca Raton (1996)

A UML Profile Oriented to the Requirements Modeling in Intelligent Tutoring Systems Projects

Gilleanes Thorwald Araujo Guedes and Rosa Maria Vicari

Instituto de Informática
Programa de Pós-Graduação em Computação (PPGC)
Universidade Federal do Rio Grande do Sul (UFRGS) - Porto Alegre - RS - Brasil
gtaguedes@inf.ufrgs.br, rosa@inf.ufrgs.br
<http://www.inf.ufrgs.br/pos>

Abstract. This paper describes a proposal for the creation of a UML profile oriented to the intelligent tutoring systems project. In this paper we shall describe the proposed profile as well as its application into the modeling of the AMEA intelligent tutoring system.

Keywords: UML Profiles, Stereotypes, Actors, Use-Cases, Agents, AMEA.

1 Introduction

In the area of Artificial Intelligence, the employment of intelligent agents as auxiliary aids to software applied to the most diverse dominions is being spread out. This practice has shown to be a good alternative for the development of complex systems, fostering a great increase of agent-supported software development in the several areas, one of those the intelligent tutoring systems, in which agents are also employed as pedagogical agents.

However, the development of this kind of system has presented new challenges to the software engineering area and this led to the surfacing of a new sub-area, blending together concepts brought over from both the software engineering and artificial intelligence areas, which is known as the AOSE - Agent-Oriented Software Engineering, whose goal is that of proposing methods and languages for projecting and modeling agent-supported software.

Despite many AOSE methods having been created for multi-agent systems (MAS) projects, Vicari [1] states these are not totally adequate for the modeling of Intelligent Tutoring Systems. Likewise, some attempts to adapt and extend UML (Unified Modeling Language) to the multi-agent systems projects were made; nonetheless, those attempts we have studied did not concern themselves into extending and applying some of the UML resources, such as the use-case diagram, which is mainly employed for requirements collecting and analyzing, an essential phase for the achievement of a good system project. Thus, we developed a UML profile for the project of Intelligent Tutoring Systems, for which end we started by adapting the use-case diagram.

2 UML, Metamodels and Profiles

According to [2], the UML is a visual language for specifying, constructing, and documenting the artifacts of systems. It is a general-purpose modeling language that can be applied to all application domains.

The UML specification is defined using a meta-modeling approach (i.e., a meta-model is used to specify the model that comprises UML) that adapts formal specification techniques. When meta-modeling, we initially establish a distinction between meta-models and models. A model typically contains model elements. These are created by instantiating model elements from a meta-model, i.e., meta-model elements. The typical role of a meta-model is that of defining the semantics for the way model elements within a model get instantiated.

A Profile is a kind of Package that extends a reference meta-model. The primary extension construct is the Stereotype, which is defined as part of the Profiles. A profile introduces several constraints, or restrictions, on ordinary meta-modeling through the use of the meta-classes defined in this package.

3 UML-Derived Languages

Some attempts have already been tried to adapt UML for the project of multi-agent systems, though nothing specific for the project of intelligent tutor systems. One of the first attempts was the AUML language [3]. Besides that, other languages, like AML [4], AORML [5], and MAS-ML [6] were also proposed.

However, neither of the above focus the matter of requirements collecting and analyzing nor on its modeling by means of the UML use cases diagram and no attempt was found in those languages to extend the metaclasses employed in that diagram for applying them on the multi-agent systems project.

4 UML Profile for the ITS Project

Considering that UML is a standard modeling language broadly accepted and understood in the software engineering area, that multi-agent systems own their proper characteristics, and that very few of works applying UML into a multi-agent systems project did care to focus the matter of requirements collecting and analyzing, we decided on creating a UML profile in which we extend the metaclasses employed by the use-case diagram, thus creating stereotypes prepared to identify the particular functionalities of this kind of system, as can be seen in the following illustration.

To develop this profile, we began by using the original metaclass, Actor. According to [7], the Actor metaclass represents a type of role played by an entity that interacts with the subject, but which is external to the subject. An actor does not necessarily represent a specific physical entity but merely a particular facet (i.e., "role") of some entity that is relevant to the specification of its associated use cases.

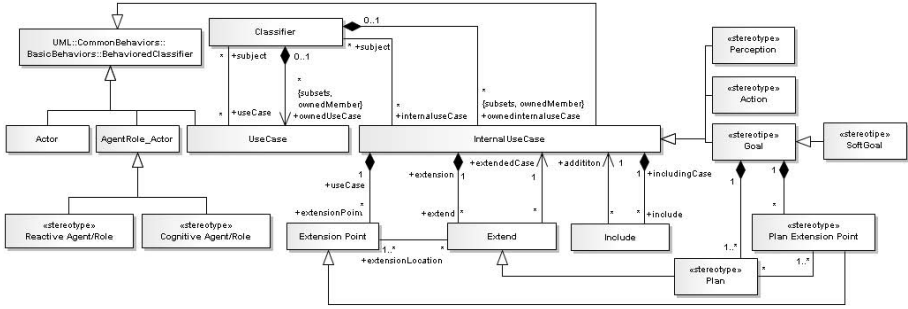


Fig. 1. UML Profile for the Project of Intelligent Tutoring Systems

However, most of the times, the software agents are not external to the software, rather they customarily are inserted in the system environment and, as they are independent, proactive, and able to interact with the software according to their goals, these should be represented as actors. So, as states [8], it is necessary to adapt this concept, considering that agents can be internal to the system, that is, an agent can be a part of the subject and, therefore, if we are to represent agents (their roles) as actors, said actors should be internally represented within the system’s borders, for they belong to the software.

The representation of agents/roles as UML actors can also be seen in [9], in which agents represent active objects and are modeled within the system as rectangle-headed actors. However [8] only suggests that the concept of actor should be adapted and [9] did not created a UML profile for agents. In our work we explicitly derived new metaclasses from the original metaclasses endeavoring for the former to allow an adequate requirements modeling in MAS.

However, according to [2], it is not possible to take away any of the constraints that were applied to a metamodel such as UML when using a profile. Thus, instead of extending the Actor metaclass, we created a new metaclass derived from the same metaclass as the Actor metaclass had been, as can be seen in figure 1, creating the Agent/Role-Actor Metaclass. In this new metaclass we copied all the characteristics of the Actor metaclass, including the same symbol, suppressing only the constraint that an actor needs must be external to the system. And, from this new metaclass we derived the Reactive Agent/Role and Cognitive Agent/Role metaclasses.

As can be observed, we applied the stereotype ”stereotype” in both metaclasses, which means that these metaclasses will be applied as stereotypes on agent/role actors while attributing them special characteristics, for said agent/role actors shall be utilized within specific domains to represent cognitive and reactive agents/roles. We opted for creating two stereotypes so as to establish a difference between reactive and cognitive agents, considering that either group presents different characteristics.

According to [7], a use case is the specification of a set of actions performed by a system. Each use case specifies some behavior that the subject can perform

in collaboration with one or more actors. Use cases can be used both for specification of the (external) requirements on a subject and for the specification of the functionality offered by a subject. Moreover, the use cases also state the requirements the specified subject poses on its environment by defining how they should interact with the subject so that it will be able to perform its services.

Thus, with the objective of adapting the use cases concept to the MAS modeling, we derived a new metaclass departing from the same metaclass as the UseCase metaclass, called "InternalUseCase" and then we created the same relationships with the Classifier metaclass for this new metaclass as those belonging to the UseCase metaclass. We made it that way because we intended to represent goals, plans, actions and perceptions as use cases, but the semantics of the UseCase Metaclass says that use cases represent external requirements, therefore, to adapt this concept to the MAS modeling we created a similar metaclass, all the while we modified its semantics to represent internal requirements.

Naturally, all Internal use cases are internal to the system and can be neither used nor seen by human users. From the InternalUseCase metaclass, we extended some metaclasses to attribute special characteristics to internal use cases; these extended metaclasses will be employed as stereotypes. Thusly, we created the metaclasses, Perception and Action, to model internal use cases that contained the necessary steps for an agent to perceive or do something, a procedure also suggested by [8] to be applied in normal use cases, though as stated before in [7], these refer to the external requirements, not to the internal ones.

A third metaclass was derived from the metaclass, InternalUseCase, to represent Goals. A internal use case employing the stereotype, Goal, shall contain a description of a desire to be attained by an agent and the possible conditions for that desire to become an intention. Besides, we evolved the first proposal of this profile presented in [10] when we included the concept of softgoal, deriving the SoftGoal metaclass from the Goal metaclass. The concept of softgoal is described in [11] and it is used to represent cloudy goals, which do not present a clear definition and/or criteria for deciding whether they are satisfied or not.

A somewhat similar proposal for representing goals as use cases can be seen in [9], but, besides using internal use cases, we went a little further beyond the proposal presented by [9], when we considered that, just like a goal represents a desire that will not necessarily become an intention, the steps for its execution should be detailed in other or others internal use cases; in the situation of a internal use case employing the stereotype, Goal, we shall only detail those perceptions and conditions necessary for that goal to become an intention. Considering a goal might eventually have more than a plan and that this or these plans will only be accomplished within certain conditions, we decided on deriving still another metaclass, Plan, from the metaclass, Extend.

According to [7] a metaclass, Extend, represents a relationship from an extending use case to an extended use case that specifies how and when the behavior defined in the extending use case can be inserted into the behavior defined in the extended use case. If the condition of the extension is true at the time the first extension point is reached during the execution of the extended use case, then

all of the appropriate behavior fragments of the extending use case will also be executed. If the condition is false, the extension will not occur. Because a plan only will be triggered after some particular condition be satisfied, we decided on extending the metaclass, Plan, from the metaclass, Extend, and associated the former to the metaclass, Goal, by means of a composite association.

Finally, we derived the meta-class, Plan Extension Point from the meta-class, Extension Point. According to [7] an extension point identifies a point in the behavior of a use case where that behavior can be extended by the behavior of some other (extending) use case, as specified by an extend relationship. We derived this last meta-class only to set up a difference between Plan Extension Points and normal Extension Points; nonetheless, this also can serve to render explicit the condition for a plan to be executed.

5 The AMEA Project

The AME-A [12] architecture - Ambiente Multi-Agente de Ensino-Aprendizagem or “Teaching-Learning Multi-agent Environment” is composed by a hybrid society of agents that cooperate to aid into the students learning process. The environment interacts with human agents that can be both the teacher or the students and owns several reactive and cognitive agents.

The teacher can create a learning activity or evaluate the students with the help of the Teacher’s Tools reactive agent. The student, on his turn, can choose between execute an unmonitored learning session or a monitored learning session. In the first option, the student only interacts with the Unsupervised Learning reactive agent that will only present him/her the contents to be learned.

The monitored learning activity is set as the main focus for the system, in which it aims to maximize the student learning by means of the aid of five cognitive agents, to wit: Student Modeling (SM), Methodology and Teaching Plan (MTP), Learning Orientation (LO), Learning Analysis (LA) and Knowledge Application Orienting (KAO). The first models the student profile in a dynamic way, while the second chooses the methodology and learning plan that are more adequated to the student profile every time it changes or whenever the student performance is lower than the expected level; the LO agent selects the contents to be taught and the way how these will be presented according to the chosen methodology; the LA agent checks on the student performance throughout the session and the KAO agent applies a general evaluation after the session ends.

Since the teacher and the student are human agents that interacts with the system but are external to it, we represent them as normal actors, out of the system boundaries and we associate to the actor that represents the teacher to the functionalities represented by normal use cases, “Create learning activity” and “Evaluate students”. Notice that in these functionalities there is also an interaction with the Teacher Tools agent, who, being a reactive agent, was represented as an agent/role actor with the stereotype “Reactive Agent/Role” and placed inside the system boundaries, because it is inserted in the software.

Now the actor student was associated to the functionalities “Execute an unmonitored learning session” and “Execute a monitored learning session”, equally represented as normal use cases. In the first functionality there is also a interaction with the reactive agent Unsupervised Learning, which, in the same way as the previous agent was represented as an agent/role actor with the estereotype “Reactive Agent/Role”.

The functionality “Execute monitored learning session” represents the more important activity of the system and it is the one that presents more complexity, involving the five cognitive agents of the software, which, because are agents, are represented as Agent/Role Actors inside of the system, containing the “Cognitive Agent/Role” stereotype, since they are cognitive agents.

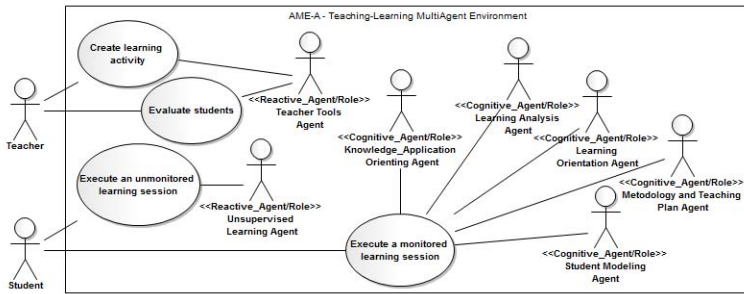


Fig. 2. Normal Actors modeling external users, normal Use Cases modeling functionalities, and stereotyped Agent/Role Actors modeling agents

These functionalities were represented as normal use cases, since they represent services that the external users can ask from the system. Figure 2 illustrates this excerpt of the modeling. This figure presents only a part of the use cases diagram. Due to the lack of space and for a matter of visibility, we suppressed the representation of the cognitive agents goals, plans, perceptions and, actions, detailing separately the requirements for each cognitive agent. To sustain the connection with this excerpt of the modeling, we shall maintain the association of the cognitive agents with the use case “Execute monitored learning activity”.

The SM agent has for its goal modeling the student in a dynamic way. This agent has two perceptions. First it must perceive that the learning session beginning and, in this case, trigger the “apply questionnaire” plan to determine the student profile. And it needs perceive when the student behavior changes, in which case it has to trigger the plan to remodel the student profile. Figure 3 illustrates the modeling of these requirements for the SM agent.

In this figure we associated to the SM agent an internal use case (IUC) with the Goal stereotype representing the goal that this agent has to model the student profile. This goal has two inclusion associations with two IUCs that represent the perceptions the agent needs to own to determine whether is necessary to trigger some of the plans associated to the goal. An inclusion association determines that

the behavior of the use case therein included should mandatorily be executed by the use case it includes. So, to reach its goal, the SM agent has to execute these perceptions. So that, we use an IUC with the stereotype Perception to represent the perception of learning session beginning and another IUC with the same stereotype to represent the perception of the student behavior.




Fig. 3. Internal use cases with Goal, Plan, and Perception stereotypes

Notice that, in the IUC that represents the goal, there are two Plan extension points, both representing the points in the goal behavior to where the plans associated to it can be extended, besides establishing the conditions for the plans to be executed. In such way, if the agent perceives the learning session is beginning, the plan to apply a questionnaire to the student will be triggered, represented by an IUC containing the Plan stereotype; and, if the agent notices a change in the student behavior, the plan to remodel the student will be triggered, equally represented by an IUC containing the Plan stereotype. Further observe that the two IUCs that represent the plans are associated to the IUC Goal by means of extension associations, that is, these IUCs will only be executed after the conditions detailed by the Plan extension points are satisfied.”

The following agent, MTP, has for its goal to choose the methodology and teaching plan which is more adequate to the student, to do so, it has to perceive when the student model changes, this includes the perception of when the student is modeled for the first time, when there is no model yet. This goal has an associated plan, “Change Learning Methodology and Teaching Plan” which will be triggered when the student model changes or when the student’s performance is low. The execution of this plan includes the sending of a message to the LO agent informing that the methodology was modified.

In figure 4, to model these requirements we associated an IUC including the Goal stereotype with the agent to represent its goal. After that, we associated, by means of inclusions, two IUCs with the Perception stereotype to represent the perceptions of student model change and student performance. After that, we created an extension association to connect the IUC with the Plan stereotype, which will represent the plan for methodology and teaching plan changes to the Goal IUC. This plan will be triggered only when the student model changes or

when the student performance is low, as demonstrated by the Plan extension points in the Goal IUC. Finally, if the plan is executed, it is need to communicate to the LO agent the chage of methodology; since this is done by means of a communication between agents, we identified this as an action and we associated it to the plan by means of an IUC with the Action stereotype.

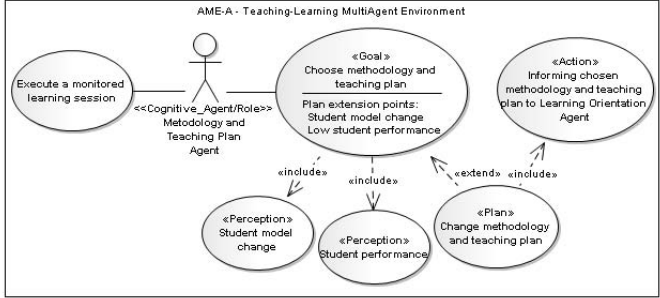


Fig. 4. Methodology and Teaching Plan agent requirements

The LO agent has for its goal to present learning contents for the student and for its perception the choice of the methodology and teaching plan. When the methodology and the teaching plan are perceived by the LO agent, it executes the plan “Select material for learning”. These requirements are modeled on figure 5, where we represent the goal, the perception and the plan as internal use cases containing respectively the stereotypes Goal, Perception and Plan. As the plan will only be triggered when the choice of a methodology is perceived, there is an inclusion association between the goal and the perception, obligating the agent to verify if it occurs. There is also an extension association between the goal and the plan, since the plan only will be triggered if the condition is satisfied.




Fig. 5. Learning Orientation agent requirements

Figure 6 presents the modeling related to the requirements of LA agent, which has for its goal to check the knowledge acquired by the student, represented by an IUC containing the Goal stereotype. To execute this goal, the agent must perceiving the student’s performance; this perception is represented by an IUC

containing the Perception stereotype. Besides, the LA agent must inform this performance to the MTP agent, which is represented as an IUC containing the Action stereotype. If the student performance is considered low, then the plan “Boosting student” is triggered, equally represented as an IUC with the Plan stereotype. This plan has for action to send motivation messages to the student; we identified this as an IUC containing the Action stereotype and we connected this to plan by means of an inclusion association.




Fig. 6. Learning Analysis agent requirements

The last agent involved in the process, the KAO agent, has for its goal to evaluate the student after the session ends. Thus, it needs to perceive the learning session ends so as to know when to trigger the plan “Apply evaluation”. Here we applied the same stereotypes used in the previous examples in internal use cases, as demonstrated in figure 7.

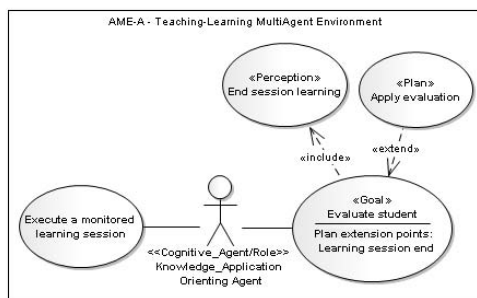


Fig. 7. Knowledge Application Orienting agent requirements

There are structural and behavioral questions that need to be best detailed, since the presented profile focus only the system requirements. All the same, even on project level, it is needed to represent information like beliefs, desires and intentions by means of class diagrams and the processes represented by the IUCs must be detailed through behavioral diagrams. Up to now we have been

directly using these diagrams in [10], though we might instead extend other UML metaclasses so as to best adequate them to the ITS Project or we can apply some of the UML derived languages to the MAS project previously described.

6 Conclusions

Throughout this paper we have presented a UML profile developed to the intelligent tutoring systems project oriented to the requirements collection and analysis. We demonstrated the applicability of said profile by means of AMEA System modeling and also that the stereotypes we have created can be used to model cognitive and reactive agents, and actions, perceptions, goals, and plans as well. However, we intend to apply this profile on some other projects for Intelligent Tutoring Systems, possibly of a more complex sort, so as to find out whether this profile is really adequate for projecting this type of systems and to pinpoint any weakness needing some improvement.

Although this profile has been developed for the Intelligent Tutoring Systems modelling, we believe it might be applied to other MultiAgent System projects oriented to other dominions. We also believe this profile can be adapted to most UML already existing extensions created for the MultiAgent Systems project.

References

1. Vicari, R.M., Gluz, J.C.: An Intelligent Tutoring System (ITS) View on AOSE. *International Journal of Agent-Oriented Software Engineering* (2007)
2. OMG - Object Management Group. Unified Modeling Language: Infrastructure Specification - Version 2.1.1. OMG (2007), <http://www.omg.org>
3. AUML Official Website (2007), <http://www.auml.org>
4. Trencansky, I., Cervenka, R.: Agent Modeling Language (AML): A comprehensive approach to modeling MAS. *Informatica* 29(4), 391–400 (2005)
5. Wagner, G.: A UML Profile for External AOR Models. In: *Third International Workshop on Agent-Oriented Software Engineering* (2002)
6. Silva, V., et al.: MAS-ML: A Multi-Agent System Modeling Language. *International Journal of Agent-Oriented Software Engineering, Special Issue on Modeling Languages for Agent Systems*, Inderscience Publishers 2(4) (2008)
7. OMG - Object Management Group. Unified Modeling Language: Superstructure Specification - Version 2.1.1. OMG (2007), <http://www.omg.org>
8. Bauer, B., Odell, J.: UML 2.0 and Agents: How to Build Agent-based Systems with the new UML Standard. *Journal of Engineering Applications of AI* (2005)
9. Flake, S., Geiger, C., Küster, J.: Towards UML-based Analysis and Design of Multi-Agent Systems. In: *Proceedings of ENAIS 2001, Dubai (March 2001)*
10. Guedes, G.T.A., Vicari, R.M.: A UML Profile Oriented to the Requirements Collecting and Analyzing for the MultiAgent Systems Project. In: *Proceedings of SEKE 2010, San Francisco Bay* (2010)
11. Bresciani, P., et al.: Tropos: An Agent-Oriented Software Development Methodology. *Autonomous Agents and Multi-Agent Systems* 8(3), 203–236 (2004)
12. D'Amico, C.B., Vicari, R.: Adapting Teaching Strategies in a Learning Environment on WWW. *WebNet* (1998)

Learning by Collaboration in Intelligent Autonomous Systems

Jorge Ierache, Ramón García-Martínez, and Armando De Giusti

Computer Science PhD Program. Computer Sc. School. La Plata National University
Institute of intelligent systems and robotics experimental teaching, FICCTE, UM
Intelligent Systems Laboratory. Engineering School. Buenos Aires University
Institute for Research on Information Technology- LIDI, School of Information Technology,
UNLP
jierache@unimoron.edu.ar

Abstract. Very few learning systems applied to problem solving have focused on learning operator definitions from the interaction with a completely unknown environment. Autonomous Intelligent Systems (AIS) deal with that issue by means of architectures where learning is achieved by establishing plans, executing those plans in the environment, analyzing the results of the execution, and combining new evidence with prior evidence. This paper proposes a selective mechanism of learning allowing an AIS to learn new operators by receiving them from another AIS in a higher stage in the Learning Life Cycle (LLC) with more cycles of interaction in the environment. The proposed collaboration mechanism also considers how to deal with theory ponderation (operators ponderation) and how to include the new operators (provided for) in the set of theories of the receiver AIS. The experimental results show how using collaboration-based learning among AIS provides a better percentage of successful plans, plus an improved convergence rate, than the individual AIS alone.

1 Introduction

We are presenting in this work the results achieved from the application of collaboration from the autonomous intelligent systems (AIS) of robots that are in a higher stage in the learning life cycle of an AIS [1]. The autonomous intelligent systems (AIS) evolve from initial theories (set of operators built in by the AIS's programmer) to ones learned from interaction with the environment or other AISs. Given unknown environments, real autonomous systems must generate theories of how their environment reacts to their actions, and how the actions affect the environment. Usually these learned theories are partial, incomplete and incorrect but can be used to plan, to further modify those theories, or to create new ones. The previous work on machine learning applied to problem solving has mainly focused on learning knowledge, whose goal was to improve the efficiency of the problem-solving task [2]; [3]. There is also a current interest in learning state transition probabilities in the context of reinforcement learning. [4] However, few researchers have approached the generalized operator acquisition problem [5], [6], described as techniques for automatically

acquiring generalized descriptions of a domain theory. This issue is crucial when dealing with systems that must autonomously adapt to an unknown and dynamic environment. LOPE (Learning by Observation in Planning Environments) is an AIS implemented architecture that integrates planning, learning, and execution in a closed loop, showing an autonomous intelligent behavior [7]. Learning planning operators (what we will call *operators* is also referred to as action models within the reinforcement learning community) are achieved by observing the consequences of executing planned actions in the environment. In order to speed up the convergence, heuristic generalizations of the observations have been used. Moreover, probability distribution estimators have been introduced to handle the contradictions among the generated planning operators [8] and how by sharing among AISs the learned operators improve their behavior [9]. In this context, a learned operator O (theory) in LOPE [8] is a tuple $\langle C, A, F, P, K, U \rangle$ where: C is the initial situation (conditions), A is the action to be performed, F is the final situation (post-conditions), P means the times that the operator O_i was successfully applied (the expected final situation F was obtained), K means the times that the action A was applied to C , U means the utility level reached by applying the action to the initial situation C of the operator.

As the natural next step this paper aims to propose the general description of the collaboration model within the framework of the Learning Life Cycle (section 2), to describe the experimentation design and to analyze the results attained (section 3) and to draw some conclusions and future research lines (section 4).

2 General Description of the Collaboration Mechanism

In this section it is presented the description of the collaboration method between an AIS that has reached a higher evolutionary stage and a receiver AIS that is in a lower stage of evolution. The mechanism proposed considers theory ponderation and the inclusion of new theories in the knowledge base of the receiver AIS, provided by the collaborator AIS. Figure 1 shows a diagram of the model oriented to collaboration, on the basis of the LLC (Learning Life Cycle) proposed by the authors [1]. In the framework of the LLC of an AIS, three learning layers are presented: [a] Layer BIO (Built-In Operators) is the layer where the operators are implanted into the “Born” AIS by their creator, [b] Layer TBO (Trained Base Operations) is the one where the operators are learned by the “Newbie” AIS, in the simulation scenario (previously designed) and by evolutionary learning techniques, [c] Layer WIO (World Interaction Operators) is the one where the operators are learned by interaction of the “Trained” AIS with the part of the world representing their operation environment and with the presence of other AISs to turn into a “Mature” AIS.

The objective of the AIS is to autonomously learn operators (action models) that may predict the effects of the actions on the environment, through the observation of the consequences of these actions and to advance in their evolution state (Born, Newbie, Trained, Mature).

The collaboration allows that an AIS may receive knowledge that has been previously acquired by an AIS that is in a higher evolutionary stage within the framework




Fig. 1. Collaboration model proposed

of the LLC (Learning Life Cycle). Figure 1 shows the possible levels of collaboration: [a] a Mature AIS that ended layer WIO with a Trained AIS that is at the beginning of layer TBO, or with a Born AIS at the end of layer BIO; [b] a Trained AIS that ended layer TBO with a Newbie AIS that ended layer BIO or a Born AIS starting layer BIO; [c] a Newbie AIS that ended layer BIO with a Born AIS that is starting layer BIO.

When an AIS collaborates, it shares the experience gained, which is the product of both right and wrong performance. Therefore, by following the theory model and its ponderation method, when collaborating with operators it is necessary to reflect this knowledge in the receiver AIS, by maintaining the consistency of its set of theories. A theory (operator) is represented by T: Initial Situation (IS), Action (A), Final Situation (FS), P,K,U. [8]; where P means the quantity of times a theory T_1 was successfully used (the predicted effects were attained) and K means the quantity of times the theory T_1 was used. Finally, U means the level of utility reached when applying the action A to the initial situation (IS) of the theory T_1 . The method employed provides for the concept that the theory successfully used must be reinforced by increasing its P and its K and that the theory unsuccessfully used must be weakened by increasing its K, without increasing its P [8].

In the collaboration process, if a theory exists in the receiver AIS, it is reinforced by adding up its P and K; on the other hand, if a theory of the collaborator AIS does not exist in the receiver AIS, but a similar theory does exist in the latter, this theory weakens, maintaining the same P and adding up the K of the collaborator AIS to its K. Furthermore, the theories of the collaborator AIS that are not in the receiver AIS or are not similar or equal do not show any change, and are only transferred to the receiver AIS with their P and K values. To maintain the soundness of the method of theory ponderation, it is verified that after the collaboration the addition of the Ps of the similar theories in the receiver AIS is equal to the K of any of them.

The collaboration method algorithm is detailed in Pseudo-code 1. The algorithm works in the following way: for each Theory of the Receiver AIS (RA), (1) it tries to find if there exists a similar theory in the set of Theories of the Collaborator AIS (CA); (1.1.) If this is the case, the Ks of similar theories are added up, (1.1.1). Then, if besides the existence of a similar theory there exists an equal theory in the Set of Theories of the Collaborator AIS (SetCA), the Ps of the equal theories are added up. (1.1.2). After repeating this process for each theory existing in the set of Theories of

the Receiver AIS (SetRA), the theories of the set of the collaborator AIS (SetAC) that have not equal theories in the set SetRA are processed.

(2). If a theory of the SetCA has a similar theory in SetRA, (2.1), it is registered in this one with the P of the theory of SetCA and the K of the similar theory of SetRA. If it does not exist a similar theory, it is registered in SetRA, without modifying the P and the K.

Pseudo-code .1 “Collaboration Algorithm”

```

INPUTS: Set Theories AIS receiver (SetAR), Set Theories collaboration
AIS (SetAC).
OUTPUT: Set Theories AIS receiver SetAR
1. FOR EACH THEORY Ti in SetAR
  BEGIN
    1.1 IF exists TJ which belongs to SetAC Similar to TI
    of SetAR THEN
      1.1.1 KTI=KTI+KTJ
      1.1.2 IF exists TJ which belongs to
        SetAC equal to TI of SetAR THEN
          PTI=PTI+PTJ
    END.
  2. FOR EACH THEORY TJ in SetAC
  IF not exists TJ which belongs to SetAC equal to TI which belongs
  to SetAR THEN.
  BEGIN
    2.1 IF exists TI which belongs to SetAR Similar to TJ THEN
  Register TJ with P=PTJ and KTI in SetAR
  ELSE
  Register TJ with P=PTJ and K =KTJ in SetAR
  END.

```

3 Experimentation and Result Analysis

In this section we are presenting the design of the experiments, the results attained and their analysis. Several experiments have been carried out on the basis of the e-puck robot, by using the Cyberbotics' Webots 6 mobile robot simulator developed in the Laboratory of Micro-Informatics of the Swiss Federal Institute of Technology, Lausanne, Switzerland (EPFL). Figure 2 shows the experimentation scenario and the AIS (e-puck robots). The labyrinth-shaped setting is composed of four joined walls that surround the external limits of the setting and internally of a three-cornered labyrinth. In turn, at the end of the labyrinth it appears the type of block obstacle (besides the wall-type).

Two types of utility functions have been applied. The first one was based on the situation perceived by the robot from the information on the distance to the obstacles given by its sensors; considering that the range of each sensor is [0 to 1024], the theory applied by the robot has more “U” utility, the farther the wall or the obstacle is; the nearer to zero (0) the values of its sensors are, the farther from the obstacle the robot will be. Finally, it was transformed the final utility of the theory into the value range [-1 a +1], the higher the value obtained by adding up the values of each sensor,

the nearer to -1 the utility will be. So as to count the successful theories, it was adopted as a criterion the “U” utility value range: $0.75 < U \leq 1$.

The second utility function was based on the action carried out by the robot (speeds applied to each wheel of the robot); the utility was determined on the basis of the speeds calculated by applying Braitenberg quotients [10], compared to the speeds indicated by the action of the theory applied by the AIS for its left and right wheels. The nearer the action carried out by the robot is with regard to the one calculated with the quotients, the higher the utility of the action taken by the AIS (nearer 1) will be; the farther the action is, the lower utility (near -1) it will be. To count the successful theories, it was adopted as a criterion the “U” utility value range: $0.75 < U \leq 1$.

The AIDs constituted by robot autonomous Systems (RASs), during the experimentation, have formulated a set of theories that allow them an “a priori” prediction of the effects of their actions. Upon this experimentation basis, the exchange (cooperation) was performed between Newbie AISs that are in the layer BIO and the collaboration from a Trained AIS that went around the layer TBO to a Newbie receiver AIS located in the LLC (Learning Life Cycle) layer BIO.

For the experimentation, it was applied the knowledge base of theories of three AISs that are represented by: the robot A (born -with 600 simulation cycles, which is starting layer BIO), the robot B (born- with 3000 simulation cycles, which is going around layer BIO), the robot B (newbie- with 6000 simulation cycles, which is finishing layer BIO). On the basis of said theories generated, it was implemented the exchange at the ‘born’ stage in layer BIO, between robot A (born) and robot B (born), the collaboration of robot B (newbie) that is starting layer TBO towards the receiver robot A (born) which is in layer BIO.

The stated theory bases were used by robot E (experimentation robot), on which the experiments with the application of the methods of plan ponderation, mutation and a combination of both were conducted. The neuter AIS was developed on the basis of a born robot that got 600 cycles (robot A). The robots A and B started with their creator (programmer), getting 400 cycles to generate their initial theory base through an initial reactive behavior. The theory bases of the different robots were stored in an XML format.




Fig. 2. Experimentation setting of the E-puck AIS

3.1 Experiment Design

The procedure applied in each experiment was performed on the basis of the simulation of the AIS robot for a number of 600 cycles (cycle = perception/action), considering the

starting position of the AIS and the selection of the environment at random. Every 20 cycles, it was generated a report on the status of the AIS variables, as shown in table 1 as an example, thus creating the row I of the matrix containing the variables of the AIS, which was generated in the cycle I*20. The general experimental procedure considers fifty experiments for each case, from which an average matrix resulting from the matrices corresponding to each experimental stage is constructed. The generation of graphs of the experiments presented is carried out on the basis of the information of each average matrix and the interpretation of the results achieved.

Table 1. Structure of the Report on the status of the AIS variables

Cycles	Situations	New Theories	Theories [0,75 <U < 1]	Quantity of Successful Plans	Quantity of Successful Cycles	% Successful Cycles
1	1	0	56	0	0	0
20	2	0	56	8	19	95
..
600	116	472,33	128,33	459,33	193,33	32

Independent Variables: The following independent variables are regarded: [a] *Time*, the experiments are developed on the basis of this independent variable; a time unit is considered as the period elapsed between one perception of the environment and another one. It is a quantitative variable whose value corresponds to the interval [0,600]; [b] *Threshold of Sensor reading*, the comparison of readings for each of the eight IR sensors in the proximity of the AIS (e-puck robot), is a set of quantitative variables [0,999] for each of the eight sensors of the AIS, from which the readings of the situations perceived throughout the time are compared. The determination of the comparison threshold of each sensor was carried out on the basis of the standard deviation of the readings of the situations registered by the robot in its initial born stage for fifty cycles, showing a reactive behavior; [c] *Threshold of Reading comparison of the effectors*, the speed values for each wheel of the AIS, is a set of quantitative variables [0,999] for each of the two wheels of the robot AIS, from which the readings of the actions (speeds applied to each wheel) are compared. The determination of the threshold for each wheel of the robot AIS is carried out on the basis of the standard deviation of the readings of the actions registered by the robot in its initial stage (born) for fifty cycles, showing a reactive behavior; [d] *Mutation* is a qualitative variable; it indicates if the theory mutation is active, generating new theories from similar theories, [8]; [e] *Plan ponderator* is a qualitative variable that indicates if the behavior of plan ponderation [8] is active; it is used for the selection of the plan to be executed by the AIS; [f] *Theory exchange* is a qualitative variable that indicates if theory exchange [9] is active and is used in the generation and assessment of theories between AISs as a product of exchange (cooperation); [g] *Collaboration* is a qualitative variable (proposed) that indicates if the collaboration between AISs is active, and is used in the generation and assessment of theories, from the theories provided by the action of collaboration of an AIS that is going around a higher LLC (Learning Life Cycle) layer towards a receiver AIS that is going around a lower LLC layer.

Dependent Variables: In the context of the experiments carried out, the following dependent variables are considered: [a] *Quantity of Situations* is a quantitative variable, whose value corresponds to the interval [0,200], and indicates the quantity of different situations that have been generated by the AIS according to the variable “time”; [b] *Quantity of successful theories* is a quantitative variable, whose value corresponds to the interval [0,600] and indicates the quantity of successful theories according to the variable “time”. Successful theories are considered to be those whose “U” utility exceeds the threshold of 0.75; [c] *Quantity of successful plans* is a quantitative variable, whose value corresponds to the interval [0,600]. This variable indicates the quantity of successful plans which have been generated by the system in the time interval: [0, time]. It is the quantity of plans that have reached their end after having gone around all the nodes of the plan chain; each node represents a situation (supposed conditions or predicted effects.) This value can be a rational number, since when a plan is not successful, its partial success is taken; [d] *Quantity of successful cycles* is a quantitative variable, whose value corresponds to the interval [0,600] and indicates the quantity of cycles that have proved to be successful according to the variable “time”; it is a value that is accumulated during the simulation and starts at zero (0) and increases by one (1) whenever the robot has carried out a cycle in which it ended with a U utility theory $> 0,75$ and $U < 1$; [e] *Percentage of successful cycles* is a quantitative variable, whose value corresponds to the interval [0,100] and indicates the percentage of cycles that have proved to be successful according to the variable “time”. This variable indicates the percentage of cycles that have ended with a new theory with U utility: $0,75 < U \leq 1$, at a given instant of the robot simulation which corresponds to the time interval [0, time]; [f] *Quantity of new theories* is a quantitative variable, whose value corresponds to the interval [0,600] and indicates the quantity of new theories according to the variable “time”. The theories considered are only those that were generated during the AIS operation and not those that were accumulated by the AIS in its theory base.

3.2 Graphs and Discussion of the Experimentation Results

To set up the architecture of each AIS that participates in every experiment, the following methods corresponding to independent variables were considered: mutation, plan ponderator, theory exchange between AISs and the proposal of the method of collaboration between AISs. Several experiments that come from the application of the different methods of learning acceleration and their combinations have been developed on the basis of the above-mentioned dependent and independent variables.

Figure 3 shows the comparison graph between situations and theories generated throughout the time, which has been obtained on the basis of the experimentation of an AIS set up with the methods of plan ponderation, mutation, and exchange and the collaboration one proposed. It is shown that even later, when it is observed a tendency to stabilization of the number of situations, the quantity of theories increases faster compared to the quantity of situations (supposed conditions or predicted effects). This shows that the system discovers relations between situations, being the AIS the theory generator.

Here follow the results of the experiments, compared to related works, by means of comparison graphs: the quantity of successful plans (fig.4), the quantity of successful cycles (fig.5), the percentage of successful cycles (fig.6), the quantity of successful

theories (fig.7) and the quantity of new theories (fig.8),under the following method configurations: AIS with plan ponderator and mutation methods, adding the collaboration method proposed (APMEC), AIS with plan ponderator, mutation and exchange methods (APME),proposed by Garcia Martinez et al [9], AIS with with plan ponderator and mutation (APM), proposed by García Martínez, R. y Borrajo, D., [8], Neuter AIS (AN) does not apply any method, proposed by Fritz, W et al [2].

Figure 4 shows the comparison graph of the quantity of successful plans; the AIS that combines the plan ponderator and mutation methods, adding the collaboration method proposed (APMEC), obtained the greater quantity of successful plans throughout the time, followed during the initial period by the AIS that applied plan ponderator, mutation (APM); after the initial period, the AIS that applied theory exchange experiences a slowing down in the quantity of successful plans generated and is surpassed by the AIS with Plan ponderator, mutation and exchange (APME). The lower quantity of successful plans throughout the time was generated by the neuter AIS (AN) that does not apply any method for learning acceleration.




Fig. 3. Quantity of Theories vs situations




Fig. 4. Quantity of successful plans

Figure 5 shows the comparison graph of “Quantity of successful cycles”, the AIS (APMEC) that applies collaboration obtained the greater quantity of successful cycles throughout the time. The same is observed in figure 6, the comparison graph of “Percentage of successful cycles”, the AIS (APMEC) that applies collaboration obtained the best average percentage of successful cycles throughout the time, with respect to the other methods applied by the AISs (APME, APM, AN).




Fig. 5. Quantity of successful cycles




Fig. 6. Percentage of successful cycles

These results contribute to the fact that the AIS that applies collaboration may generate a greater quantity of successful plans, on the basis that it has a greater quantity of accumulated successful theories (theories whose utility is $0.75 < U \leq 1$), being the product of the application of the collaboration method proposed, as it can be observed in figure 7: *Quantity of Successful Theories*.




Fig. 7. Quantity of Successful Theories




Fig. 8. Quantity of New Theories

Figure 8 shows the graph “Quantity of New Theories”; the greater quantity of new theories generated throughout the time corresponds to the neuter AIS (AN) and the smaller quantity to the AIS (APMEC) that applies collaboration. This is due to the fact that the knowledge base of theories of the neuter AIS is less representative regarding experience than that of the AIS that received collaboration; that is the reason why by having more experience in the operation environment (more quantity of successful theories), the AIS that received collaboration generated the smallest quantity of new theories compared to the rest of the methods that generated new theories during the experimentation.

Other series of experiments were carried out. They considered the individual comparison of each method applied on the basis of a neuter AIS, plan ponderator-mutation AIS, Exchange AIS, Collaboration AIS; the results were satisfactory for the AIS that received collaboration.

4 Preliminary Conclusions and Future Research Lines

According to the results achieved during the experimentation, it is observed that the collaboration produces a greater acceleration in the learning of a receiver AIS, with respect to the methods compared, with an increase of its successful plans and its successful theories throughout the time, as well as the quantity of successful cycles, the percentage of successful cycles and attaining a performance that required the generation of a smaller quantity of new theories, because of having a more representative theory base of its operation environment, attaining the best performance that results from the learning by collaboration. As new research lines we are exploring the mechanisms that may allow to invoke dynamically the collaboration of an AIS that is at a higher evolutionary stage, considering the selection metrics of theory bases of the AISs, based on percentages of successful cycles, quantity of theories with a utility

over 75 %, quantity of successful cycles and quantity of successful plans. Finally, we consider collaboration experimentation from an AIS that comes from a creator or developer different from that of the receiver AIS, given the proposed extension of the LLC (Learning Life Cycle). Possible real life examples of applicability of autonomous systems, are guided robot to avoid obstacles, for example to access the rescue of victims of earthquakes, these robots will share their theories and receive collaboration of robots that are in an upper layer of Learning Life Cycle.

References

1. Ierache, J., García Martínez, R., De Giusti, A.: Learning Life Cycle in Autonomous Intelligent Systems in IFIP International Federation for Information Processing. In: Artificial Intelligence and Practice II, Max Bramer, vol. 276, pp. 451–455. Springer, Boston (2008)
2. Fritz, W., García-Martínez, R., Blanqué, J., Rama, A., Adobbati, R., Sarno, M.: The autonomous intelligent system. *Robotics and Autonomous Systems* 5, 109–125 (1989)
3. Borrajo, D., Veloso, M.: Lazy incremental learning of control knowledge for efficiently obtaining quality plans. *AI Review Journal* 11, 371–405 (1997)
4. Sutton, R.: Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In: Proc. 7th Int. Conf. on ML, pp. 216–224. Kaufmann, San Francisco (1990)
5. Carbonell, J.G., Gil, Y.: Learning by experimentation: The operator refinement method. In: Michalski, Kodratoff (eds.) *Machine Learning: An AI Approach*, vol. III, pp. 191–213. Morgan Kaufmann, San Francisco (1990)
6. Wang, X.: Planning while learning operators. PhD thesis, School of Computer Science, Carnegie Mellon University (1996)
7. García-Martínez, R., Borrajo, D.: Planning, learning, and executing in autonomous systems. In: Steel, S. (ed.) *ECP 1997. LNCS*, vol. 1348, pp. 208–220. Springer, Heidelberg (1997)
8. García Martínez, R., Borrajo, D.: An Integrated Approach of Learning, Planning and Executing. *Journal of Intelligent and Robotic Systems* 29, 47–78 (2000)
9. Maceri, D., García-Martínez, R., Borrajo, D., Britos, P.: Learning by Knowledge Sharing in Autonomous Intelligent Systems. In: Sichman, J.S., Coelho, H., Rezende, S.O. (eds.) *IBERAMIA 2006 and SBIA 2006. LNCS (LNAI)*, vol. 4140, pp. 128–137. Springer, Heidelberg (2006)
10. Braitenberg, V.: *Vehicles: Explorations In Synthetic Psychology*. MIT Press, Cambridge (1984)

PART VII
AI Techniques

Context-Based Probabilistic Scene Interpretation

Bernd Neumann and Kasim Terzic

Department of Informatics, University of Hamburg
{neumann,terzic}@informatik.uni-hamburg.de

Abstract. In high-level scene interpretation, it is useful to exploit the evolving probabilistic context for stepwise interpretation decisions. We present a new approach based on a general probabilistic framework and beam search for exploring alternative interpretations. As probabilistic scene models, we propose Bayesian Compositional Hierarchies (BCHs) which provide object-centered representations of compositional hierarchies and efficient evidence-based updates. It is shown that a BCH can be used to represent the evolving context during stepwise scene interpretation and can be combined with low-level image analysis to provide dynamic priors for object classification, improving classification and interpretation. Experimental results are presented illustrating the feasibility of the approach for the interpretation of facade images.

Keywords: Scene interpretation, probabilistic scene models, context-based interpretation.

1 Introduction

Scene interpretation has the goal to understand visual scenes beyond single-object recognition. The basic architecture of a knowledge-based scene interpretation system is shown in Fig. 1.




Fig. 1. Basic architecture for scene interpretation

Image analysis procedures first determine a low-level description of the image (or image sequence) in terms of patches or segments which at best correspond to meaningful primitive objects. The task of high-level scene interpretation is to determine meaningful object configurations, situations and activities based on the primitive objects provided by low-level analysis, and on models in a conceptual

knowledge base. For many application domains, *compositional hierarchies* provide a natural structure for the conceptual knowledge base. A compositional hierarchy consists of *aggregates* which describe a scene in terms constituents at a lower hierarchical level, meeting certain conditions. Typical structures of a compositional hierarchy are shown in Fig. 2 for the facade domain. Although the facade domain is used for examples in this contribution, the approach is generic in nature. Note that the conceptual knowledge base distinguishes between models of physical objects (e.g. *B-Door*) and models of object appearances in images (e.g. *B-Door-View*).

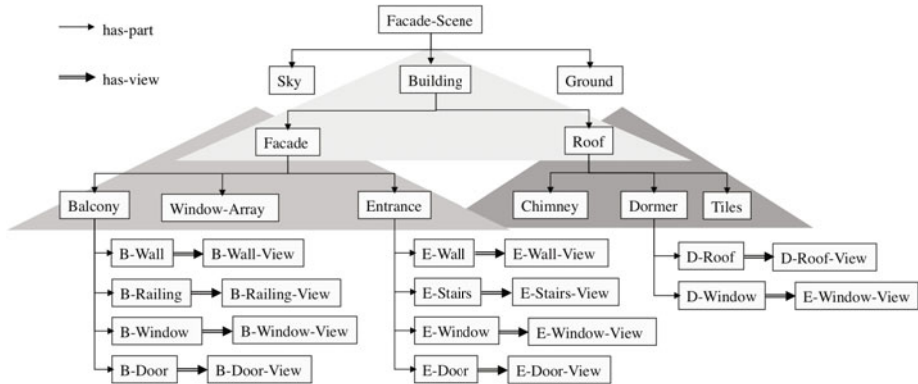


Fig. 2. Compositional hierarchy for facade scene (not all components shown). Triangles mark aggregate structures.

In this contribution, we present a new approach to uncertainty management for stepwise high-level scene interpretation. As evidence is incorporated step by step, alternative interpretation possibilities arise which can be ranked effectively by the probabilistic high-level context based on preceding interpretation steps.

The paper is structured as follows. In the next section, we discuss related work. In Section 3, we present the conceptual probabilistic framework developed for high-level scene interpretation. This framework has been implemented in an operational interpretation system called SCENIC. System details and experimental results are described in Section 4. We conclude with a summary of the results.

2 Related Work

Compositional hierarchies have been employed for high-level scene interpretation by many researchers [1, 2, 3, 4, 5] with basically non-probabilistic (crisp) frame-based representations, as commonly used in AI. Rimey [6] was the first to model compositional hierarchies with tree-shaped Bayesian Networks (BNs), requiring parts of an aggregate to be conditionally independent. Koller and coworkers [7, 8] extended BNs in an object-oriented manner for the representation of structured objects. Their Probabilistic Relational Models allow augmenting a crisp relational structure with an arbitrary probabilistic dependency structure. Gyftodimos and Flach [9] introduce hierarchical BNs to allow multiple levels of granularity. While these contributions

improve the expressive power of BNs, they do not specifically support compositional hierarchies of aggregates as required for context modelling in scene interpretation. For this purpose, Bayesian Compositional Hierarchies (BCHs) were developed [10]. This structure mirrors the formal logic-based ontologies of [5] but allows for a probabilistic ranking of alternative interpretations. An interesting alternative approach is provided by Mumford and Zhu [11] where a grammatical formalism takes the place of hierarchical knowledge representation and parsing algorithms are applied for scene interpretation, leading to efficient processing but complicating the integration with large-scale knowledge representation.

Context through visual clues has been in the focus of recent work [12, 13]. Nearby visual clues are used to reduce the number of false positives. However other high-level knowledge can not be integrated directly, and visual context will usually be too weak to provide conclusions. Li et al. [14] presented a way to incorporate a global context: If the type of picture is known (e.g. the type of sport shown), then the priors for the classes can be changed correspondingly to increase the chance of correct classification. However, details of the contextual relationships between scene objects are not exploited. This is possible in our approach where dynamic context exploitation is integrated into a general probabilistic framework for scene interpretation.

3 Probabilistic Models for High-Level Interpretation

In this section, we will describe the probabilistic framework developed for scene interpretation. We will first give a formal specification in a general form (Section 3.1) and suggest beam search as a suitable interpretation algorithm. We then show how probabilistic scene interpretation can be realized efficiently with a probabilistic model organized as a Bayesian Compositional Hierarchy (BCH) (Section 3.2). We finally present a further specialization where probabilistic information is separated into a structural and an appearance part. This way, object classification can be based on an appearance-based low-level component and a structure-based high-level component which provides dynamic expectations based on the evolving scene context.

3.1 General Probabilistic Framework

In a general form, probabilistic scene interpretation can be modelled as evidence-based reasoning with large joint probability distributions (JPDs). Let $\underline{R}_{\text{all}} = \underline{R}_1 \dots \underline{R}_N$ be vector-valued random variables describing *scene objects* of a conceptual knowledge base (properties of primitive objects and aggregates of a compositional hierarchy), and let $\underline{V}_{\text{all}} = \underline{V}_1 \dots \underline{V}_K$ be random variables describing *view objects* associated with the scene objects (see Fig. 2), i.e. entities which can be observed in an image, such as groups of SIFT features or other low-level descriptors. For the purposes of this contribution, we will assume that each primitive scene object is related to exactly one view object which describes its possible appearances in an image. Different objects may look alike, especially under noisy conditions, hence observed evidence may not be unambiguously assigned to a view object, causing alternative interpretation candidates.

A probabilistic model for a scenario with a particular set of objects is given by a JPD $P(\underline{R}_{\text{all}} \underline{V}_{\text{all}})$ where $\underline{R}_{\text{all}}$ are hidden variables and $\underline{V}_{\text{all}}$ are observables. This JPD can be used to express all relevant probabilistic dependencies:

- how the properties of observations are related to properties of scene objects,
- how the properties of a specific scene object are related to each other,
- how properties of one scene object are related to properties of other scene objects.

To represent scenarios which may occur with different object configurations (sets of objects), we combine the corresponding JPDs as alternative models using a random variable S with domain $1..M$ and distribution $P_S(1)=p_1 \dots P_S(M)=p_M$ to describe the selection of an alternative model:

$$P_{\text{scenario}} = P_S(S) P^{(S)}(\underline{R}_{\text{all}}^{(S)} \underline{V}_{\text{all}}^{(S)}) \quad (1)$$

Typically, the objects of a scenario will only be a subset of all objects present in a realistic scene, hence we have to account for "noise objects". To simplify the following presentation, we will only consider scenes where the observed evidence exactly matches one or more of the alternative models in P_{scenario} . A treatment of noise objects and missing evidence can be added without affecting the basic procedure.

Let $\underline{e}_{\text{all}}$ be all evidence about a scene, then scene interpretation amounts to determine the most probable instantiation of one of the alternative models based on this evidence. This can be achieved by two maximization steps.

1. Determine the model m with JPD $P^{(m)}$ and the view objects $\underline{V}_{\text{all}}^{(m)}$ which are most probable given the observed evidence.

$$\max_{m, \underline{V}_{\text{all}}^{(m)}} P_{\text{scenario}}(m \underline{V}_{\text{all}}^{(m)} | \underline{e}_{\text{all}}) = \max_{m, \underline{V}_{\text{all}}^{(m)}} p_m P^{(m)}(\underline{e}_{\text{all}}) \quad (2)$$

2. For the winning model m and view objects $\underline{V}_{\text{all}}^{(m)}$, determine the instantiations of hidden variables $\underline{R}_{\text{all}}$ which maximize

$$\arg \max_{\underline{R}_{\text{all}}} P^{(m)}(\underline{R}_{\text{all}}^{(m)} = \underline{r}_{\text{all}} | \underline{V}_{\text{all}}^{(m)} = \underline{e}_{\text{all}}) \quad (3)$$

The underlying assignment problem is formidable as illustrated in Fig. 3. In a stepwise process, noisy and ambiguous evidence must be assigned to one of several models in one of a large number of permutations, therefore a global maximization as defined by Eqs. 2 and 3 is rarely feasible for real scene interpretation tasks. Even in restricted domains such as the facade domain, the probabilistic model P_{scenario} may consist of hundreds of alternative models, each comprised of hundreds of random variables representing scene object and evidence properties with potentially large value ranges. Hence the search space may be very large and approximate methods are required.

As a general strategy for approximate probabilistic scene interpretation we propose beam search [15], where only promising alternative partial interpretations are kept (in the "beam"), while improbable ones are discarded. There are three reasons for this choice. First, beam search is based on incremental evidence assignment which allows to generate dynamic expectations, indispensable for the analysis of time-varying scenes and real-time recognition tasks. Second, beam search permits scalable parallelism




Fig. 3. Evidence assignment to view objects of alternative models

by selecting an appropriate beam width. Third, beam search corresponds well to our formulation of probabilistic scene models in terms of alternatives.

For scene interpretation with beam search, the following steps have to be executed:

- A Initialise the beam with all alternative models given by Eq. 1.
- B Select unassigned evidence \underline{e} .
- C Determine possible assignments of \underline{e} for each model in the beam, clone models in case of multiple assignments, discard models where no assignment is possible.
- D Rank models using Eq. 2 (where $\underline{e}_{\text{all}}$ now stands for all evidence observed so far), discard unlikely models from beam.
- E Repeat B to E until all evidence is assigned.
- F Select highest-ranking model, determine final interpretation using Eq. 3.

Note that evidence selection in Step B can be crucial in the initial phase where non-discriminating evidence could give rise to a possibly large number of alternative models. A selection which maximizes information gain is advisable, as in decision tree methodology [16].

Efficient storage of P_{scenario} and computation of the marginalisations in Eq. 2 may easily become a bottleneck for realistic tasks. Therefore, Bayesian Network technology is required and the dependency structure of object properties plays an important part. In the following section, we will present Bayesian Compositional Hierarchies [10] which allow arbitrary dependencies within an aggregate but provide a tree-shaped dependency structure between aggregates and thus efficient computational procedures in tune with compositional hierarchies.

3.2 Bayesian Compositional Hierarchies

Probabilistic dependencies in compositional hierarchies for high-level scene interpretation have characteristics which can be exploited for efficient storage and computation. For one, strong mutual dependencies are often confined to the parts of an aggregate and less stringent between parts of different aggregates. Second, higher-level representations often abstract from lower-level details. We therefore propose a probabilistic model, called Bayesian Compositional Hierarchy (BCH), where aggregates (i) are modelled by arbitrary JPDs, (ii) possess an external description abstracting from details about their parts, and (iii) depend on other aggregates only via the part-of relations of the compositional hierarchy.

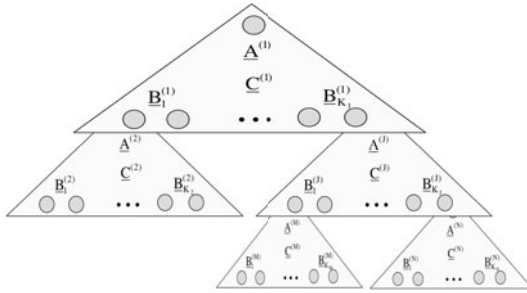


Fig. 4. Structure of a BCH. Triangles represent aggregates, circles represent parts.

Fig. 4 shows the schematic structure of a BCH. Each aggregate is described individually by a JPD $P(\underline{A} \ \underline{B}_1 \dots \underline{B}_K \ \underline{C})$ where \underline{A} is the external description, $\underline{B}_1 \dots \underline{B}_K$ are descriptions of the parts, and \underline{C} expresses conditions on the parts. For example, a *Balcony* can be described by bounding boxes \underline{B}_1 , \underline{B}_2 , \underline{B}_3 for *B-Door*, *B-Window* and *B-Railing*, distances \underline{C} between these parts, and a bounding box \underline{A} for the *Balcony* as a whole. The hierarchy is constructed by taking the external aggregate descriptions at a lower level as part descriptions at the next higher level, hence $\underline{B}_1^{(1)} = \underline{A}^{(2)}$ etc.

The JPD of the complete hierarchy is given by

$$P(\underline{A}^{(1)} \dots \underline{A}^{(N)}) = P(\underline{A}^{(1)}) \prod_{i=1..N-1} P(\underline{B}_1^{(i)} \dots \underline{B}_{K_i}^{(i)} \ \underline{C}^{(i)} \mid \underline{A}^{(i)}) \tag{4}$$

This remarkable formula shows that the JPD of a BCH can be easily constructed from individual aggregate representations, and belief updates can be performed by propagation along the tree structure. Let $P'(\underline{B}_i)$ be an update of $P(\underline{B}_i)$, by evidence or propagation from below. Then the updated aggregate JPD is

$$P'(\underline{A} \ \underline{B}_1 \dots \underline{B}_K \ \underline{C}) = P(\underline{A} \ \underline{B}_1 \dots \underline{B}_K \ \underline{C}) P'(\underline{B}_1) / P(\underline{B}_1) \tag{5}$$

A similar equation holds when $P(\underline{A}^{(m)})$ is updated by propagation from above.

If the aggregate JPDs are Multivariate Gaussians, propagation can be performed by closed-form solutions [10]. Multivariate Gaussians are also very convenient for implementing the interpretation by beam search described in Section 3.1. The marginalisations required for ranking alternative interpretations according to Eq. 2 are directly available from the aggregate covariances, and the final maximising interpretation according to Eq. 3 can be given in terms of the mean values.

3.3 Separating Appearance and Structural Context

In the probabilistic framework presented so far, evidence assignment is governed by a distribution which captures appearance properties of objects as well as the influence of the structural context, e.g. position and orientation of other objects. Considering the rich methodology for appearance-based object classification and the need to modularize, it is useful to separate the two knowledge sources. In this section, we present an approach where a high-level probabilistic model is used to represent the structural context providing structure-based dynamic priors for an appearance-based probabilistic classifier.

We assume that each evidence \underline{e} can be decomposed into two parts $\underline{e} = [\underline{e}_a \ \underline{e}_s]^T$ where \underline{e}_a describes appearance properties (e.g. size, shape, colour, texture) and \underline{e}_s describes structurally relevant properties (e.g. position, distance to reference objects). Furthermore, we assume that both components are statistically independent, given knowledge about which view object (and hence which scene object) has caused this evidence. Hence

$$P^{(m)}(\underline{e}_a \underline{e}_s \mid \underline{V}_i^{(m)}) = P^{(m)}(\underline{e}_a \mid \underline{V}_i^{(m)}) P^{(m)}(\underline{e}_s \mid \underline{V}_i^{(m)}) \quad (6)$$

In beam search, let \underline{e} be the evidence to be assigned to some view object $\underline{V}_i^{(m)}$ and let $\underline{e}_{\text{past}}$ be past evidence already assigned to view objects $\underline{V}_{\text{past}}^{(m)}$. Then to perform the ranking in Step D, we have to compute

$$\begin{aligned} P(m \underline{V}_i^{(m)} \mid \underline{e}_a \underline{e}_s \underline{e}_{\text{past}} \underline{V}_{\text{past}}^{(m)}) &= \frac{P(\underline{e}_a \mid m \underline{V}_i^{(m)} \underline{e}_{\text{past}} \underline{V}_{\text{past}}^{(m)}) P(m \underline{V}_i^{(m)} \underline{e}_s \mid \underline{e}_{\text{past}} \underline{V}_{\text{past}}^{(m)})}{P(\underline{e}_a \underline{e}_s \mid \underline{e}_{\text{past}} \underline{V}_{\text{past}}^{(m)})} \quad (7) \\ &= \frac{P^{(m)}(\underline{e}_a \mid \underline{V}_i^{(m)}) p_m P^{(m)}(\underline{V}_i^{(m)} = \underline{e}_a \mid \underline{e}_s \underline{e}_{\text{past}} \underline{V}_{\text{past}}^{(m)})}{\sum_{m, \underline{V}_i^{(m)}} P^{(m)}(\underline{e}_a \mid \underline{V}_i^{(m)}) p_m P^{(m)}(\underline{V}_i^{(m)} = \underline{e}_a \mid \underline{e}_s \underline{e}_{\text{past}} \underline{V}_{\text{past}}^{(m)})} \quad (8) \end{aligned}$$

The numerator of Eq. 8 specifies the probability that a specific view object (and hence a specific scene object) in model m has the appearance \underline{e}_a . This is a likelihood function which can be provided by low-level image analysis, for example, based on a probabilistic decision tree. The second term of the numerator specifies the probability of observing \underline{e}_a given the structural information \underline{e}_s and past evidence assignments. This is information which can be provided by high-level interpretation, for example, in terms of a BCH. It provides a dynamic prior for the appearance-based likelihood function of the first term. The denominator, finally, is a normalizing factor.

In the following section, we will report about scene interpretation experiments where Eq. 8 is used to rank evidence assignment to primitives of a structural facade model represented by a BCH. It is shown that this improves classification and interpretation performance as compared to single-object classification based only on appearances.

4 Evaluation of Context Influence

In this section, we present two experiments illuminating the probabilistic interaction of low-level and high-level interpretation processes. In the first experiment, described in more detail in [17], the effect of probabilistic context on the classification rate was tested in the domain of building floors. In the second experiment, probabilistic models of entire facades were used to interpret and combine conflicting and error-prone segmentations from several low-level algorithms.

A major problem in this domain is the similar appearance of objects which play different roles in structural models and hence must be distinguished, e.g. balcony doors vs. entrance doors. This makes structural context a very important factor for disambiguation and complete understanding of a scene, while motivating appearance models to be oriented at coarser groupings than the roles in structural models. For our

experiments, we therefore developed an appearance model in terms of a probability distribution for classes C such as doors and windows instead of balcony doors or balcony windows. The probabilities $P(\underline{C}|\underline{E}_a)$ were estimated by decision trees learnt from a training set of annotated facade images, as described in [18]. Using Bayes' rule with priors determined from the dataset, $P(\underline{E}_a|C)$ can be obtained as required for Eq. 8. A BCH modelling facade floors was learnt from the same training examples.

In the first experiment, the system was evaluated on a database of 393 floors from the facade domain, which were extracted from the eTRIMS [19] annotated image database. In this database, each object in the image is annotated with a bounding polygon and a class label. Aggregates above the floor level, such as *Facade* and *Building* were not used for this experiment. Fig. 5 shows several examples from the set of floors used.



Fig. 5. Examples of the variety of floors within the domain

The appearance features extracted from segmentations were area, aspect ratio, rectangularity and compactness. The context was determined based on the position of the object in the image and its probabilistic relation to the positions of previously integrated objects. The task was to classify the annotated polygons.

The performance was evaluated using ten-fold cross-validation. The results are shown in Fig. 6 left. An overall 13% increase in the classification rate shows that scene context can significantly improve classification. The beam search algorithm was used with beam width 1, i.e. only the best alternative was kept at each step. This led to some compounded errors when wrong context was established early, following a wrong initial classification. We expect that the results will be further improved with a beam keeping several alternatives.

In the second experiment, the system was evaluated as a source of context for segmenting a complete facade image. In this instance, the BCH contained six different probabilistic models of facade structures. The goal was to recognise the structure of a facade image and to use the evolving context to guide the segmentation process. Three separate segmentation algorithms were used for this experiment (i) a trainable window detector based on Adaboost [20], (ii) an interest region detector [21], and (iii) a detector of stable regions in scale space [22]. Each one of them was run on 190 images, and the results compared to the hand annotations. Each detected polygon was labelled with the correct class from the corresponding annotation in the case of sufficient overlap (50% in this case), or as a ‘‘False Positive’’ if there was no correspondence. Then a

10-dimensional feature vector consisting of the four shape descriptors from the previous experiment and the means and standard deviations of the three colour channels (RGB) was computed for each detected region. Then the class likelihood distribution was modelled by a decision tree for each detector separately, including the likelihood that a false positive produced the detection.

fold	BCH	DT	BCH & DT
0	0.65	0.70	0.75
1	0.66	0.65	0.78
2	0.62	0.64	0.77
3	0.69	0.72	0.78
4	0.68	0.68	0.75
5	0.62	0.65	0.76
6	0.67	0.71	0.83
7	0.63	0.70	0.83
8	0.66	0.68	0.76
9	0.63	0.71	0.79
Average	0.652	0.685	0.780
Std. Dev.	0.025	0.029	0.029



Fig. 6. left (Exp. 1): Classification rate obtained by 10-fold cross-validation. "BCH" shows the probabilistic context based on the position alone, "DT" the class probability based on appearance features estimated by the decision tree, the last column shows the combined result. Fig. 6 middle and right (Exp. 2): Facade image (middle) and result of context-based segmentation (right). The BCH used six structural models, out of which the most fitting one was automatically selected and used to interpret the polygons provided by the low level.

As in the first experiment, the region with the highest classification confidence was classified and the context propagated, changing the class posteriors of the remaining regions. Although the combination of the three detectors found many false positives (over 80%) in addition to correct detections, the strong probabilistic context managed to enable a complete segmentation of the image, as shown in Fig. 6 right. Most regions were discarded as false positive detections, and the remaining detections interpreted according to the most fitting of the six structural models.

5 Conclusion

A new general framework for stepwise probabilistic scene interpretation has been presented, where evidence classification is supported by dynamic expectations generated from the evolving high-level context. Experiments with learnt models of facade images demonstrate significant improvements over context-free classification and show that weak low-level results can be compensated by high-level context.

References

1. Tsotsos, J.K., Mylopoulos, J., Covvey, H.D., Zucker, S.W.: A Framework for Visual Motion Understanding. In: IEEE PAMI-2, pp. 563–573 (1980)
2. Nagel, H.-H.: From Image Sequences towards Conceptual Descriptions. Image and Vision Computing 6(2), 59–74 (1988)

3. Neumann, B.: Description of Time-Varying Scenes. In: Waltz, D. (ed.) *Semantic Structures*, pp. 167–206. Lawrence Erlbaum, Mahwah (1989)
4. Georis, B., Mazière, M., Brémond, F., Thonnat, M.: Evaluation and Knowledge Representation Formalisms to Improve Video Understanding. In: *Proc. IEEE International Conf. on Computer Vision Systems ICVS 2006*, p. 27. IEEE Computer Society, Los Alamitos (2006)
5. Neumann, B., Moeller, R.: On Scene Interpretation with Description Logics. In: Christensen, H.I., Nagel, H.-H. (eds.) *Cognitive Vision Systems*. LNCS, vol. 3948, pp. 247–275. Springer, Heidelberg (2006)
6. Rimey, R.D.: Control of Selective Perception using Bayes Nets and Decision Theory. TR 468, Univ. of Rochester, Computer Science Department, Rochester, USA 14627 (1993)
7. Koller, D., Pfeffer, A.: Object-oriented Bayesian Networks. In: *The Thirteenth Annual Conference on Uncertainty in Artificial Intelligence*, pp. 302–313 (1997)
8. Getoor, L., Taskar, B.: *Introduction to Statistical Relational Learning* (eds, pp. 129–174. The MIT Press, Cambridge (2007)
9. Gyftodimos, E., Flach, P.A.: Hierarchical Bayesian Networks: A Probabilistic Reasoning Model for Structured Domains. In: de Jong, E., Oates, T. (eds.) *Proc. Workshop on Development of Representations, ICML*, pp. 23–30 (2002)
10. Neumann, B.: Bayesian Compositional Hierarchies - A Probabilistic Structure for Scene Interpretation. TR FBI-HH-B-282/08, Univ. of Hamburg, Department Informatik (2008)
11. Mumford, D., Zhu, S.-C.: *A Stochastic Grammar of Images*. Now Publishers (2007)
12. Yu, J., Luo, J.: Leveraging Probabilistic Season and Location Context Models for Scene Understanding. In: *The 2008 International Conference on Content-based Image and Video Retrieval (CIVR 2008)*, pp. 169–178. ACM, New York (2008)
13. Perko, R., Wojek, C., Schiele, B., Leonardis, A.: Integrating Visual Context and Object Detection within a Probabilistic Framework. In: Paletta, L., Tsotsos, J.K. (eds.) *WAPCV 2008*. LNCS, vol. 5395, pp. 54–68. Springer, Heidelberg (2009)
14. Li, L.-J., Socher, R., Fei-Fei, L.: Towards Total Scene Understanding: Classification, Annotation and Segmentation in an Automatic Framework. In: *Computer Vision and Pattern Recognition, CVPR* (2009)
15. Lowerre, B.: *The Harpy Speech Recognition System*. Ph.D. thesis, Carnegie Mellon University (1976)
16. Mitchell, T.M.: *Machine Learning*. The Mc-Graw-Hill, New York (1997)
17. Kreuzmann, A., Terzić, K., Neumann, B.: Context-aware Classification for Incremental Scene Interpretation. In: *Proc. Workshop on Use of Context in Vision Processing (UCVP 2009)*, Boston (2009)
18. Terzić, K., Neumann, B.: Decision Trees for Probabilistic Top-down and Bottom-up Integration. Technical Report FBI-HH-B-288/09, University of Hamburg (2009)
19. etrim - e-Training for Interpreting Images of Man-made Scenes, <http://www.ipb.uni-bonn.de/projects/etrim/>
20. Terzić, K., Hotz, L., Šochman, J.: Interpreting Structures in Man-made Scenes; Combining Low-Level and High-Level Structure Sources. In: *International Conference on Agents and Artificial Intelligence (ICAART 2010)*, Valencia (2010)
21. Jahangiri, M., Petrou, M.: Fully Bottom-Up Blob Extraction in Building Facades. In: *Proc. Pattern Recognition and Image Analysis, PRIA* (2008)
22. Drauschke, M.: An Irregular Pyramid for Multi-Scale Analysis of Objects and Their Parts. In: Torsello, A., Escolano, F., Brun, L. (eds.) *GbrRPR 2009*. LNCS, vol. 5534, pp. 293–303. Springer, Heidelberg (2009)

Full Text Search Engine as Scalable k-Nearest Neighbor Recommendation System

Ján Suchal and Pavol Návrat

Slovak University of Technology
Faculty of Informatics and Information Technologies
Ilkovičova 3, 842 16 Bratislava, Slovakia
suchal@fiit.stuba.sk, navrat@fiit.stuba.sk

Abstract. In this paper we present a method that allows us to use a generic full text engine as a k-nearest neighbor-based recommendation system. Experiments on two real world datasets show that accuracy of recommendations yielded by such system are comparable to existing spreading activation recommendation techniques. Furthermore, our approach maintains linear scalability relative to dataset size. We also analyze scalability and quality properties of our proposed method for different parameters on two open-source full text engines (MySQL and SphinxSearch) used as recommendation engine back ends.

Keywords: full text search, recommendation systems.

1 Introduction

Recommendation systems are becoming very common in many domains, but their results can be mostly seen as recommendations in online shops (e.g. Amazon.com), news (e.g. Google news) or social networks sites (e.g. Facebook). In current era of web applications, to create a good and also a scalable recommendation system is not an easy task. Typically very specialized systems are developed to deal with the problem of high quality recommendations on large datasets.

In this paper we present a new method that allows to use a generic and generally available full text engine back ends (MySQL and SphinxSearch) to generate recommendations based on a k-nearest neighborhood approach. First we explain how data needs to be preprocessed into virtual documents to exploit full text engine capabilities for recommendation generation and present our k-nearest neighbor algorithm. Next, we discuss how neighborhood size affects precision of such recommendations. We show that long non-scalable queries are not present in real-world datasets by exploiting power-law distributions in datasets and by showing negligible precision gains from long queries. We also show that our method yields recommendations with comparable precision as spreading activation techniques normally used in recommendation systems. Furthermore we address scalability of our method showing that it yields linear scalability relative to

dataset size. For our experiments two medium-sized real world datasets have been used.

2 Full Text Recommendation Algorithm

In order to be able to use a generic full text engine as recommendation system we need to transform data into special structures - virtual documents. We create two classes of such virtual documents - per-user and per-item virtual documents. Suppose we have data available from an e-shop consisting of records who bought which items. A per-user virtual document consists of user identifier and concatenation of item identifiers that this user is connected to. Per-item virtual document contains item identifier and space separated user ids. Tables 1 and 2 show examples of such virtual documents.

Table 1. Per-item virtual documents

<i>Item identifier</i>	<i>Users</i>
item1	user1 user2 user3
item2	user1 user2 user4
item3	user4

Table 2. Per-user virtual documents

<i>User identifier</i>	<i>Items</i>
user1	item1 item2
user2	item1
user3	item1 item2
user4	item3

Issuing specialized full text queries on such virtual documents enables a quick retrieval of similar items or users. For example finding top-k similar users to a given *user1* can be rewritten to a simple full text query: "*item1*|*item2*" where | represents the OR operator and *item1* and *item2* are items already seen by *user1*. Furthermore such full text query returns results ordered by relevance, which is in our case similarity. Finding similar items based on a list of users can be done by querying per-item virtual documents.

Since full text engines typically use weighting schemes[1] by using full text indices on per-user documents, lower weights are automatically given to items that are common for all users and higher weights for items that are connected to the same user several times (such as items bought by the same user on a regular basis). Similarly for per-item documents, users that tend to be connected to more items get lower weight since they are less discriminative.

By using these virtual documents and similarity queries we can now formulate our top-k nearest neighbor recommendation algorithm using pseudo code as follows:

```
Inputs: user_id, k, neighborhood size N
Output: top-k recommended items for user user_id

function recommend_with_fulltext(user_id, k, N)
    items = find_items_connected_to(user_id)
    items_query = create_query(items.ids)
    n_similar_users = query_per_user(items_query, N)
    n_similar_users.remove(user_id) # remove current user
    user_query = create_query(n_similar_users.ids)
    # find items not already seen, based on similar users
    similar_items = query_per_item(user_query, k, items)
    return similar_items
end
```

At first, user history is retrieved by finding all items connected to a given user. Next a “full text” query is constructed by concatenating item identifiers with OR-operator. This query is used to retrieve top N similar users (with similar history) from per-user virtual documents table. Next a specialized “full text” query is again constructed by concatenating user identifiers with OR-operator. Finally this query is used to retrieve top k items, but excluding items already in user history, since we do not want to recommend items that user already is connected to. In short, this algorithm work in such a way that it first finds similar users based on user’s history and then finds top-k items that these similar users have in their history.

This algorithm however makes an important and hidden assumption that needs to be addressed. Since full text engines are heavily optimized for short query strings (in our case low number of item/user identifiers in a query) and generally do not scale well for very long query strings, we need to make sure that *find_items_connected_to* does return only a limited number of items. This is a problem in theory, but in practice and thanks to nature of real-world datasets, we are mostly dealing with power-law distributions where most users are connected only to a very limited subset of items. In those rare cases where users are connected to a large number of items, a filtering based on a secondary heuristic (e.g. information gain or time-based weighting) should be considered. For the sake of simplicity, we ignore this detail in our further evaluation since implicitly used full text tf-idf weighting already addresses this problem at an intermediate level.

3 Evaluation

In order to test recommendation quality, speed and scalability we picked two real datasets for evaluation. Githubcontest data published by github.com and 15 days of server access logs from top Slovak online news publisher www.sme.sk.

Github.com is an online social collaborative repository hosting service based on git source control management, where users can create and share source code. Users can even fork and watch repositories of other users. A subset of github.com database was released as a challenge to build a recommendation system that recommends users

which repositories they should also watch based on repositories they are watching already, on their programming language preferences and other attributes. Github.com evaluated recommendation systems by querying top10 repository recommendations for 4788 users giving one point for each correctly guessed repository. In our evaluation we use the same approach, but for clarity we use ratio of correct guesses to all guesses and refer to this value as precision on top 10 – P10.

As a part of collaboration with largest Slovak online news publisher we have access to server logs of their online news portal www.sme.sk having roughly one million page views a day (excluding discussion forums). To evaluate our recommender system we are recommending articles to users based on their reading history. For evaluation we have used a 15 day window, splitting it into “train” and test sets. Similarly as in github.com contest, we generate top10 recommendations for each user in test set and add a point for each article that this user has actually visited in this test set. Again we evaluate precision on top 10 – P10. Table 3 shows size characteristics of these two datasets.

In evaluation we silently ignore the fact that such an evaluation cannot be used for testing the real precision or quality of any recommendation system since we are unable to simulate how users would actually react given such top-k recommendations. For example in evaluation using a train/test dataset in domain of news articles, a sole fact that a user has not seen a recommended article does not make it a bad recommendation, it could simply mean that the user has not found this article. On the other hand a sole fact that a user has seen a recommended article does not make this recommendation a good one. This could also mean that this recommendation was not necessary since the user found this article even without our recommendation system[2] [3].

However for our purposes (scalability analysis and parameter sensitivity) such train/test dataset evaluation is sufficient.

Table 3. Dataset size characteristics

<i>Dataset</i>	<i>github.com</i>	<i>sme.sk</i>
# of users	56 519	1 023 407
# of items	120 867 (repositories)	162 455 (articles)
# of entries	440 237 (followings)	11 996 530 (pageviews)

For all experiments we used two generic and freely available full text engines (MySQL 5.0 full text and SphinxSearch 0.9.8) as recommendation backends.

3.1 Recommendation Quality

To evaluate the precision of our top-k nearest neighborhood based method we use precision on top-10 metric. Since generic full text engines typically offer different tf-idf weighting schemes we also experimented with possible configurations of engines.

Figure 1 shows recommendation precision as a function of neighborhood size (N) and different engines and their configurations. The big drop for *Sphinx default* at $N=10$ is caused by additional Sphinx's full text word proximity weighting that gives more weight if words in query are close to each other. This is crucial for word counter-productive for our purpose, since we are not interested in which order of users/items appear in virtual documents.

We can see that for $N > 40$ precision gain starts to be negligible, showing that we do not need long queries for recommendation generation. Also Sphinx's *bm25* precision is superior to MySQL for lower N , but for higher N MySQL starts to be slightly more precise. This might seem as an advantage for MySQL, but quite the contrary is true because higher N means longer queries and scalability problems for large datasets as we show in experiments in section Scalability.

Figure 1 also shows that difference between weighting schemes *wordcount* and *bm25* (Sphinx *bm25* and MySQL) is 3-4%. Here, *wordcount* refers to precision that could be yielded by creating basic inverted indices on user/item data and using our k-nearest neighborhood algorithm.




Fig. 1. Recommendation precision as function of neighborhood size and full text backend engine for github.com dataset

Since this evaluation only compares different configurations of the same algorithm we compared our method also with an existing spreading activation based recommendation method. Spreading activation is a recursive energy distribution algorithm with a threshold stopping criterion [4] [5]. Our results show that spreading activation on

github.com dataset yields best results ($P10 = 19.6\%$) for starting energy = 500 and threshold = 0.1 which is slightly worse than best result of our k-nearest neighborhood based method ($P10 = 20.0\%$).

3.2 Scalability

We have shown that our recommendation method yields comparable results as an existing spreading activation based approach. Figure 2 shows a scatterplot of time needed for recommendation generation versus precision on top10 for such recommendation. Since these are very different implementations to compare in an absolute manner (e.g. recommendation speed) we are only interested in scalability characteristics. Results show that while recommendation time for spreading activation grows exponentially due to neighborhood traversal explosion, k-nearest neighborhood method maintains a low footprint.




Fig. 2. Recommendation time versus precision on top-10 for spreading activation and k-nearest neighborhood based method

Figure 3 shows scalability analysis of recommendation generation time as a function of neighborhood size on github.com dataset for both recommendation engine backends and Sphinx variants. All configurations maintain linear scalability for varying neighborhood size, with *MySQL* being superior to *Sphinx bm25*.

Figure 4 shows scalability as a function of dataset size for larger sme.sk dataset showing that both algorithms maintain linear scalability relative to dataset size. However this analysis also shows that *MySQL* scales much worse, being superior to *Sphinx* only for smaller datasets and having a steeper linear characteristic beyond 250 thousand entries mark.




Fig. 3. Scalability analysis showing time needed for recommendation generation as a function of neighborhood size for various full text backend engines and settings on github.com dataset




Fig. 4. Scalability analysis showing time needed for recommendation generation for different backend full text engines as a function of dataset size on sme.sk dataset (k = 35)

4 Related Work

Our nearest neighborhood recommendation method uses standard tf/idf weighting scheme typically used for document vectorization in content-based recommendation methods [6], but we do not use cosine similarity to find similar vectors.

There are two major approaches to scalability problems when dealing with nearest-neighborhood-based recommendation systems. One approach is to reduce the number of item similarity comparisons required to retrieve most similar items. This can be done by pre-clustering similar items into sets, which are used for pruning possible item candidates [7]. The second, but more challenging approach is to rewrite existing algorithms in such a way that scalable frameworks (e.g. map-reduce) can be effectively utilized [8].

5 Conclusions

In this paper we have proposed a k-nearest neighborhood based recommendation method that exploits generic and generally available full text engines as back ends for quick similarity search. Using experiments on two real world datasets we have also shown that this method not only yields comparable results to spreading activation based recommendation methods, but is superior in means of linear scalability relative to dataset size. We have also addressed the following drawback of full text engines: they are optimized only for short queries. We have shown that a) in real world datasets power-law distributions are present and most users will be connected only to limited subset of items, thus finding similar users quickly is possible, and b) by showing that recommendation precision gains are negligible for high neighborhood sizes that cause long queries.

6 Future work

Our work is planned to be extended by

- adding additional columns to virtual documents, we would be able to match similar users/items based on more than one attribute. For example on github.com dataset we could also use preferred programming languages as an additional weighted similarity metric, or similarly on sme.sk dataset article categories.
- evaluating scalability and precision characteristics for additional backed engines such as PostgreSQL fulltext and Apache Solr.
- evaluating scalability in distributed environments on large datasets.

Acknowledgments. This work was partially supported by the grants VEGA 1/0508/09, KEGA 345-032STU-4/2010 and it is the partial result of the Research & Development Operational Programme for the project Support of Center of Excellence for Smart Technologies, Systems and Services, ITMS 26240120029, co-funded by the ERDF.

References

1. Buckley, C.: The importance of proper weighting methods. In: Human Language Technology Conference (1993)
2. Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.T.: Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems (TOIS)* 22, 1 (2004)
3. Karypis, G.: Evaluation of Item-Based Top-N Recommendation Algorithms. In: Conference on Information and Knowledge Management (2001)
4. Crestani, F.: Application of Spreading Activation Techniques in Information Retrieval. *Artificial Intelligence Review* 11(6), 453–482 (1997)
5. Salton, G., Buckley, C.: On the use of spreading activation methods in automatic information. In: Annual ACM Conference on Research and Development in Information Retrieval (1988)
6. Pazzani, M., Billsus, D.: Content-Based Recommendation Systems. *The Adaptive Web*, 325–341 (2007)
7. Sarwar, B.M., Karypis, G., Konstan, J., Riedl, J.: Recommender Systems for Large-scale E-Commerce: Scalable Neighborhood Formation Using Clustering (2002)
8. Das, A.S., Datar, M., Garg, A., Rajaram, S.: Google news personalization: scalable online collaborative filtering. In: International World Wide Web Conference (2007)

Developing Intelligent Environments with OSGi and JADE

Davide Carneiro¹, Paulo Novais¹, Ricardo Costa², and José Neves¹

¹ Department of Informatics, University of Minho, Braga, Portugal
{dcarneiro, pjon, jneves}@di.uminho.pt

² College of Management and Technology - Polytechnic of Porto, Felgueiras, Portugal
rfc@estgf.ipp.pt

Abstract. The development of intelligent environments poses complex challenges, namely at the level of device heterogeneity and environment dynamics. In fact, we still lack supporting technologies and development approaches that can efficiently integrate different devices and technologies. In this paper we present how a recent integration of two important technologies, OSGi and Jade, can be used to significantly improve the development process, making it a more dynamic, modular and configurable one. We also focus on the main advantages that this integration provides to developers, from the Ambient Intelligence point of view. This work results from the development of two intelligent environments: VirtualECare, which is an intelligent environment for the monitoring of elderly in their homes and UMCourt, a virtual environment for dispute resolution.

Keywords: Ambient Intelligence, Online Dispute Resolution, Multi-agent Systems, Service-Oriented Architectures, OSGi, Jade.

1 Introduction

Ambient Intelligence is a relatively new field of Artificial Intelligence. In this paradigm, computers are seen as a proactive tool that assists us in our day to day. For the first time, the user is placed at the center of the computer-human interaction, which constitutes a major shift in the traditional paradigm [9]. In fact, in the past we had to move ourselves to the vicinity of a computer in order to interact with it, using old-fashioned interaction means. Now we are interacting with environments that are built on numerous and distributed small computers that communicate, embedded in our common devices.

Evidently, the development of such environments poses complex challenges, namely because these are highly dynamic environments, include heterogeneous devices, are very complex to model, and need to be reliable. In this paper we address some of these challenges by bringing together two fields from computer science: Multi-agent Systems and Service-Oriented Architectures. Multi-agent Systems (MAS) [1] emerged from the combination of Artificial Intelligence with distributed computational models, generating a new paradigm: distributed artificial intelligence.

From the Ambient Intelligence point of view a MAS can be defined as a group of entities (software or hardware) which will “read” the environment they are in and take decisions in order to achieve some common goal (like the comfort or safety of persons in the environment) based on knowledge from every agent in the system. In to the so-called Service Oriented Architectures [2], functionalities are provided in the form of technology-independent services, based on three key concepts: the service provider, the service user and the service registry.

In order to develop an intelligent environment that can incorporate insights from these two trends, we consider in this paper two well known technologies: Jade agent platform [3] and OSGi service platform [4]. Jade (Java Agent Development Framework) is a software framework that significantly facilitates the development of agent-based applications in compliance with the FIPA specifications. FIPA (Foundation for Intelligent Physical Agents) promotes standards that aim at the interoperability and compatibility of agents [10]. The use of OSGi (Open Services Gateway Initiative) allows developers to build java applications on a modular basis. The resulting modules are called bundles, which are not only competent to provide services, but also to use services provided from other bundles. In OSGi, a bundle can be installed, started, stopped or un-installed at run-time and without any kind of system reboot, making OSGi-based technologies very modular and dynamic.

Each of these two technologies has already been used successfully to implement a variety of projects in this field following two separate approaches (e.g. [11, 12, 13]). In fact, each one has, as will be seen ahead, characteristics that can be incorporated in intelligent environments, resulting in advantages that are noticed in all the phases of the development, ranging from the specification and design of the system to the actual deployment and use. Moreover, there are approaches that use these two technologies together. [17] for example presents one for using Jade and OSGi, together with a methodology for integration. However, a recent development in these technologies has made this step a much simpler one: the recently released version 3.7 of Jade agent platform, which integrates the OSGi technology. Using this integration it is now possible to run JADE agents inside an OSGi environment, package agents code inside separated bundles, update it by means of the bundle update feature of OSGi and give agents access to all typical OSGi features such as registering and using OSGi services. This constitutes an opportunity that must be exploited in order to simplify the process of developing complex intelligent environments.

In this paper we do not present a new methodology for developing agents nor for developing service-based applications. Instead, we present an approach that builds on this new important development and that can make use of existing approaches for agent and service development individually. As an example, a developer could make use of VisualAgent [18] to easily develop the agents of the system and rely on a Model-Driven Architecture like the one presented in [19] to implement the OSGi services. It results in an iterative and scalable process for developing highly modular applications targeted at virtual environments.

Specification of the Problem

Intelligent environments can be considered in many different domains, namely domestic, medical, legal, public spaces, workplaces, among others. In this paper, given our previous experience, we consider the development of domestic environments for

the enhancement of the comfort and security of the user, with an emphasis on home care. In order to highlight the applicability of the presented approach, we will also use as example the development of UMCourt, a platform-independent virtual environment in which parties in dispute find alternatives to a disadvantageous dispute in court. This delimitation of the problem helps to define which devices and services can be considered.

In the case of UMCourt, as it is a virtual environment, it can be accessed from any regular internet-enabled device. The services considered are used transparently by the agents that represent the users and these users do not know with which services they are interacting. They simply request tasks that are delivered to them. In the case of VirtualECare, however, users explicitly interact with different physical devices, with specific capabilities. These devices can be grouped into three categories or layers: computing, communication and interface. Each of these layers has different devices and objectives and only with a close integration of all of them the system can be built.

In the computing layer one can have a multitude of devices which have computing capabilities and that can work independently from each other. In that sense we can consider small and common devices such as the mobile phone, mobile computers, watches, displays, photo machines, televisions or PDAs. In our modern homes, even window blinds, coffee machines or refrigerators have computing and networking capabilities. We can also consider environmental sensors such as temperature, luminosity and humidity sensors, smoke and flood sensors.

Concerning the communication layer, the main challenge comes from the heterogeneity of devices with communication capabilities and the different protocols and communication means that exist. A common home setting nowadays can have several networks, namely Ethernet, Wi-Fi, Bluetooth, power line, among others. There is a whole infrastructure of communicating devices that need to be compatible so that they can be integrated. The interface layer is also a very important part of the system since the devices through which the user interacts with the system will create the image that the user will have of the whole system. This means that this interaction must be very user friendly. In that sense, devices such as video cameras, microphones or touchscreens must be considered. In the opposite side we can enumerate the devices through which the system interacts with the environment, namely the actuators that allow controlling home appliances such as lights, heaters, air conditioning devices, among others.

It is a fact that nowadays a multitude of interesting devices exist that can improve our day-to-day living. However, these are very hard to integrate as each one has different characteristics. To address this problem there are already some defined standards but it is necessary that manufacturers follow them and that all of them follow a single one. The solution is then to find effective ways to integrate these different devices by creating a compatibility layer that not only allows the devices to communicate but also makes it easier to develop applications to make use of the services provided by the devices.

2 Nature of the Architecture

Given the already mentioned characteristics of the devices that may inhabit an intelligent environment and given the dynamic nature of the interaction paradigm (e.g. user

can interact with many eventually mobile devices, user can start or stop devices, user can incorporate new devices), we can state that an architecture for Ambient Intelligence should be dynamic, modular, expansible, flexible, scalable and compatible [14,15]. In order to implement such an architecture the two previously mentioned technologies will be considered together, joining the advantages of Service Oriented Architectures and Multi-agent Systems in order to build a dynamic environment that can incorporate a multitude of heterogeneous devices.

The Role of Jade. Agent-based technologies have been used for the most different purposes. Basically, this paradigm intends to solve problems through the interaction of simple entities called agents. Each agent has its knowledge about the world and its objectives to accomplish, which may be individual or collective objectives. Likewise, agents may either cooperate or compete in order to achieve their objectives [8]. One of the most interesting research trends in this field is in argumentation theory. In argumentation, agents debate, defend their beliefs and try to convince the other agents into believing the same they do in order to achieve their objectives [5]. Argumentation is also suited for solving conflicts that are usual in these environments. The most common example is a situation in which two agents have conflicting objectives (e.g. maintaining comfort versus saving energy). Besides argumentation, negotiation techniques [6] can and have also been considered to address these challenges.

One important issue here is the one of communication between the agents as these must respect a common standard that ensures that all the agents make use of the same ontology and message syntax. In that sense, the use of FIPA-ACL standard is highly useful. By doing so, many drawbacks concerning communication are solved and the compatibility of the architecture with external agents that respect the defined standard is assured, increasing the expansibility.

Given their features, agents will be used for all the high level decision making processes. It is thus necessary to define agents or groups of agents according to their roles in the architecture. This is, evidently, a task that is domain-dependent. An example focused on a domestic environment is described in section 4.

The Role of OSGi. In a common Aml architecture there is a group of components that must be connected. The objective of using an OSGi platform is to create a compatibility layer that can interconnect all the previously mentioned devices with software components such as databases or external service providers. The approach to be followed will consist in hiding each of these different components behind a different OSGi bundle. The key idea of this approach is to hide the singularities of each device and confine them to the respective bundle. This way, depending on the component being controlled, each bundle will be responsible for the interaction logic and the registration and request of necessary services for the correct execution of the component. Developers of other bundles thus do not need to know the specifications of a given component they intend to interact with. They need only to request the services being provided by the bundle that controls the component. In that sense, OSGi bundles can be seen as black boxes providing services: it does not matter what is inside the box as long as we know how to use the service.

3 Developing Intelligent Environments with OSGi and Jade

The roles of OSGi and Jade in the architecture have already been briefly depicted. We will now detail how these two tools can be used together in order to improve the process of developing intelligent environments. As stated before, the MAS is in charge of the high level decision processes, relying on tools like negotiation and argumentation to take globally optimum solutions. Additionally, OSGi is used to build a service layer that ensures the compatibility between all the different components of the architecture.

There are four main components in the architecture that allow a logical high level organization and a modular style of development: Jade container (the virtual location where the agents execute), Jade platform (may hold several Jade containers), OSGi bundle (able to provide and to use services from other bundles) and OSGi platform.

The process of creating an intelligent environment supporting the intended features is organized into a group of sequential and eventual iterative phases, as seen in figure 1.




Fig. 1. The three phases of the development process

Initial Specification. The process starts with the drawing of an initial specification. When defining this, some important factors must be considered such as the target environment of the system, its context of application, the devices that are likely to be used, among others. It is also important here to draw the first sketch of the functionalities of the system in terms of high level components. When drawing this first sketch it is usually useful to have some insights on the final architecture and the organization of the components, so that the first high level components and their organization do not differ significantly from the expected final version. However, as this approach is highly modular, the first sketch of the architecture may be defined without knowing how the final version will be: the intermediary versions can be easily reorganized by changing functionalities (e.g. moving bundles or agents between platforms). This significantly lifts the pressure on the development teams that have to define the architecture of the system without having a clear picture of how the final system will look like, solving one of the biggest challenges in the development of complex systems [7]. Note that in this first stage there is no need to declare which component is of which type, thus giving more freedom to make future changes to the system.

Development. When the process reaches the development phase, the first sketch of the architecture is implemented. This might consist in simply creating high level components according to the initial specification and optionally creating some services with no functionality, only to define the connections between the several components and the way that information will be shared. Note that in this case we define only the several OSGi platforms and the direction of the services that will be used, configuring only high level components.

Another important task that can be performed here is to implement simulated bundles instead of “real” ones. Let us call simulated bundle to a bundle that simulates the services it should provide, for example, a bundle that simulates the value of the temperature instead of reading that value from an actual sensor. This is important as it allows developing a prototype version of the system without having all the necessary devices, thus reducing the costs and allowing for more easily and rapidly developing the prototype in initial stages. This means that in an eventual intermediary phase in the process, the system might be constituted only by simulated bundles or a mixture of simulated and real bundles. With OSGi, the whole architecture can be built out of simulated bundles that are then gradually replaced with real ones. The only main concern here is that the bundles that are being replaced have the same name and the same services signature, i.e., for the remaining services they are the same, although the reality is that they provide the same services in different ways.

Assessment. Having implemented this first prototype version, the process moves on to the Assessment phase. In this phase, the system is tested in terms of its efficiency, robustness, usability and scalability. In initial versions, the interest in the assessment is to improve architecture-related parameters. Therefore, one of the important tests to perform is to invoke all the services in order to determine if they are correctly implemented and their signatures respected. It is also vital to evaluate the logical organization of the architecture. This may result in operations like dividing a complex bundle into simpler bundles, dividing an OSGi platform or Jade platform into several ones of the same type if they hold many components, group bundles or agents that are scattered in different platforms into one common platform, among others. In later phases of the development process, the assessment stage may receive as input the experience of interaction with target users. This might result in recommendations for changes suggested by experts in several fields, final users, etc. All these changes are then compiled and passed to a new iteration of the development stage for implementation. This process goes on until a satisfactory architecture is achieved.

By allowing to perform such tasks, this approach allows to develop intelligent environments by following several existing prototyping techniques, namely Throwaway, Evolutionary and Incremental prototyping.

4 Example Settings

4.1 VirtualECare

Having described the nature of the architecture that we achieve with this approach, let us now present an example of how the development process may occur. This example is a simplification of what was the development process of the VirtualECare project [15]. In this example we will assume that we have only temperature and luminosity sensors. According to the proposed process, the first step is to define an initial specification for the system, in terms of the high level roles and functionalities that one intends to implement. In this case we will consider four objectives: we want to monitor the environment in terms of environmental parameters, control the devices present in the environment, incorporate intelligent decision mechanisms and model the user’s preferences and needs.

In the first iteration of the development phase, we implement four bundles, each one representing one of the functionalities enumerated before (Figure 2). We also know that the bundle responsible for the decision mechanisms will contain agents so this will be a Jade-OSGi bundle.




Fig. 2. Result of iteration 1 with four bundles

This is a rather simple process, consisting in creating only the activators and manifests for each bundle. In the first assessment phase all the bundles are started in order to determine if they have been well defined. In this phase we also define which agents will be necessary in order to be implemented in the next iteration. In this phase we also decide that the Monitorization bundle should be divided into two bundles, one for interacting with real sensors and another one for simulating parameters for which we do not have sensors.




Fig. 3. Result of iteration 2 with five bundles and three agents

In the second iteration, we implement the decisions that have been taken in the previous one, arriving at an architecture as the one seen in Figure 3. Analyzing this system, we conclude that it makes more sense that the preferences and needs are embodied in two agents so that they also take part in the decision process. Therefore the decision is to move the code from a bundle to two agents. While developing the Real Sensors bundle, it was concluded that, due to the differences in the iteration logic of each type of sensor, this bundle should be divided. The same is decided for the Actuator bundle.

In further iterations similar operations can be performed, always without interfering with the components already present. One possible result is the architecture defined in Figure 4, considering a Fault Check bundle that restarts agents that have failed and a database.

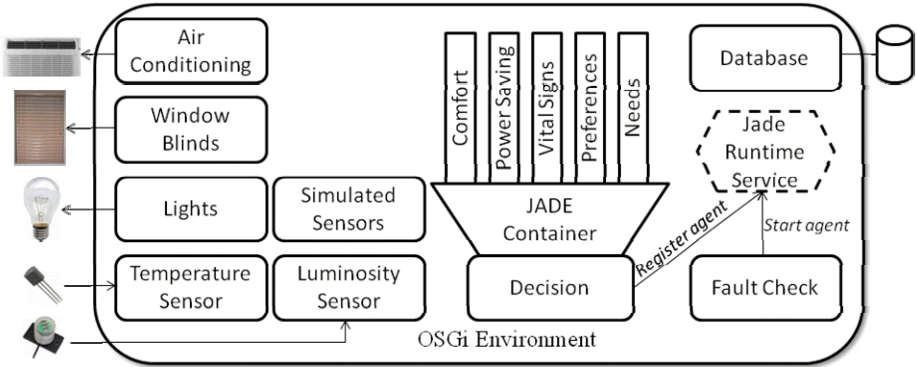


Fig. 4. A simple AmI architecture with 9 bundles, 5 agents and the external components considered

4.2 UMCourt

UMCourt [16] is the second project in which this approach is being successfully applied. This project aims at the creation of an Online Dispute Resolution platform that uses insights from the Artificial Intelligence field, namely Case-based Reasoning (CBR), to implement a group of services intended to help parties in dispute. These services include the estimation of the most likely outcomes of a dispute, the generation of strategies and solutions, a negotiation environment and general tools for legal documentation management.

Also here, the two technologies have been used with different purposes. Agents are used here in tasks that require significant context information. Examples are the negotiation and Case-based Reasoning modules. In these modules, FIPA-ACL messages are used not only to model information of the legal cases but also to model the control messages that define the negotiation protocol and the CBR process. An example of an ACL Message is shown below.

Example of an ACL message from agent Coordinator to agent Retriever requesting the cases similar to 1263491000923, assuming the default settings.

```
Sender : ( agent-identifier
  :name Coordinator@davide-desktop:1099/JADE
  :addresses (sequence http://davide-desktop:7778/acc ))
Conversation-ID : 1263492569251
Reply-To : Coordinator@davide-desktop:1099/JADE
Ontology : CBR_LABOUR
Content : RETRIEVE_SIMILAR DEFAULT 1263491000923
```

OSGi, however, is used differently. It is used to implement the low level tasks that lighten the execution of the agents. Namely, bundles implement behaviors for selecting cases according to given criteria, transparently accessing the database, loading and indexing cases, among others. Adopting this approach has the advantage of decreasing the complexity of the agents by removing these tasks from the scope of the agent. At the same time, it increases code reuse as it is common that different agents make use of some of these tasks that are, this way, encapsulated inside bundles. Figure 5 shows a simplified view of the UMCourt architecture.




Fig. 5. A simplification of the UMCourt architecture

5 Conclusions

Jade and OSGi can be used independently to create intelligent environments following two different approaches. However, their integrated use provides a much more powerful solution that can cut development time and requisites. Following the highly modular approach presented here, it is possible to develop these environments gradually, making changes as needed. Hence, it is easy to add new functionalities (in the shape of new bundles or agents), it is easy to rearrange the existing architecture (e.g. sub-dividing or integrating components) always without interfering with the components already present. Functionally, Jade ensures the development of agents and provides all the advantages of a complete messaging service, facilitating the implementation of complex negotiation and argumentation protocols.

OSGi, in the other hand, allows effective integration of very different components, creating a compatibility layer that exposes the functionalities of each component and hides the unnecessary complexity. Concluding, this approach allows to develop intelligent environments following the different prototyping techniques mentioned, fastening the whole process. Moreover, the advantages are present during the development phase and are reflected in the final architecture that is always ready to be improved with a new bundle or agent.

Acknowledgments. The work described in this paper is included in TIARAC - *Telematics and Artificial Intelligence in Alternative Conflict Resolution* Project (PTDC/JUR/71354/2006), which is a research project supported by FCT (Science & Technology Foundation), Portugal.

References

1. Wooldrige, M.: *An Introduction to Multiagent Systems*. John Wiley & Sons, Chichester (2002)
2. Perrey, R., Lycett, M.: Service-oriented architecture. In: *Proceedings of Applications and the Internet Workshops*, pp. 116–119, 27-31 (2003)
3. Bellifemine, F., Poggi, A., Rimassa, G.: *Developing Multi-agent Systems with JADE*. Springer, Heidelberg (2008)
4. Alliance, O.S.: *Osgi service platform*, release 3 (2003)
5. Amgoud, L., Parsons, S., Maudet, N.: Arguments, dialogue, and negotiation. In: Horn, W. (ed.) *Proc. 14th European Conf. on AI*, pp. 338–342. IOS Press, Berlin (2000)
6. Brito, L., Novais, P., Neves, J.: The logic behind negotiation: from pre-argument reasoning to argument-based negotiation. In: Plekhanova, V. (ed.) *Lit. - Intelligent Agent Software Engineering*, pp. 137–159. London Idea Group Publishing (2003)
7. Edwards, W.K., Grinter, R.E.: At Home with Ubiquitous Computing: Seven Challenges. In: Abowd, G.D., Brumitt, B., Shafer, S. (eds.) *UbiComp 2001*. LNCS, vol. 2201, pp. 256–272. Springer, Heidelberg (2001)
8. Olson, G.M., Malone, T.W., Smith, J.B. (eds.): *Coordination Theory and Collaboration Technology*. Erlbaum, Mahwah (2001)
9. Dix, A., Finley, J., Abowd, G., Beale, R.: *Human-computer interaction*, 3rd edn. Prentice-Hall, Inc., Upper Saddle River (2004)
10. FIPA: FIPA ACL Message Structure Specification, <http://www.fipa.org/specs/fipa00061/SC00061G.html> (accessed in January, 2010)
11. The Amigo Project: Amigo – Ambient Intelligence for the networked home environment. Short project description (2004)
12. Haigh, K., Kiff, L., Myers, J., Guralnik, V., Krichbaum, K., Phelps, J., Plocher, T., Toms, D.: *The Independent LifeStyle Assistant: Lessons Learned*. Honeywell Laboratories (2003)
13. Camarinha, L., Afsarmanesh, H.: Virtual Communities and Elderly Support. In: *Advances in Automation, Multimedia and Video Systems, and Modern Computer Science* (2001)
14. Carneiro, D., Novais, P., Costa, R., Gomes, P., Neves, J.: EMon: Embodied Monitorization. In: Tscheligi, M., et al. (eds.) *AmI 2009*. LNCS, vol. 5859, pp. 133–142. Springer, Heidelberg (2009)
15. Costa, R., Novais, P., Lima, L., Carneiro, D., Samico, D., Oliveira, J., Machado, J., Neves, J.: VirtualECare: Intelligent Assisted Living, in Electronic Healthcare. In: Weerasinghe, D. (ed.) *Social Informatics and Telecommunications Engineering*. Series Institute for Computer Sciences, pp. 138–144. Springer, Heidelberg (2009)
16. Carneiro, D., Novais, P., Andrade, F., Zeleznikow, J., Neves, J.: The Legal Precedent in Online Dispute Resolution. In: Guido (ed.) *Legal Knowledge and Information Systems*, pp. 47–52. IOS Press, Amsterdam (2009)
17. Spanoudakis, N., Moraitis, P.: An Ambient Intelligence Application Integrating Agent and Service-Oriented Technologies. In: *Research and Development in Intelligent Systems XXIV, Proc. 27th SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence (AI-2007)*, pp. 393–398 (2007)
18. De Maria, B.A., Silva, V.T., Lucena, C.J.P., Choren, R.: VisualAgent: A Software Development Environment for Multi-Agent Systems. In: *Proc. of the 19th Brazilian Symposium on Software Engineering (SBES 2005)*, Tool Track, Uberlândia, MG, Brazil (2005)
19. Model-driven Approach to Real-Time Embedded Systems development, MARTES (2007), <http://www.martes-itea.org>

PART VIII
AI Applications

Bee Hive at Work: Following a Developing Story on the Web

Pavol Navrat and Anna Bou Ezzeddine

Faculty of Informatics and Information Technologies, Slovak University of Technology
Bratislava, Slovakia
navrat@fiit.stuba.sk, ezzeddine@fiit.stuba.sk

Abstract. Problem solving methods, inspired biologically or by nature, are gaining in importance. Their potential for devising new, but also for increasing efficiency of some well-known methods for solving different kinds of problems, is becoming more and more obvious. The behaviour of social insects provides a lot of inspiration. We took inspiration from honey bees, in particular the way they behave in a hive. We proposed a modified model devised for searching parts of the web online. In a series of experiments, we attempted to follow a developing story with the aim to automatically find documents (web pages) bringing news on the story. This service, working as a beehive, delivered documents as expected.

Keywords: bee hive model, bee swarm intelligence, web search, on-line web crawling, developing story.

1 Introduction

When seeking new ways how to solve various classes of difficult problems, researchers more and more often seek inspiration from several animal species, observing as they behave in nature. At first sight, behaviour of particular animals may seem primitive, but when they act in a swarm, they can achieve states that can be interpreted as solutions to quite complicated problems. Moreover, the problem solving process is often quite effective. Processes such as collecting food by ants or foraging honey by honey bees are an instant inspiration. From observing them, simple rules can be inferred. Rules constitute basis for algorithms to solve various classes of interesting problems.

Among the animal species that have been attracting the interest of researchers, social insects occupy a dominant position. It is especially so because of the simple behaviour of an individual in relation to a complicated collective behaviour. A society (or swarm) of social insects is able to evolve means for collective problem solving. Complexity of such problems is well above the abilities of an individual. They can solve problems without central management and without predetermined fixed structures despite the massive internal whirr.

At first sight different kinds of social insects, such as ants, honey bees, wasps or termites, have different kinds behaviour. However if we study their behaviour and activities in greater detail, we can find many similar principles.

The behaviour of honey bees is a study subject of researchers in various disciplines, mainly biologists and social biologists in particular.

Their studies have proven to be extremely useful [1],[3],[4],[5],[24],[27]. They provide the underlying information on how honey bees behave in nature. We dare to identify one additional new dimension for problem solving: the web. It brings new opportunities, since it poses new kinds of problems. Nowadays we have a great wealth of information instantly retrievable. Therefore, we identify web problem solving as another excellent area of research.

The area is of course very broad. We have chosen one interesting information gathering scheme, which has become quite typical nowadays, when news appear on the web continuously. An event is reported. However, it induces other events and reactions, so in effect a story is developing. All this leaves traces on web pages of media portals reporting on it. We attempt to tackle this problem, moreover, we attempt to tackle it genuinely on-line. This is a true challenge for a bee hive.

The rest of the paper is organized as follows. In Section 2, we briefly describe the concept of self-organization in biological systems. We continue with briefly mentioning related works, then in Section 4 we describe the bee hive model and in Section 5 the web search it performs, concentrating on explaining the way a bee evaluates quality of a source. In Section 6, we elaborate how our bee hive tracks a story on the web. Section 7 concludes and hints future work.

2 Self-organization in Biological Systems

Intelligent behaviour of social insects begins through indirect communication among the individuals of the community known as the stigmergy principle.

The term of stigmergy was coined in the late fifties of the last century [7], to denote types of correlation among social insects that arose from for example, building mounds by termites and ants, or hives by bees. Stigmergy is a mechanism of spontaneous indirect coordination among the individuals in a community. It is a form of self-organisation, in which the community forms complex, declaredly intelligent, structures without the necessity of either planning or management. It supports effective coordination among very simple animals, that do not possess any memory or intelligence whatsoever.

This term was later transferred also to other fields including computer science. The term swarm intelligence was according to [3] used in artificial intelligence in 1989 by mobile robot systems.

Swarm intelligence describes collective behaviour in a decentralized self-organised system that can be natural or artificial. Swarm intelligence, in terms of artificial intelligence, represents the ability of systems composed of non-intelligent agents with limited individual abilities, to achieve general intelligent behaviour of the conjuncted swarm [26]. The agent is at this time known as an individual that is able to sense and observe its surroundings and consequently undertake some of the available actions. These actions also include changes in the surroundings that are occupied by the agent. Individuals do not possess information about how to solve the existing problem and

the intelligent behaviour appears as a consequence of actions of the swarm represented by these agents.

3 Related Work

Description of self organization of a honey bee system is discussed in [6]. Simple rules help the bees select the best source of nectar. Nectar is essential for their survival, so they fly out to their surroundings and look for it. Finding a food source compels the bee to fly back to the bee hive and bring the information about the source.

Models of bees collectively looking for collecting nectar are dissertated in [25]. Individually oriented simulation is constructed to simulate collective behaviour of bees in time of foraging. Each bee has to follow the same set of behavioural rules. Building up a simulation model, which would achieve similar results as those presented in [24], has been the main goal.

Camazine in [5] presents a description of dynamical interaction among bees in the process of carrying nectar from two sources using mathematical simulation.

The mathematical model of [5] has been an inspiration for a prototype of multi-agent recommending system, which was proposed in [10] and presented in [11] and [12]. They show the way of utilizing the bee hive metaphor. They drew from works of [23] and [22] applying different approaches to solving problems. Lorenzi worked in [9] on this question as well. By allowing more than two sources of food, they generalised the model of [5]. The foible of their model is that it presupposes the existence of as many bees as there are sources.

4 Bee Hive Model

In the process of devising the model, we inspired ourselves by a model introduced in [12]. Some limitations of the model were identified so the main aim was to eliminate them.

We considered as perhaps the main limitation of the model [12] that it included the initial assignment of one honey bee to one food source. Using this type of assignment for achieving valid results requires using as many honey bees – agents - as there are sources. Existence of a greater number of sources (e.g. web pages) incapacitates the model.

Parameters of the model are:

- total number of honey bees in the model N ,
- maximum dancing time for a particular food source MDT ,
- maximum time that the honey bee spends in the auditorium OT ,
- information whirr $NOICE$, accuracy in exchanging information between the dancing bee and the observing bee,
- error in evaluating the quality of the source ERR , its value is from the interval $\langle 0,1 \rangle$.

Our model consists of three parts: dispatch room, dance floor and auditorium.

Every honey bee starts from the dispatch room, chooses a source randomly, evaluates its quality q and returns back to the hive. There it decides whether it remains with the discovered source or moves to the auditorium to observe dancing bees.

If it decides to remain with the source, it will get to the dancing room with the probability q to promote the given source, or it will return to its source with the probability $(1-q)$.

The length of the dance time in the dancing room is proportional to the quality of the source, but it cannot be greater than MDT. After terminating its dance the bee returns back to its source.

If the bee decides to abandon its source, it enters to the auditorium, where it chooses to watch some of the dancing bees randomly. With the probability that is assigned by the ratio of the total number of bees dancing for the same source as the selected bee and the number of all bees in the dancing room, the observing bee would follow the recommended source. Otherwise it would remain in the dancing room and attempt to choose a different source. The stay of the bee in the auditorium must not overpass the value OT. The bee proceeds to the dispatch room after the expiration of the OT period.

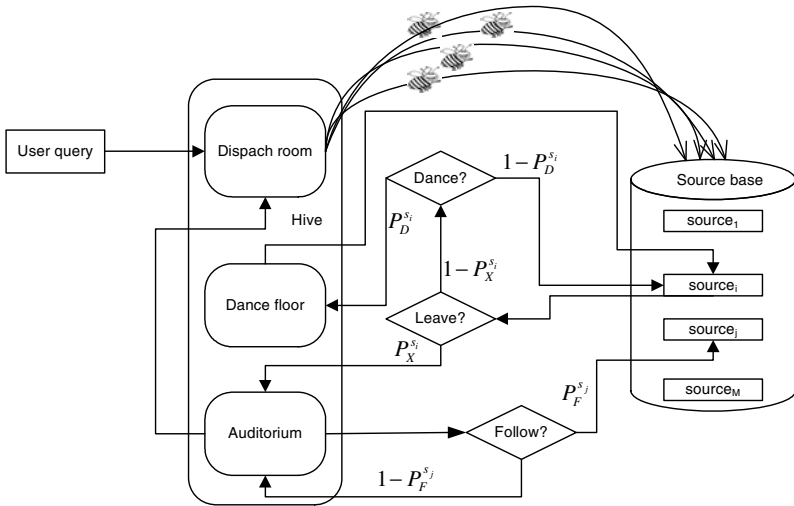


Fig. 1. Behaviour model of honey bee community according to [14]

The mechanism of our model can be found in [14], which contains a more detailed elaboration. In [15] and [17] we present additional experiments with the parameters of the model. The original model is universal and it can be used in various domain fields. It defines the behaviour of honey bees in the hive. In [16] we explicated the behaviour of bees when being outside the hive for Internet search.

In the process of searching on the Internet, the web page is considered to be the source, and the URL address is regarded as its location.

5 Web Search

In [8], different methods of honey bee swarm utilization in information systems are described in greater detail. As we mentioned, the priority of our research is web search, in particular on-line search inspired by the behaviour of a swarm.

Nowadays majority of locators work on off-line database, which includes indexed web pages from some parts of the Internet. The contents of these web pages are updated constantly and new web pages are accumulated.

The volume of Internet is enormous therefore it is impossible to search through all web pages. This is why it is necessary to focus the web search only on the most promising way to search. Every user possesses some preferences when searching the web. They can be expressed by selecting a cluster of preferred web pages. Web search begins from a selected cluster of web pages. So a web page represents a source and the accessibility of another source represents hypertext reference to another web page (source).

One of the most important functions of the model is detecting the quality of the found web page. Quality of the source is from interval $\langle 0,1 \rangle$. The resultant quality is made of partial qualities. In particular cases it is possible to select partial qualities and their importance in the process of determining the whole quality of the source.

Partial qualities:

1. Quality of the range between two web pages – it is defined as a number of different domains that a bee visits in order to get from one web page to another.

2. Incidence number quality – quality based on the number of word incidence. It is computed in agreement with the formula:

$$q_{count} = \frac{-1}{2 \left(n + \frac{1}{2Q_{COUNT}} \right)} + Q_{COUNT} \quad (1)$$

where n is the number of incidence of the given word and Q_{COUNT} is the maximal value defined for this partial quality.

3. Quality of incidence in the caption – in this case it is important to trace the hierarchical level of word incidence in captions. Numbers from 0 to 6 are associated to every level. Calculation of this partial quality is given by formula:

$$q_{header} = Q_{HEADER} - h * \frac{Q_{HEADER}}{HEADER_{MAX} + 1} \quad (2)$$

where h represents the minimum of all caption values that include the searched word; Q_{HEADER} is the maximum value of this partial quality; $HEADER_{MAX}$ is the maximal depth of the caption hierarchy that the calculation takes into account.

4. Readability index by Flesch – this test evaluates the text according to a 100 degree scale. The higher the score is, the more comprehensible the text. The formula for Flesch's readability index is:

$$206.835 - (1.015 \times ASL) - (84.6 \times ASW) \quad (3)$$

where ASL = Average Sentence Length (the number of words divided by the number of sentences), ASW = Average number of Syllable per Word (the number of syllables divided by the number of words). In [18] we have shown that using the combination of key words and readability index makes the process of obtaining promising results possible, especially in professional text search.

6 Web Story

When an inquisitive person (IP), a curious fellow, takes interest about a particular event, which is being monitored by media, he or she wants to read news reporting on that event. Often, an event is not an isolated one. There is some follow up, some reactions, induced events. A story develops. IP wants to follow the story. Sometimes, the news itself is the most important information. But usually, reports are soon followed by commentaries, in-depth analyses, blogs or discussions. IP may be interested in some of this as well. It is fair to assume all this data is available on the web. To access just the right documents requires searching, filtering, sorting and other data processing, which is time-consuming. Moreover, since the story develops in time, it would require almost continuous monitoring of the story as it progresses.

Thus our aim was to devise a method implementable on a personal computer that would be capable of supplying documents related to the developing story as they are emerging on the web for several hours or days.

It seems that any such method must include searching the web. After collecting a set of related documents, it is necessary to filter and cluster the data set. News portals contain mainly articles, annotations, discussions, blogs and symposia. Articles are important from the aspect of information tracing. Annotations, blogs and symposia are elaborated in order to gain opinions on particular events, eventually on offered products [21].

Information that was found and sorted was classified according to the publication date. We attempted to approach this problem using the bee hive model [19]. Inspired by the work of [13] we used the concept of energy from agent InfoSpider.

Source quality is calculated as the sum of partial qualities (incidence number quality, caption incidence quality, Flesch's readability index).

Biological observations show [1] that bees are specialized as either scouts or recruits in the process of collecting food. Scouts look for food independently; recruits gather food with the use of information from other bees.

In our model a bee in the auditorium can be considered as a free bee. If it chooses a bee from the dancing room and will follow that bee's source, it becomes a recruit.

If it does not choose a bee and passes to the dispatch room, it becomes a scout. Scouts and recruits act similarly when being inside the hive, but their behaviour changes as soon as they leave the hive. After finding a source with non zero quality, the scout returns to the hive to promote its source.

Parameters of our model of the hive are given in Table 1, and of the quality calculation in Table 2.

Table 1. Parameters of the hive

Number of bees	30
MDT maximal dancing time	7 iteration
OT maximal time in auditorium	4 iteration

Table 2. Parameters of quality calculation

Initial energy of a bee	1
Energy increase	source quality
Energy decrease because of transition to new source	0.05
Maximum quality of incidence number	0.7
Maximum caption incidence quality	0.15
Maximum caption number	3
Flesch's readability index	0.15

Experiments

We used the event of earthquake in Haiti that has been widely monitored by media, at the time as a developing story. In our experiment, three news portals represent the start web pages: www.pravda.sk, www.sme.sk and www.ta3.com. The key words we looked for were: earthquake and Haiti.

The process of following the story was divided into three parts.

The first part is represented by an experiment that took place from 13 January 2010 10:00 am to 14 January 2010 4:00 pm 9327 web pages were found, out of which 1066 were of non zero quality. Web pages of non zero quality were divided into 5 classes:

Table 3. Division of related pages into classes, 13-14 January 2010

Informative page	493
List of articles	348
Discussions	73
Blogs	65
Graphic content (pictures, videos)	87

Let us assume the IP seeks new raw information on this story, that has been developing so dramatically. For the IP we have in mind, discussions and blogs are irrelevant because they represent only reactions to the event.

The list of articles has no informative value, but is important for page discovery.

Because of sufficient number of informative web pages, we selected only web pages with quality higher than 60 percent. (There were 217 such web pages). These web pages were classified according to the published date extracted from the page.

The most frequently used words on those 217 informative web pages were: disaster, tragedy, victims, UN, chaos, help.

The second part of the experiment took place on 16 and 17 January 2010 always at the same time: from 8:00 am to 5:00 pm 11439 web pages were found, 1193 with non zero quality. Web pages of non zero quality were divided into these classes:

Table 4. Division of related pages into classes, 16-17 January 2010

Informative page	552
List of articles	385
Discussions	83
Blogs	72
Graphic content (pictures, videos)	101

298 informative web pages with the quality higher than 60 percent were found.

The most frequently used words which occurred on the web pages after 16 January 2010 8:00 am were: cadavers, indigence, looting, despair, water, help.

The last part of the experiment took place from 18 to 21 January 2010, always at the time from 6:00 pm to 11:00 pm and 12576 web pages were found, out of which 1271 of non zero quality. Web pages of non zero quality were divided into these classes:

Table 5. Division of related pages into classes, 18-21 January 2010

Informative page	598
List of articles	399
Discussions	87
Blogs	78
Graphic content (pictures, videos)	109

327 informative web pages possessing a quality of above 60 percent were found.

The most frequently used words on the web pages after 18 January 2010 6:00 pm were: water, collections, help, charity, putrefaction, physicians.

Due to the results of the experiments we can deduce, that our designed system is able to follow a story as it develops. When repeating the experiment twice or three times, the algorithm found almost all the web pages that were marked as relevant by the previous algorithm run, hence supporting a hypothesis that our method based on the modified bee hive model is quite robust.

7 Conclusion and Future Work

The bee hive model has been used in our previous works mainly in web search and in function optimisation [20]. In this work we present a modified bee hive model and employ it in a system for tracking a developing story.

From the experiments we can conclude the following:

- the system is able to look for web pages and evaluate the quality of the found web pages automatically,
- it can collect relevant pages,
- it can reconstruct the story backwards in time,

- it can monitor the story that is developed during the search,
- it provides statistical results about the searching process,
- by means of this system we can obtain the most frequently used words on time distinguishable web pages.

According to the acquired statistical results we can conclude that after the initial wide interest of the population about the mentioned disaster (the first part of the experiment), the interest stabilizes gradually. Number of discussions and blogs did not change largely through the whole duration of the experiment.

Due to the obtained key words we can deduce the changing contents of the relevant articles.

In the forthcoming period we plan to process the acquired discussions and blogs and to evaluate the opinions of discussants on possible ways of help in disasters.

Acknowledgement

This work was partially supported by the grants VEGA 1/0508/09, KEGA 345-032STU-4/2010 and it is the partial result of the Research & Development Operational Programme for the project Support of Center of Excellence for Smart Technologies, Systems and Services, ITMS 26240120029, co-funded by the ERDF.

References

1. Biesmeijer, J.C., de Vries, H.: Exploration and exploitation of food sources by social insect colonies: A revision of the scout-recruit concept. *Behavioral Ecology and Sociobiology* 49(2-3), 89–99 (2001)
2. Bonabeau, E., Theraulaz, G., Deneubourg, J.L., Aron, S., Camazine, S.: Self-organization in social insects. *Trends in Ecology and Evolution* 12, 188–193 (1997)
3. Bonabeau, E., Thraulaz, G., Dorigo, M.: *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, Oxford (1999)
4. Bullock, H.E., Dey, P., Reilly, K.D.: A “Bee Hive” Model For Heterogeneous Knowledge. In: *Expert Systems*, pp. 417–417. ACM, New York (1986)
5. Camazine, S., Sneyd, J.: A model of collective nectar source selection by honey bees: Self-organization through simple rules. *Journal of Theoretical Biology* 149(4), 547–571 (1991)
6. Camazine, S., Deneubourg, J.L., Franks, N.R., et al.: *Self-organization in Biological Systems*. Princeton University Press, Princeton (2003)
7. Grasse, P.-P.: La Reconstruction du nid et les Coordinations Individuelles chez *Bellisotermes Natalensis* et *Cubitermes* sp. la Theorie de la Stigmergie: Essai d’interpretation du Comportement des Termites Constructeurs. *Insectes Sociaux* 6, 41–80 (1959)
8. Karaboga, D., Akay, B.: A survey: algorithms simulating bee swarm intelligence. *ArtifIntellRev*, SpringerScience Busines Media B 31(1-4), 61–85 (2009)
9. Lorenzi, F., Ricci, F.: Case-based reasoning and recommender systems. Technical report, IRST (2004)
10. Lorenzi, F., Sherer Dos Santos, D., Bazzan, A.L.C.: Case-based Recommender Systems: a unifying view. In: *LNCS/LNAI State-of-the-Art Survey book on Intelligent Techniques in Web Personalization* (2005)

11. Lorenzi, F., Sherer Dos Santos, D., Bazzan, A.L.C.: Case-based recommender systems inspired by social insects. In: XXV Congresso da Sociedade Brasileira de Computacao, July 22-29 (2005)
12. Lorenzi, F., Sherer Dos Santos, D., Bazzan, A.L.C.: Negotiation for task allocation among agents in case-based recommender systems: a swarm-intelligence approach. In: IJCAI 2005 Workshop on Multi-Agent Information Retrieval and Recommender Systems, Edinburgh, pp. 23–27 (2005)
13. Menczer, F., Monge, A.E.: Scalable Web Search by Adaptive Online Agents: An InfoSpiders Case Study. In: Intelligent Information Agents. Springer, Heidelberg (1999)
14. Navrat, P., Kováčik, M.: Web Search Engine as a Bee Hive. In: 2006 IEEE/WIC/ACM International Conference on Web Intelligence, pp. 694–701. IEEE Computer Society, Los Alamitos (2006)
15. Navrat, P., Kovacik, M., Bou Ezzeddine, A., Rozinajova, V.: Information retrieval using honey bees. In: Znalosti 2007, Ostrava, pp. 63–74 (2007)
16. Navrat, P., Jastrzemska, L., Jelínek, T., Bou Ezzeddine, A., Rozinajová, V.: Exploring Social Behaviour of Honey Bees Searching on the Web. In: Li, Y., Raghavan, V.V., Broder, A., Ho, H. (eds.) 2007 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology (Workshops), Silicon Valley, USA, November 2-5, pp. 21–25. IEEE Computer Society, Los Alamitos (2007)
17. Navrat, P., Kováčik, M., Bou Ezzeddine, A., Rozinajová, V.: Web search engine working as a bee hive. In: Web Intelligence and Agent Systems: An International Journal, vol. 6, pp. 441–452. IOS Press, Amsterdam (2008)
18. Navrat, P., Taraba, T., Bou Ezzeddine, A., Chudá, D.: Context Search Enhanced by Readability Index. In: Artificial Intelligence in Theory and Practice II. IFIP 20th World Computer Congress. IFIP Series, vol. 276, pp. 373–382. Springer Science+Business Media, LLC, New York (2008)
19. Navrat, P., Jastrzemska, L., Jelinek, T.: Bee Hive At Work: Story Tracking Case Study, Milan, Italy, September 15-18, pp. 117–120. IEEE Computer Society, Los Alamitos (2009)
20. Navrat, P., Jelinek, T., Jastrzemska, L.: Bee Hive At Work: A Problem Solving, Optimizing Mechanism. In: World Congress on Nature and Biologically Inspired Computing (NaBIC 2009) & 8th International Conference on Computer Information Systems and Industrial Management Applications (CISIM2009), December 9-11. PSG College of Technology, Coimbatore (2009)
21. Navrat, P., Bou Ezzeddine, A., Slizik, L.: Mining Overall Sentiment in Large Sets of Opinions, AWIC 2009. In: Advances in Intelligent Web Mastering - 2. AISC, vol. 67, pp. 167–173. Springer, Heidelberg (2010)
22. Resnick, P., et al.: An open architecture for collaborative filtering of netnews. In: Proceedings ACM Conference on Computer-Supported Cooperative Work, pp. 175–186 (1994)
23. Schafer, J.B., Konstan, J.A., Riedl, J.: E-commerce recommendation applications. *Data Mining and Knowledge Discovery* 5(1/2), 115–153 (2001)
24. Selley, D., Camazine, S., Sneyd, J.: Collective decision-making in honey bees: How colonies choose nectar sources. *Behavioral Ecology and Sociobiology* 28, 277–290 (1991)
25. Vries, H., Biesmeijer, J.C.: Modelling collective foraging by means of individual behaviour rules in honey-bees. *Behav. Ecol. Sociobiol.* 44, 109–124 (1998)
26. White, T., Paturek, B.: Towards Multi-Swarm Problem Solving in Networks. In: Multi Agent Systems - International Conference on Volume (1998)
27. Zhang, S., Schwarz, S., Pahl, M., Zhu, H., Tautz, J.: Honeybee Memory: A Honeybee Knows What To Do And When. *Journal of Experimental Biology* 209, 4420–4428 (2006)

Computer-Aided Estimation for the Risk of Development of Gastric Cancer by Image Processing

Yoshihiro Sasaki¹, Ryukichi Hada¹, Tetsuro Yoshimura², Norihiro Hanabata²,
Tatsuya Mikami², and Shinsaku Fukuda²

¹ Department of Medical Informatics and Gastroenterology

² Hirosaki University Hospital, Honcho-53, Hirosaki 036-8563, Japan
gahiro@cc.hirosaki-u.ac.jp

Abstract. The aim of this study was to establish a computer-aided estimating system for determining the risk of development of gastric cancer, achieved by image processing on an ordinary endoscopic picture. Digital endoscopic pictures of the background gastric mucosa in 26 *Helicobacter pylori* (*H. pylori*) positive patients with early intestinal type gastric cancer and age-gender-matched *H. pylori* positive subjects without cancer were used. The pictures were processed for 15 pictorial parameters. Out of the 15 pictorial parameters, 3 parameters were found to characterize the background gastric mucosa with gastric cancer against that without. Based on the Bayes decision theory, the computer-aided estimating system has been established. Sensitivity, specificity, positive predictive value and negative predictive value of the Bayes classifier were found to be 0.64, 0.64, 0.65 and 0.63, respectively. This method may permit an effective selection of the high risk population of gastric cancer needing follow-up endoscopy.

Keywords: *Helicobacter pylori*; Gastric cancer; Endoscopy; Image processing; Computer-aided diagnosis.

1 Introduction

Most studies have found *Helicobacter pylori* (*H. pylori*) to be closely associated with the development of gastric cancer [1], [2], [3], [4], [5]. Among the patients with *H. pylori* infection, histologic findings at higher risk for gastric cancer have included severe gastric atrophy, corpus-predominant gastritis and intestinal metaplasia [1]. In Japan, early detection of gastric cancer has been a crucial clinical problem for the mortality control because the incidence of gastric cancer has been about 9-fold higher compared with that in the USA [6]. When gastric atrophy advances, the submucosal vasculature becomes to varying extent traceable on endoscopy. Although it is quite easy for endoscopists to recognize and describe the endoscopic features of atrophy as a risk factor for gastric cancer in empiric terms, pictorial information has been extremely refractory to precise definition and to analysis by computers.

Recently, computer aided diagnosis (CAD) methods have been developed for image interpretation in large-scale breast or lung cancer screening studies. CAD methods are

important for: (a) improving the sensitivity of cancer detection, (b) reducing interobserver variation in image interpretation, (c) improving efficiency of screening by identifying suspect lesions [7]. We have developed a quantitative method to characterize endoscopic features in a processed endoscopic picture [8]. The aims of this study were to characterize background gastric mucosa at higher risk for gastric cancer and to establish a CAD system for determining the risk of development of gastric cancer in a population with positive *H. pylori*.

For this purpose, digital endoscopic pictures of the background gastric mucosa in subjects with or without gastric cancer were processed for 15 pictorial parameters characterizing endoscopic features. Out of the 15 pictorial parameters, 3 parameters were found to characterize the background gastric mucosa with gastric cancer against that without. Sensitivity, specificity, positive predictive value and negative predictive value of this CAD system were found to be 0.64, 0.64, 0.65 and 0.63, respectively.

2 Materials and Methods

In this section, we present an image processing method to extract endoscopic features of the background gastric mucosa in subjects with or without gastric cancer for CAD system determining the risk of development of gastric cancer, and evaluate the performance of this system.

2.1 Subjects

A total of 2584 subjects underwent gastroscopy between January 2002 and February 2003 at Hirosaki University Hospital. Out of the subjects, we enrolled 26 consecutive patients with early well differentiated gastric carcinoma (19 males and 7 females; mean age, 65.1 years; range, 34 to 77). Out of 185 consecutive patients with non-ulcer dyspepsia, 26 age-gender-matched patients served as controls (19 males and 7 females; mean age, 65.1 years; range, 35 to 77). A total of 52 subjects were enrolled in this study. All subjects gave written informed consent. The study protocol was approved by the Ethics Committees of Hirosaki University Hospital and was reviewed annually.

2.2 Endoscopy

Endoscopy was performed with local anaesthesia (lidocaine). Endoscopic pictures were taken under the neutral chromatic parameters of an endoscopy system (EVIS 240 for a control unit and XK240 for an endoscope; Olympus, Tokyo, Japan), and saved in a digital filing system (EVIS-F; Olympus, Tokyo, Japan). Two biopsy specimens were taken from the greater curvature of the antrum and the middle body of the stomach for a rapid urease test (Helicocheck, Ohtsuka Pharmaceutical, Tokyo, Japan). In patients with suspected gastric cancer, additional biopsies were performed for histological diagnosis.

Table 1. Textural features

Number	Name	Characteristic
f1	Angular second moment	Homogeneity
f2	Contrast	Contrast
f3	Correlation	Linearity
f4	Sum of square	Heterogeneity
f5	Inverse difference moment	*
f6	Sum average	*
f7	Sum variance	*
f8	Sum entropy	*
f9	Entropy	*
f10	Difference variance	*
f11	Difference entropy	*

*See detail in the reference 10.

2.3 Pictorial Features of the Endoscopic Images

The endoscopic picture of the lesser curvature of the gastric body distended by a sufficient amount of air (Figure 1) was used. In electronic endoscopy, red (wavelength > 600 nm), green (500 to 600 nm) and blue light (<500 nm) are sequentially irradiated on the surface of the mucosa. Part of the lights incident on the mucosa, after suffering scattering and absorption, is re-emitted in a nearly diffuse state [9]. A single endoscopic picture is composed of the red, green and blue diffuse reflectance images. The color of any object in the digitized picture can be quantified by the intensity of red (R), green (G) and blue (B) reflectance. As hemoglobin is the predominant chromophore in the gastrointestinal mucosa, a colorant endoscopic picture is equivalent to a grey scale picture with index of hemoglobin ($IHB=32\log_2[R/G]$) assigned to each pixel (Figure 2) [8]. The grey scale picture with blackout or halation excluded was processed for 4 basic statistical parameters (mean IHB, SD of IHB, skewness of IHB, kurtosis of IHB) and 11 textural features that are related with spatial arrangements of the grey scale intensities (Table 1) [10]. Criteria of parameters for a reliable CAD system can be stated as follows: 1) Significant difference is demonstrated between the two subject groups and 2) the parameters are mutually independent to the other ones in a set of parameters with a significant difference. Out of the 15 pictorial parameters, SD of IHB, skewness of IHB and f3 (correlation feature) have been found to meet the criteria, thereby uniquely characterizing the background mucosa with cancer against that without cancer.

2.4 Bayes Classifier for the Computer Aided Diagnosis System

A computer-aided estimating system for determining the risk of development of gastric cancer was established on the basis of Bayes decision theory [11]. Let ω_1 and ω_2 be the state of the background gastric mucosa with or without gastric cancer, respectively. Let the feature vector \mathbf{x} be a 3-component column vector, where the component x_1 , x_2 and x_3

designate SD of IHB, skewness of IHB and f3, respectively, and let $p(\mathbf{x}|\omega_i)$ be the state-conditional probability density function for \mathbf{x} . Then, the likelihood for the background gastric mucosa with gastric cancer $f(\mathbf{x})$ can be computed from $p(\mathbf{x}|\omega_i)$ by Bayes rule:

$$f(\mathbf{x}) = p(\mathbf{x}|\omega_1) / (p(\mathbf{x}|\omega_1) + p(\mathbf{x}|\omega_2)).$$

$p(\mathbf{x}|\omega_i)$ ($i=1, 2$) is calculated as

$$p(\mathbf{x}|\omega_i) = \exp[-1/2(\mathbf{x}-\boldsymbol{\mu})^t \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})] / [(2\pi)^2 |\boldsymbol{\Sigma}|^{1/2}],$$

where $\boldsymbol{\mu}$ is the 3-component mean vector in ω_i , $\boldsymbol{\Sigma}$ is the 3 by 3 covariance matrix, $(\mathbf{x}-\boldsymbol{\mu})^t$ is the transpose of $\mathbf{x}-\boldsymbol{\mu}$, $\boldsymbol{\Sigma}^{-1}$ is the inverse of $\boldsymbol{\Sigma}$, and $|\boldsymbol{\Sigma}|$ is the determinant of $\boldsymbol{\Sigma}$. When a given gastric mucosa with positive *H. pylori* has a feature vector \mathbf{x} , the risk for the development of gastric cancer is given by $f(\mathbf{x})$, which increases from 0 to 1 as the risk increases.

2.5 Assessment for the Performance of CAD System

Five-fold cross-validation method was employed. A total of 26 data samples were divided into 5 test samples and 21 training samples. The training samples were utilized to obtain $f(\mathbf{x})$. The test samples were used to evaluate the prediction performance (including sensitivity, specificity, positive predictive value and negative predictive value) of $f(\mathbf{x})$ at varying cut-off values from 0.1 to 0.9. The averaged sensitivity and specificity over a total of 5 test sets were plotted against the cut-off values (Figure 3). The optimal cut-off, which can be defined as the crossing point of the two curves, was found to be 0.42.

2.6 Statistical Analysis

Pictorial parameters were expressed as mean \pm SD. Statistical significance was evaluated using Student's or Welch's unpaired t-test. The individual p-values were corrected for multiplicity effects due to multiple testing [12]. Differences between means were considered significant if the corrected-p value < 0.05 .

3 Experimental Results

The pictorial parameters in the background mucosa with or without gastric cancer were shown in Table 2. The background mucosa with cancer had a significantly larger SD of IHB (6.8 ± 1.1 , $p=0.013$), skewness of IHB (-0.17 ± 0.63 , $p=0.021$) and f3 (4.6 ± 1.0 , $p=0.015$) than that without cancer (6.0 ± 0.9 for SD of IHB; -0.69 ± 0.51 for skewness of IHB; 3.7 ± 1.2 for f3). In a set of the three parameters, anyone was independent to the other ones. The other pictorial parameters f1, f4, f7, f8 and f9 have also significantly differed between the two categories. But, they have not been independent to anyone of the three parameters (SD of IHB, skewness of IHB and f3).

An example of the background mucosa with gastric cancer was 6.4 in SD of IHB, -0.5 in skewness of IHB, 5.0 in f3 and estimated at high risk ($f(x)=0.85$) (Figure 4), while that without cancer was 4.5 in SD of IHB, 0.1 in skewness of IHB, 3.6 in f3 and estimated at low risk ($f(x)=0.24$) (Figure 5). Under the cut-off value of 0.42, sensitivity, specificity, positive predictive value and negative predictive value in estimating whether at high risk or not were found to be 0.64, 0.64, 0.65 and 0.63, respectively.






Fig. 1. The endoscopic picture of the lesser curvature of the gastric body distended by a sufficient amount of air

Fig. 2. The grey scale picture composed of pixel arrays with the index of hemoglobin assigned to each pixel for pictorial parameters. The black (0 intensity) area designates blackout or halation area excluded from analysis.

Fig. 3. Sensitivity (open circle) and Specificity (filled circle) against cut-off values





Fig. 4. An example of the background mucosa with gastric cancer (SD of IHB=6.4, skewness of IHB= -0.5, f3=5.0 and estimated at high risk or $f(x)=0.85$)

Fig. 5. An example of the background mucosa without gastric cancer (SD of IHB=4.5, skewness of IHB=0.1, f3=3.6 and estimated at low risk or $f(x)=0.24$)

Table 2. Basic statistical parameters and textural features in the background mucosa with or without gastric cancer

	With cancer	Without cancer	p value*
mean*	56.4±7.6	58.6±8.5	0.350
SD*	6.8±1.1	6.0±0.9	0.013
Skewness	-0.2±0.6	-0.7±0.5	0.021
Kurtosis	5.9±3.4	7.3±3.8	0.263
f1	1.2±0.3	1.4±0.4	0.026
f2	2.0±0.6	1.8±0.4	0.270
f3	4.6±1.0	3.7±1.2	0.015
f4	1.9±0.6	1.4±0.4	0.014
f5	4.7±0.5	4.8±0.4	0.337
f6	2.6±0.4	2.7±0.5	0.339
f7	2.2±0.8	1.6±0.5	0.015
f8	2.6±0.2	2.4±0.2	0.028
f9	3.9±0.3	3.7±0.2	0.025
f10	1.7±0.5	1.6±0.4	0.290
f11	3.1±0.3	3.0±0.2	0.316

mean*, the mean of gray scale intensity of index of haemoglobin; SD*, the standard deviation of gray scale intensity of index of haemoglobin; p value* (corrected according to Benjamini & Hochberg method, see detail in the reference [12]), the background gastric mucosa with cancer vs. without cancer; the parameters were expressed as mean±SD.

4 Discussion

We have developed a reliable method to quantify endoscopic features in a processed grey scale picture with IHB assigned, which has led to a computer-aided grading system for endoscopic severity in ulcerative colitis [7]. In the present study, the background gastric mucosa with gastric cancer has been demonstrated to be uniquely characterized against that without cancer by three pictorial parameters. With the use of the parameters, a computer-aided estimation system for the risk of development of gastric cancer has been established.

In Japan, endoscopy is performed frequently for gastric cancer screening even in subjects without any symptom because of the high incidence of the disease [6]. Clinical interest has been particularly focussed on quantitative endoscopic features for the risk of development of gastric cancer among subjects with *H. pylori* infection to establish a screening or follow-up protocol with a higher cost-performance.

Among the patients with *H. pylori* infection, histologic findings at higher risk for gastric cancer have included severe gastric atrophy, corpus-predominant gastritis and intestinal metaplasia [1]. Out of the three histologic findings, gastric atrophy is characterized as the presence of characteristic patterns or mesh-like or river-like submucosal vasculature. The degree of mucosal atrophy (the clarity or complexity of the patterns) or the area of the atrophic mucosa have been considered to correlate with the three pictorial parameters (SD of IHB, skewness of IHB and f3).

Pictorial information has been extremely refractory to precise definition and to analysis by computers, although it is quite easy for human observers to recognize and describe in empirical terms. With the advent of the electronic endoscope, digital filing system and image processor, it has become possible to quantify any element composing a digitized endoscopic picture by mathematical processes. We have noticed that endoscopic features for diagnosis are concerned with spatial arrangements of IHB. The spatial arrangements expressed by empirical terms as fine, coarse, smooth, rippled, irregular or lineated are quantified by the basic statistical parameters or textural features. The parameters of diagnostic value may vary with diseases or organs studied and are not necessarily independent to each other [10]. For a reliable CAD system, it is essential to determine the parameters of a diagnostic value on the basis of above stated criteria.

In gastrointestinal endoscopy, image interpretation is inherently associated with interobserver variation. We have previously established CAD for grading endoscopic severity in patients with ulcerative colitis [8] or CAD for detecting flat type early esophageal cancer [13]. In the present study, it has been demonstrated that CAD methods can be utilized as a means for an effective screening of gastric cancer by identifying the high risk population needing follow-up endoscopy. The challenge is now to construct a database of the endoscopic images for the optimization of CAD software, leading to the creation of a more scientific basis for evaluation and comparison of image processing methods [7].

References

1. Uemura, N., Okamoto, S., Yamamoto, S., Matsumura, N., Yamaguchi, S., Yamakido, M., Taniyama, K., Sasaki, N., Schlemper, R.J.: *Helicobacter pylori* infection and the development of gastric cancer. *N. Eng. J. Med.* 354, 784–789 (2001)
2. Forman, D., Newel, I.D.G., Fullerton, F., Yarnell, J.W., Stacey, A.R., Wald, N., Sitas, F.: Association between infection with *Helicobacter pylori* and risk of gastric cancer: evidence from a prospective investigation. *Brit. Med. J.* 302, 1302–1305 (1991)
3. Huang, J.Q., Sridhar, S., Chen, Y., Hunt, R.H.: Meta-analysis of the relationship between *Helicobacter pylori* seropositivity and gastric cancer. *Gastroenterol* 114, 1169–1179 (1998)
4. Danesh, J.: *Helicobacter pylori* infection and gastric cancer: systemic review of the epidemiological studies. *Aliment. Pharm. Therap.* 13, 851–856 (1999)
5. Tabata, H., Fuchigami, T., Kobayashi, H., Sakai, Y., Nakanishi, M., Tomioka, K., Nakamura, S., Fujishima, M.: *Helicobacter pylori* and mucosal atrophy in patients with gastric cancer: a special study regarding the methods for detecting *Helicobacter pylori*. *Dig. Dis. Sci.* 44, 2027–2034 (1999)
6. Lambert, R., Guilloux, A., Oshima, A., Pompe-Kirn, V., Bray, F., Parkin, M., Ajiki, W., Tsukuma, H.: Incidence and mortality from stomach cancer in Japan, Slovenia and the USA. *Int. J. Cancer* 97, 811–818 (2002)
7. New NCI Initiatives in Computer Aided Diagnosis,
<http://imaging.cancer.gov/reportsandpublications/publications/NewNCIInitiativesinComputerAidedDiagnosis>

8. Sasaki, Y., Hada, R., Munakata, A.: Computer-aided grading system for endoscopic severity in patients with ulcerative colitis. *Dig. Endosc.* 15, 206–209 (2003)
9. Wyszecki, G., Stiles, W.S.: *Color science: concepts and methods, quantitative data and formulae*. John Wiley & Sons, New York (1982)
10. Haralick, R.M., Shanmugam, K., Dinstein, I.: Textural features for image classification. *IEEE Trans. Syst. Man. Cybern.* 6, 610–621 (1973)
11. Duda, R.O., Hart, P.E.: *Bayes Decision Theory: In Pattern Classification and Scene Analysis*. John Wiley & Sons, New York (1973)
12. Benjamini, Y., Hochberg, Y.: Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Statist. Soc. Ser. B* 57, 298–300 (1995)
13. Mikami, T., Sasaki, Y., Fukuda, S., Hada, R., Munakata, A.: Computer-aided diagnosis of superficial type early oesophageal cancer by image processing on ordinary endoscopic pictures. *Digestive Disease Week, Gastrointest. Endosc.* 59, 258 (2004)

Intelligent Hybrid Architecture for Tourism Services

Ana Almeida^{1,2}, Bruno Coelho¹, and Constantino Martins^{1,2}

¹ GECAD - Knowledge Engineering and Decision Support Group

² Computer Science Department

Institute of Engineering – Polytechnic of Porto

R. Dr. António Bernardino de Almeida, 431. 4200-072 Porto, Portugal

{amn,besc,acm}@isep.ipp.pt

Abstract. This paper presents an artificial intelligence-based architecture to be used in various tourism services, namely the user modeling and recommender system components. Through the use of several machine learning techniques, such as linear models, neural networks, classification and even text mining, a hybrid and complete approach at understanding tourism application users is attained. To assess the advantages of the model, a specific prototype was created, including a complex user modeling process and an advanced and innovative recommender system. The system was tested in the scope of Oporto city, in Portugal.

Keywords: Recommender systems; tourism; user modeling.

1 Introduction

The tourism domain, and in particular, the holiday choosing process, represent very complex decision-making matters. The user is generally faced with the endless group of existing options, apart from the heterogeneity of the places to visit, for example, the type of Points of Interest (POI) available, hosted events and so on.

Such a scenario is clearly electable for the use of Artificial Intelligence (AI), and in particular, Recommender Systems (RS) [11], in order to filter the item search space. Such systems require the use of a coherent user model in order for results, and the overall system, to be customized and targeted to him / her: that's where User Modeling (UM) techniques come into play [1]. After an assessment of the current tourism systems, as well as other systems which use RS, there is the belief that significant work can still be done, regarding a more complex modeling of users, the introduction of innovative filtering techniques and ultimately more intelligent domain item recommendations.

This paper is organized as follows: section 2 will present a brief approach regarding RS in tourism, while section 3 presents the RS itself. Within this section, all different techniques that come into play when the RS is fired are thoroughly depicted. In section 4, the developed tourism application will be presented regarding its advantages and testing results. Finally, in section 5, some conclusions will be outlined, along with future work perspectives.

2 Recommender Systems in Tourism

Despite recent evolutions in RS, current systems still do not explore the huge potential of this area, namely in the tourism domain. User information, needed for the filtering effort, is extensively requested at startup, without necessarily being used throughout the remaining life-cycle of the application. Most of the times, though, systems rely on single and / or poor representations of the assumptions they make about users, which ends up in incoherent recommendations being performed [3][15]. A few examples follow:

- FilmTrust [14]: FilmTrust is a web-based system that explores the concept of trust in a movie related social network. In FilmTrust, users can not only express their particular opinion about a movie, but also define a trust degree for other users and their opinions. This follows the principle of basing predictions on reliable peers, instead of solely on similar ones.
- TripAdvisor [18]: this tourism website advises locations and activities for users, also containing a highly social component. Although this application contains one of the most hyped Recommender Systems (RS). UM does not seem to make a great part of the system's philosophy, being replaced by purely social theories;
- WAYN [19]: this application is an evident proof of a Web 2.0 social endeavor regarding tourism. This application featuring a very complete tourist profile; however, that profile is not used for any kind of serious reasoning, such as a RS.

Since the data available is not very extent, it was not chosen to perform a formal comparison between the different systems. Still, informal considerations can be extrapolated. In a very broad statement, the current main flaw regarding tourism RSs is the *poor UM* backing them up. Most systems rest their efficiency on a single UM technique; even if such modeling is not incorrect, it is certainly not enough, considering the complexity of the human being in various aspects. Another approach surfacing in the latest years is the overrated preference for social filtering methods [19]. While the use of this kind of recommendations is not uninteresting (we also embrace them in this very work), it is clear that the user itself is still the most important and primary source of recommending material, one whose deep analysis has not yet been performed in such a complete manner as will be presented. The growth of social networks and the current Web 2.0 is making RS *too social-targeted*, providing the UM architecture with a void that can still be exploited. To finish, and in such a domain where both personal and social interests are at stake, it is with surprise that systems do not rely on information related to the *psychological / behavioral* side of users, a key component of decision-making processes.

Given this action space for improvement, the proposed work tries to evolve the filtering of RS and the way systems can benefit from different kinds of user information, by using a complex UM framework. In the next pages, a RS architecture will be presented.

3 Development

3.1 Recommender System

Although tourism systems are more than a mere RS, that is the most important step regarding the user point of view, since the underlying mechanisms are invisible to them. Contrary with what the state of the art may make us think, adapting the system to accommodate user's preferences and needs can be accomplished using a variety of techniques. Figure 1 presents all techniques contemplated in the devised RS, as well as their importance.




Fig. 1. Recommender System's Techniques

Given the internal nature of techniques (Fig 1) it is clear that, for example, socialization techniques are less important than stereotypes, since the assumptions taken for granted socially might be completely wrong, a situation with much less probability in the other technique. The criteria used for such an order was the absolute relation between its results their intimacy with the user, also according with some tests executed. Another interesting aspect that can be experienced is that, in opposition with this intimacy order, lower techniques offer a much higher probability of giving refreshing results. Although out of the scope results are considered by many authors as a bad subject, promoting over-specialization, in this work we do not abide by those opinions. In fact, we express that such a point of view is much necessary in order to evolve user tastes, showing a sign of intelligence that users can interpret. Table 1 relates the presented techniques with the classical literature approaches. The proposed system makes a commitment into balancing (1) the acceptance of traditional techniques and (2) introducing innovations into each one of them and also proposing a new one [9][10][15][16][17].

Table 1. Comparison Between Literature Techniques and Proposed System

System Technique	Literature Technique
Likelihood Matrix	Knowledge-based Filtering
Keywords	Content-based Filtering
Socialization	Collaborative Filtering
Stereotypes and Psychological Model	Behavioral / Psychological Filtering

All of the above techniques will be explained regarding their internal mode of operation and evolution. In fact, those mechanisms are independent of the RS, being able to function without this component on top of them. Following is an explanation of how these play their role, in fact, in order to help the RS.

3.2 User Modeling Techniques

As it was earlier implied, a powerful RS is endorsed by several techniques that maximize the accuracy of the necessary user information in a variety of ways. Our UM methodology (which backs up the RS) must therefore be a collaborative effort of several sub-systems, each of them responsible for the retrieval of part of user data. Furthermore, it is also believed that knowing a certain user information space by using more than one method simultaneously successfully increases confidence in existent assumptions and divides responsibility amongst various techniques, which ultimately results in a system with more solutions and contingency plans [2][3][4]. This knowledge retrieval spirit represents the true evolution of our system against the considered state of the art. The system’s knowledge discovery components can be observed in Fig. 2.



Fig. 2. Knowledge Discovery Mechanisms’ Architecture

It can easily be seen that tourism applications’ upper-level core functions (namely the RS) gather information throughout all sub-systems and merge that data into a coherent user profile, in order to generate new information. These representation formalisms represent an advanced view of the user profile and allow for the generation of value added (possibly new) knowledge.

3.2.1 Likelihood Matrix

The likelihood matrix is responsible for linking the user with each one of the categories created in the POI taxonomy, being classified as a linear model in what literature techniques are concerned. It ranges from -1 to 1, where -1 means total unlikelihood and 1 represents complete interest. The use of likelihoods between types of POI and

the user is not new, as it has already been used in other systems and using different formats [2]. However, the techniques employed here allow rarer analysis such as positive / negative likelihoods, thus the choice for the -1 to 1 floating number. Plus, the underlying taxonomy is much richer and structured than the majority of other tourism systems [19]. This mechanism is the basis of the stereotype module and thus both components work together in order to provide an over-confident representation of user interests. Although one technique is based on the other, their abstraction level is different, triggering different results by both components. The proposed approach coherently represents both user likes and dislikes, by maintaining a negative and positive action space in which assumptions can diverge within. Moreover, by defining optimal thresholds, it is possible to identify the most important POI categories within users and trigger adequate response. The likelihood matrix is fed by Application Interaction Triggers (AITs), which monitor important and relevant actions executed by users within activity sessions..

3.2.2 Stereotypes

Stereotypes within the tourism domain are not inexistent, but they do not provide evolutionary capabilities as those that were developed and employed here, although they are described in other interesting manners, such as demographic attributes [7]. Stereotypes represent a widely used information abstraction mechanism used to group users into categories. The work made in [7] was fundamental in our stereotype component definition, as several ideas were mirrored in our system. Our stereotype system can be explained through a set of development guidelines which originated it. First of all, the POI taxonomy was *re-conceptualized* into hierarchical terms which would better serve as the basis for the stereotype construction; then, an initial set of *stereotypes* was created, each of them being fully described using the previous terms which will form the comparison basis for that stereotype to be linked to a user. Finally, *mechanisms* were created to compensate for an eventual insufficiency that might describe the initial set of stereotypes, as well as the suitability of their terms.

The concept hierarchy's creation process resulted in adequate terms to be used in the stereotype creation. An example of a mapping between the original taxonomy and these concepts is the following: Pubs + Nightlife + Festivities = Partying. Six initial stereotypes were created, along with the activation conditions to trigger the link between a user and a particular stereotype. Upon new user arrival, the registration form will be responsible for matching him / her to one of the existing stereotypes within the system. From that point on, different user actions will put him into different stereotypes, which will trigger different approaches in several system components, such as the RS [7]. The activated stereotypes for a given user in a certain point in time will always be those whose value surpasses the Activation Threshold, which will determine the nature of recommendations given to that user. A very important feature of stereotypes is, as said before, the group of mechanisms for coping with stereotype decay that were defined. Those are: (1) propose underused stereotypes for removal, (2) propose underused stereotype conditions for removal, (3) propose overused conditions not included in stereotypes and finally (4) propose new stereotypes based on user profiles (eliminates the grey sheep individual issue).

It's in these four processes that stereotypes can obviously be inserted into the category of clustering and classification techniques [20].

3.2.3 Keywords

The concept behind keywords is socially very powerful, addressing knowledge about items in a fashion more intimate to the user. Plus, by employing a knowledge representation with significant value to people, it also embraces the social power of the system. Keywords require significant user participation in order to be fully profitable; however, in recent years, due to the social web appearance, the idea got more doable [8]. The advantages of keywords are numerous:

- The power of keywords surpasses any other kind of controlled information representation, like the POIs taxonomy, granting value-added knowledge;
- It means the featuring of items in the user point-of-view, which enhances the application's real world / social sense and brings it closer to the user;
- Slowly and partially remove the importance of system-defined approaches (more strict and controlled), putting the power of the application in the users hands;
- By storing and relating keywords with users, tags can propose themselves as yet another means of filtering items to the user, along with all other methods.

Several approaches for initially inserted tags within the items were devised, without having to take an intensive cognitive process of cataloging them. By automatically gifting items with tags, one of the few downsides of this kind of social media, the cold-start problem, can be diminished. The following types of tags are initially and automatically set for POI:

1. Keywords that relate to the POIs class which classifies the POI. For instance, a zoo would have the tags *cultural*, *natural_parks* and *animal_preserves* automatically associated upon creation;
2. Keywords that pertain to all features related with that POI, which were previously added to the respective POIs class. For example, a certain restaurant might have the keywords *spicy* and *Mexican*;
3. Keywords that pertain to special words found within the name and description of the item, by using text mining techniques. For example, the bridge "D. Luis I" might have the keywords *bridge*, *Gustave_Eiffel* and *iron*.

3.2.4 Psychological Model

The user psychological model is in constant evolution, as the user interacts with the system and gives it traces of his personality evolution. This information might also be initially given by the user, to propel the system with a more coherent start. The features selected to be part of this module were based on psychological models devised by authors along the years [5][12][13]. Four attributes were elected: liveliness, perfectionism, outdooriness and creativity. They range from 0 to 1, representing the two extremes of that feature. The user psychological evolution proceeds as follows: each POI category is labeled by one or more of the four psychological attributes; then, interactions with POIs or POIs categories will trigger analysis between psychological models of both the user and the corresponding POIs classes. The analysis results will feed and evolve the user behavioral model towards the comparison o object (the POIs class), therefore adjusting it and changing the input of all components depending on it, such as the RS.

The effective and useful use of psychological data is very rare in the current computational scene in general, let alone in the tourism or RS [10]. The user psychological model, along with the use of stereotypes, represents a new approach - *behavioral-based* - in modeling and recommending items to users and contributes to the innovative nature of this work. This technique is also an example of a predictive statistical linear model, since recommendations are based on top of past experiences.

3.2.5 Socialization

The recommendation of items based on user's social data has been greatly used in the past years. However, as was already explained, it occupies the sixth position within the order of the filtering techniques, due to the criteria chosen for that order. In fact, one of the main objectives of the devised project is the exploitation of intra-user filtering techniques, given the space for improvement. Even so, this technique presents innovations of its own, profiting from the complex UM architecture that the system uses. Indeed, the manner by which users are matched similar against the main persona is made much more coherent, by using the last four explained techniques. This way, despite the technique being given less importance, it presents, in fact, a much more sustained theory for matching users, in respect with the current state of the art. Actually, the same comparison method is also used in a friend RS existent within the yet superficial social network which gifts the system.

3.2.6 Community Models

Our UM architecture makes use of a group of two Community Models (CMs). CMs are an adaptation of Neural Networks (NN), specifically a rather simple form of those. CMs generate two-dimensional representations of the data previously fed into its mechanisms [6], evolving and activating its nodes and links. CM's output is mirrored by other techniques, such as dynamic neural networks and Bayesian networks. However, the amount of increased work necessary to deploy those kinds of techniques compared to the value-added knowledge they provide, as opposed to these CMs, was not profitable [20]. CMs, when properly fed, can give us interesting patterns and associations between items that were previously invisible. Both CMs represent POIs in their nodes, while the relations between nodes represent POIs co-occurrence. The first CM is about system navigation (user sessions, *clickstream* analysis, etc.), while the second is concerned about effective POI commitment, such as the generation of a visiting route. Fig. 3 displays a fictional portion of one of the CMs. The color of each node is a reference to geographical position, meaning that CMs might act as a clustering technique, naturally grouping POIs, physically related or not.

Although much of the presented CM's usefulness is already attended in current systems, some of them are unprecedented, given the current state of the art:

1. Discover POI associations, by analyzing patterns contained within selected groups of items. Found associations will, apart from triggering important strategic decisions by related entities, become an effective help in understanding user personalities and also improving the RS's efficiency. This kind of operation is a very simple but effective kind of Rule Association (RA) analysis;




Fig. 3. Community Model's Example

2. If, ideally, the most used POI would be the same that were previously viewed, that might not always be the case. By searching for abnormalities between viewed versus used items, mechanisms can be triggered, which may include an increase in system visibility of some least selected items, for instance.

4 Evaluation

In this section, some evaluation guidelines will be presented in order to demonstrate the capabilities of the devised system, as well as a survey that targeted the so-far users of the application.

Startup quality: the devised prototype delivers an increased startup quality of response concerning filtering features and personalization mechanisms. By making use of a clever and abstracting model for initially asked information, the application can instantly achieve a coherent user profile and subsequent sustained suggestions. It can, technically, fill half of the user model information within the initial form. This startup quality is not performed by other systems [18].

Transparency: besides profiting from an automatic UM platform which does everything in an automatic fashion, the user can also be invited into viewing, in a transparent manner, everything that the system believes about him. By making use of a friendly interface, users can increase confidence of such critical information and enhance RS's results immediately.

Recommender System: the RS state of the art (see [9][10][15][16]) has reached a very critical state regarding innovation. The UM platform here presented forms a very diverse basis for RS computation and introduces a new way of filtering complex-domain items: behavior-based. This technique merging causes the RS to output items with diverse sources and assumptions, increasing likelihood for item success. Plus, by making use of several techniques, the RS has the ability to overcome problems that certain techniques might have, such as under-confidence or the cold start problem.

On-the-fly profile evolution: almost all of the UM building blocks are propagated immediately, causing the consequent RS to always output fresh results. Thus, without damaging any interaction flow during user sessions, the user is immediately gifted with

up-to-date responses from any application-level process. Such instant adequacy of results is not performed, for example, by the TripAdvisor application.

Diverse knowledge: the UM components that comprise the system use knowledge representation formalisms of diverse sources. One of the most important ways of analyzing these different sources is through their degree of control. While controlled knowledge (like the likelihood matrix) allows for sustained and guaranteed outputs, uncontrolled knowledge, existent in keywords, for example, grants users freedom and control in evolving the way the system works.

Survey: regarding the user survey that was made, the following conclusions were taken: (1) users were pleased by the short initial form; (2) users enjoyed the transparent spirit of the system and show interest in using the User Area; (3) users have found the RS's results satisfactory, lacking only a higher degree of speed, due to the complex reasoning tasks that are put to work when the RS fires off.

5 Conclusions

Machine learning reasoning hasn't been fully explored in the tourism domain, particularly in what the tourist model is concerned [20]. Tourism applications are indeed in a widespread state, but they still haven't taken the evolution step of intelligence, as well as usability. Users still have to perform several actions to achieve what they really want, and a substantial amount of smart reasoning from the system doesn't seem still present [18][19]. Moreover, user information has been misused and falsely utilized. On one hand, current systems ask more information than they really use, which is a downside both usability-wise and resource-wise [19]. On the other hand, UM and user adaptive-related mechanisms are used in a very poor manner, based on very weak assumptions. In the proposed system, a modeling platform that requires only a minimal amount of user effort and still manages to create a very capable user image is deployed. Next is a list of advantages that, when compared with current projects (see section 2), dictate the importance and innovative nature of this project: (1) advanced and innovative UM, (2) on-the-fly user profile update, (3) hot-start results quality, (4) behavioral-filtering introduction and (5) multi-technique and heterogeneous RS. From our point of view, the user model present in this paper (stage one), along with the more complex knowledge inference mechanisms that constitute the devised system (stage two) are two excellent basis for any tourism-focused system. The architecture presented can actually be applied to completely different scenarios and domains, given some content minor content adjustments.

References

1. Gay, G., Hembrooke, H.: Activity-Centered Design: an Ecological Approach to Designing Smart Tools and Usable Systems. MIT Press, Cambridge (2004)
2. Fink, J., Kobsa, A.: User Modeling for Personalized City Tours. Artificial Intelligence Review 18(1) (2002)
3. Kobsa, A.: Generic User Modeling Systems. In: User Modeling and User-Adapted Interaction, vol. 11(1-2). Kluwer Academic Publishers, Dordrecht (2001)

4. Kobsa, A.: User Modeling: Recent Work, Prospects and Hazards. In: Adaptive User Interfaces: Principles and Practice. North Holland/Elsevier (1993)
5. Jung, C.G.: Psychological Types. Princeton University Press, Princeton (1971)
6. Jennings, A., Higuchi, H.: A Personal News Service Based on a User Model Neural Network. In: User Modeling and User-Adapted Interaction, vol. 3(1). Springer, Netherlands (1992)
7. Rich, E.: User Modeling Via Stereotypes. *Cognitive Science: A Multidisciplinary Journal* 3 (1979)
8. Mathes, A.: Folksonomies - Cooperative Classification and Communication through Shared Metadata. In: Computer Mediated Communication - LIS590CMC (2004)
9. Berka, T., Plößenig, M.: Designing Recommender Systems for Tourism. In: ENTER 2004 (2004)
10. Felfernig, A., Gordea, S., Jannach, D., Teppan, E., Zanker, M.: A Short Survey of Recommendation Technologies in Travel and Tourism. *ÖGAI Journal* 25(2) (2006)
11. Coelho, B., Martins, C., Almeida, A.: ATOMICO - Arquitetura para Turistas com vista à Organização, Modelação e Inferência de Conhecimento - Representação do Utilizador. In: IADIS Ibero-Americana - WWW / Internet 2008 (2008)
12. Oliver, P.J., Srivastava, S.: The Big-Five Trait Taxonomy: History, Measurement and Theoretical. In: Handbook of Personality: Theory and Research. Guilford Press, New York (1999)
13. Cattell, R.B.: Handbook for the Sixteen Personality Factor Questionnaire. The 16 P.F. Test - Forms A, B, and C. Institute for Personality and Ability Testing (1962)
14. FilmTrust (2010), <http://trust.mindswap.org/FilmTrust/>
15. Burke, R.: Knowledge-Based Recommender Systems (1999)
16. Pazzani, M.J., Billsus, D.: Content-Based Recommendation Systems. In: Brusilovsky, P., Kobsa, A., Nejdl, W. (eds.) Adaptive Web 2007. LNCS, vol. 4321, pp. 325–341. Springer, Heidelberg (2007)
17. Porter, J.: Watch and Learn: How Recommendation Systems are Redefining the Web (2006)
18. TripAdvisor (2010), <http://www.tripadvisor.com/>
19. WAYN (2010), <http://www.wayn.com>
20. Zukerman, I., Albrecht, D.W.: Predictive Statistical Models for User Modeling. In: User Modeling and User-Adapted Interaction, vol. 11(1-2), pp. 5–18 (2000)

Knowledge-Based Geo-risk Assessment for an Intelligent Measurement System

T. Vicovac¹, A. Reiterer¹, U. Egly², T. Eiter², and D. Rieke-Zapp³

¹Institute of Geodesy and Geophysics, Research Group Engineering Geodesy, Vienna University of Technology, Gusshausstraße 27-29, 1040 Vienna, Austria

²Institute of Information Systems, Research Group Knowledge-Based System, Vienna University of Technology, Favoritenstraße 9-11, 1040 Vienna, Austria

³Institute of Geological Sciences, Research Group Exogene Geology, University of Bern, Baltzerstrasse 1+3, 3012 Bern, Switzerland

{tanja.vicovac, alexander.reiterer}@tuwien.ac.at,
{uwe, eiter}@kr.tuwien.ac.at,
zapp@geo.unibe.ch

Abstract. Rockfalls and landslides are major types of natural hazards worldwide that kill or injure a large number of individuals and cause very high costs every year. Risk assessment of such dangerous events requires an accurate evaluation of the geology, hydrogeology, morphology and interrelated factors such as environmental conditions and human activities. It is of particular importance for engineers and geologists to assess slope stability and dynamics in order to take appropriate, effective and timely measures against such events. This paper presents a decision-tool for geo-risk assessment on the basis of a knowledge-based system. The integration of such a tool with novel measurement sensors into an advanced system for geo-risk monitoring, which performs data fusion on-line, is innovative. To enable such a system, a knowledge base capturing domain knowledge formally is developed, which to the best of our knowledge is unprecedented; the completed part for initial risk assessment works quite well, as extensive experiments with a number of human experts have shown.

Keywords: Knowledge-Based System, Alerting System, Rockfall and Landslide Monitoring.

1 Background and Motivation

In the last years, damage caused by rockfalls and landslides has been increasing, as well as the number of persons that were killed or injured, due to a spread of settlements in mountain areas. In addition, many global climate change scenarios predict an increase in the probability of heavy rain, which is a primary trigger for rockfalls and landslides. This causes an urgent need for highly effective and reliable tools for monitoring rockfalls and landslides at an operational level. The increasing importance of rockfall and landslide monitoring is clearly also reflected by a large number of research projects. For example, in its last two framework programs, the European

Commission has positioned research about “Natural Hazards” and “Disaster-Management” as a priority topic. The core of geo-risk management consists of identifying, understanding and mitigating risk by reducing the probability or consequences of rockfalls and landslides. In the literature, several geo-risk management and geo-monitoring systems can be found; most notable are [5, 10, 11, 13, 14]. Examples for systems used in practice are GOCA [7] and GeoMoS [6]. The main application field of these tools is monitoring and analyzing deformations; however, they offer no possibility for deformation interpretation. Currently this is done by human experts from geology and civil engineering, who are interpreting deformations on the basis of a large number of data records, documents and knowledge of different origin.

Given the increasing number of problematic sites and the limited number of human experts, automated intelligent monitored interpretation systems are required. The implementation of a knowledge-based system enables an automatic process of interpretation and determining of the risk potential. In contrast to the mentioned monitoring tools (e.g., GOCA), it is possible to perform deformation interpretation with our system. Based on the measured deformation vectors, a measurement preprocessing is performed (mainly clustering to detect areas of similar movement). On the basis of this information and additional data about velocity and orientation, some conclusions about the kind of occurring movement can be drawn. Additionally, data of different, heterogeneous sources, such as geodetic deformation measurements, geotechnical measurements, geological maps, geomorphological maps, in-situ investigations, and numerical modeling methods have to be included in such a system.

It should be emphasized that the integration of a knowledge-based system for solving this task represents an innovative method.

At the Vienna University of Technology (Institute of Geodesy and Geophysics), the interdisciplinary research project *i-MeaS* (“An Intelligent Image-Based Measurement System for Geo-Hazard Monitoring”) [8] has been launched with the purpose of research, develop and implement an interpretation tool for geo-risk objects. The system gives on-line information about ongoing deformations and supports issuing alerts in case of excessive deformation behavior.

Making conclusions about incidents is a not-trivial problem; by using artificial intelligence techniques, via the integration of a knowledge-based system, new directions are opened up. Such decision systems are currently used in medicine [2] as well as image mining [1]. This new system is a complex intelligent system, working with several different data sets in real-time. Deformation measurement data will be delivered by a novel type of measurement system, which consists of two image-based sensors. Inside the captured images so-called *interest points* are detected. The calculation of the 3D coordinates is done by classical geodetic forward intersection¹. By means of such a high precision measurement system, 3D object points can be detected with an accuracy of about 2-3 mm (object distances up to 1000 m). Subsequently a geodetic deformation analysis can be performed that yields as a result deformation movement vectors, which constitute the input for later interpretation.

In this paper, we report on the architecture and functionality of the respective interpretation system and its development stage. In particular, we present a knowledge base for risk assessment, which to the best of our knowledge is unprecedented, and as

¹ Forward intersection is a standard method for determining 3D object coordinates from 2D image points.

comparative tests with a number of domain experts indicate, works well compared to human experts.

2 System Concept and Architecture

Remote monitoring of unstable slopes is a typical multidisciplinary problem incorporating a network of sensors of different kinds. Movements and deformations can be measured, for instance, with geo-technical sensors (e.g., inclinometers, tilt-meters, extensometers, etc.), or optical measurement systems (e.g., tacheometers, laser scanners, etc.). Most of these sensors must be placed on-site; in hazardous terrain this is very often not possible. It is thus also necessary to apply remote monitoring methods, some of which are based on photogrammetric methods or terrestrial synthetic aperture radar (SAR). Both yield multi-temporal images that contain distances to the scene in each pixel.

Current Systems. Recently, the interest in image-based measurement systems has increased. *Leica Geosystems* [16] developed a prototype of an “image-assisted total station” with the purpose of defining a hybrid or semi-automatic way to combine the strength of the traditional user-driven surveying mode with the benefits of modern data processing. Furthermore, *Sokkia* [15] introduced a prototypical tacheometer which provides focused color images.

The central task of all image-based deformation measurement systems is the calculation of 3D object coordinates from 2D image coordinates for a subsequent deformation analysis or object reconstruction. The basic idea of deformation measurements is capturing a zero state of the object (measurement epoch 0) and one or more subsequent object states (measurement epoch n). The time interval between the measurements depends on the type and the estimated behavior of the objects.

All these measurement systems are based on a permanent user interaction. Selection, pointing and measurement of the relevant object points have to be operated by a measurement expert. Most of the relevant processing steps are fully manual. The challenging task of the mentioned *i-MeaS* project is to develop a fully automated system (user interaction will be possible at different decision levels). Data capturing, data analysis and data interpretation should be performed as an automated process.

System Concept. In our system, we are using a new kind of optical measurement system which is based on a traditional tacheometer system², namely an image-based tacheometer system. In comparison with laser scanners, this system measures objects with higher accuracy; compared to photogrammetric systems, they are easier to use for on-line measurement processes (e.g., object monitoring), especially because measurements can be done with a high degree of automation.

The processing chain of the new measurement concept starts with the capturing of geo-referenced images, followed by image point detection and by 3D point measurement. The final output of this measurement process is a list of 3D deformation vectors of the object (deformations captured between two or more measurement/time epochs).

² A tacheometer is a surveying instrument for distance determination and measurement of horizontal and vertical angles.

Measurement data is one of the basic elements of decision-making – however, many other factors are used by the system (more details are given below). The system architecture can be divided into several components:

- the measurement sensors (e.g., geodetic, geotechnical and on-site placed meteorological sensors),
- an image analysis system (which is needed because some sensors are working on the basis of captured images),
- a system control component,
- a knowledge base,
- a system for deformation analysis, and
- a system for alerting.

Furthermore, the system includes an user interface. The simplified architecture of the system, with the knowledge base and the system control component as core units, is shown in Fig. 1. A description of the measurement system can be found in [8].

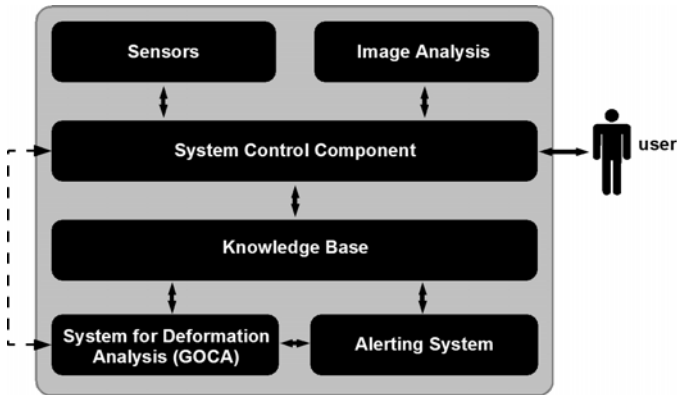


Fig. 1. Simplified architecture of the system

As mentioned above, such a complex system works on the basis of heterogeneous information. We are using the following information sources:

- generic domain knowledge (i.e., knowledge about coherences of influence factors and general deformation behavior),
- case-specific knowledge (i.e., domain knowledge collected via historical notes, morphological and geological maps), measurement data (geodetic, geotechnical, geophysical measurement data, etc.),
- local weather data (like local temperature, the amount of precipitation, the kind of precipitation, etc.), and
- global meteo data, which are provided by meteorological services.

In order to test the optical measurement system under realistic field conditions the sensor system was installed over several days on the “Pasterze” glacier, the largest glacier in the Eastern Alps. The end of the glacier which is covered by debris and a

geologically stable rock face was chosen as test site. The main purpose of the test was the evaluation of the point detection and the consecutive calculation of 3D coordinates under realistic environmental (especially illumination) conditions. The results show that the measurement system works very well and moreover could measure 3D object coordinates with an accuracy of 2-3 mm (the distance to the object was about 1,000 m).

3 Risk Assessment

As mentioned above, there is a high demand for reliable prediction of hazardous events. Apart from monitoring, the detection of possible causes is a matter of particular interest. According to these preconditions the requirements on a geo-risk assessment system are quite high. Eventually, knowledge-based interpretation should be able to draw conclusions about the type of occurring movements as well as providing possible causes for them.

The concept of data interpretation is based on the “calculation” of risk factors for critical cause variables and on the elaboration of an interpretation for the deformation. Examples for cause variables can be precipitation, permafrost, critical slope angle, etc. The range value of the risk factor is divided into six classes (low, low-medium, medium, medium-high, high, very-high). This definition is based on the discussion results with experts.

The challenging problem in developing such an alerting system is (1) to identify relevant factors and (2) subsequently to capture the interlinkage of these influence factors. The latter are described in the next section.

In our system, the process of risk assessment is divided into two steps: (1) the determination of the “*Initial Risk Factor*” and (2) the determination of the “*Dynamic Risk Factor*”. The first step estimates the plausibility of an occurring moving event. Furthermore the zero state of interpretation and the observation is defined.

The second step is focused on the processing of the temporal development of the risk factor. Therefore additional data have to be included into the decision process, e.g., measured data captured by the image-based monitoring system. Measurement data represent the 3D object deformations (data is captured in defined time periods resulting in movement/displacement vectors). As mentioned above, the system is also able to access local and global meteorological data in real-time, which can be used by the dynamic system as a basis for deformation prediction.

This process is leading to a detailed description of the deformation and an actual estimation of the risk factor, standardized on a predefined scale, which can be directly used as a global indicator for the likelihood that a landslide or a rockfall will occur.

In practice, the estimation of the risk factor is a continuous process, in which the dynamic risk factor has to be determined in a periodic feedback cycle.

In the following, we focus on the determination of the “*Initial Risk Factor*”. Beside difficult technical requirements related to sensor and data fusion, the most challenging tasks in developing such a system is the implementation of the knowledge base and, in a preliminary step, the knowledge acquisition. This problem was solved using a two-step approach: in the first step, a single expert was consulted, while in the second step an extensive system evaluation by many experts was carried out and their feedback was incorporated into a system refinement (details are given below).

4 Construction of the Knowledge Base

As described, the challenging problem is the realization of the knowledge-based part, especially the acquisition of knowledge from human experts. For this, we adopted a common methodology [3, 12]³. In order to estimate the initial risk factor, influence factors had to be identified whose values increase the likelihood of deformation. During a period of extensive discussions about the domain problem, more than thirty-five relevant influence factors were identified (e.g. vegetation, granular material, subsoil, pieces of rock, indicates, slope angle, slope profile, slip surface, material underground, saturation of soil, leaning trees, leaning rocks, crack, rock joint, joint orientation, insolation, permafrost, stone chips, frost-thaw-cycle, depth of movement, local temperature, etc.). About thirty of them are used for the determination of the “Initial Risk Factor”. Some factors including examples for possible consequences are listed in Table 1.

Table 1. Examples for influence factors for the initial risk factor and possible consequences

Influence Factor	Examples for Consequences
<i>vegetation</i>	The vegetation has an influence on the slope stability and soil saturation.
<i>granular material</i>	In interaction with slope angle and the kind of subsoil a conclusion about the slope stability can be done.
<i>slope angle</i>	In interaction with slope angle and the kind of subsoil a conclusion about the slope stability can be done.
<i>slip surface</i>	Existing slip surfaces are an indicator for slope movements.
<i>soil saturation</i>	The degree of saturation is dependent on the vegetation on the surface.
<i>leaning trees / rocks</i>	Leaning trees and rocks are indicators for slope movements.
<i>insolation</i>	Insolation can affect factors like soil saturation and in combination with the influence caused by granular material and slope angle a conclusion about the slope stability can be done.
<i>permafrost</i>	The existence or absence of permafrost has an influence on the slope stability.

On the basis of the identified factors and the discussion, we have developed an *online-questionnaire*, which serves as a makeshift for assessing the “*Initial Risk Factor*” of the object to be observed. The questionnaire comprises questions ranging from the geological and morphological characterization, the vegetation, and the hydrology of the mountainside to administrative issues. The expert may answer all object-relevant questions in-situ/online, usually using multiple sources to find the answers; geological and geomorphological maps, historical documents, data of earlier geotechnical or geodetic measurements of the observed slope, and last but not least inspection of the endangered area.

Discussions with several experts revealed that estimating a risk factor on the basis of many influence factors and extensive domain knowledge is highly complicated. Moreover, experts sometimes largely disagree. Thus, a system which incorporates the opinion of more than one expert is indispensable to guarantee continuous high-level quality of decisions.

³ Currently we are using no automated knowledge acquisition methods – for the future the use of such methods is envisaged.

For estimating the mentioned risk factor, we developed a knowledge-based system, adopting a rule-based approach, more specifically using production rules. This is because the connection between influence factors and possible causes or deformation behavior can be naturally formulated by rules, and this representation is more accessible to domain experts than other representations.

For the implementation, we have chosen JESS [4, 9], which is a rule engine and scripting environment entirely written in JAVA. JESS is easy to learn and use, is well documented and supported; moreover it supports fuzzy rules and uncertainty. Furthermore JESS integrates well into the Eclipse software development environment which is widely used in industry.

In order to make the collected numerical features (like measurement data, meteorological data, etc.) more suitable for the rule-based decision system, we use an abstraction procedure that is in line with the expert view. It translates the numerical input values (image features) into linguistic concepts which are represented by abstraction (“fuzzy”) sets. More specifically, they form an *ordinalization*, i.e., the sets are characterized by non-overlapping step membership functions; hence, this translation is a pre-stage of full fuzzification. The use of such an abstraction enables decision rules in terms of easily-understood word descriptors instead of numerical values. Furthermore, all data sets are synchronized by a common time basis.

An example of a simplified initial state rule (IS) (JESS syntax) is shown in the following. The LHS of the rule ‘IS_riskpot_rockfall’ checks whether there are elements of type IS_ROCKFALL in the working memory, fact1 and fact2, where in fact1 certain slots (i.e., attributes) have certain values (danger_recent_deformation has value ‘high’, etc., and frost_thaw_cycle has value either ‘NO’ or ‘NN’), and fact2 states that a risk factor is not defined. The RHS includes the instruction to update the status of fact2 to have risk_defined true and to set risk_pot to high_4. The values of the used slots (danger_recent_deformation, danger_slope_angle, danger_bedding, danger_fine_grit) are determined from combination of input elements (separated rules).

```
(defrule IS_riskpot_rockfall
  (declare (salience 0))
  ?fact1 <- (IS_ROCKFALL
             (danger_recent_deformation==high) &&
             (danger_slope_angle == very_high) &&
             (danger_bedding == low) &&
             (danger_fine_grit == very_high) &&
             (frost_thaw_cycle == NO || frost_thaw_cycle == NN))
  ?fact2 <- (IS_ROCKFALL{risk_defined != YES})
=>
  (modify ?fact2 (risk_defined YES)(risk_pot high_4))
```

Generally, the rule base is divided into two groups of rules: (1) rules regarding the connections between facts and consequences (e.g., rain and the consequent possible deformation of the object), and (2) rules determining the initial risk factor. The mentioned example is part of the second group. The whole “initial-risk-factor-system” consists of about 70 rules. It is also notable that we have developed a tool for visualization of rules and their firings, which helps in grasp and analyzing data dependencies.

5 Evaluation and Experiments

After developing and implementing a prototype for determining the “Initial Risk Factor”, we performed an evaluation in a two step approach: (1) evaluation by one expert, followed by an improvement of the system, and (2) an exhaustive evaluation by eight experts.

In step 1, 30 different data sets were prepared by one expert, where each data set models a test site with particular characteristics concerning slope profile, vegetation, insolation, etc. Facts like soil material, slope angle, hydrological properties, information about indicators for movement, etc. were predetermined. Then, the system processed the data sets and the risk factors determined were compared with the decisions by the single expert. The discrepancies were analyzed and the explorations were used for extending and upgrading the prototype system.

In step 2, eight geological experts had to appraise independently the 30 test cases. The resulting risk factors were compared with the result of the prototype system. It is striking to note that the risk factor between the different experts varies up to two classes. In exceptional cases, the difference is more than three. A statistical overview of the differences (Δ) between the eight experts and the system is shown in Fig. 3.

Also at the level of the individual test cases, the answer conformity was high; it is remarkable that about 42 % of the answers agree completely with the system, and 30% of the answers are at distance 1. For example, expert 1 agrees completely with the system decision in 18 test cases ($\Delta=0$), in ten test cases the decision has distance 1 and in two cases it has distance 2. From case to case even the disagreement between the experts can be quite high. This mainly depends on the individual experience and appraisal of each expert. For all the test case, there was 102 times complete agreement ($\Delta=0$) of the experts and the system, while the maximal disagreement ($\Delta=4$) was in 3 cases only.




Fig. 2. Statistic overview of the processed expert evaluation

The answers of the system and the experts completely agree in one case, and they span an interval of i classes, $i=1,2,3,4,5$, in 4,7,13,5 and 0 cases, respectively. Only in one case, the system answer is outside the interval of the expert answers. Furthermore, in 15 cases (i.e. 50%) the system answer is the median of the expert answers and in 8 cases (27%) one of the middle answers (e.g., in case 6 the system answer is 2, while the middle answers in the sorted expert answers 1,1,1,1,2,2,3,4 are 1 and 2, thus the median is 1.5); in 5 cases (17%), it is one class to the median, and only in 2 cases the discrepancy is higher, with a maximum of two classes in the outlier case.

As the above statistical data indicates, the risk assessment is a difficult task where expert experience is required to obtain a reasonable solution. This is witnessed by the fact that there is no sole “correct” assessment for many of the different test cases, and expert opinions on the risk can vary for the same test case. The system can compete with the human experts; the differences between the result of the system on one side and the experts’ results vary in a similar way as the results vary between different experts. In order to further improve the quality of risk assessments of the system and to test its usability, we initiated an even broader evaluation where we asked additional experts for their opinion on our test cases. These new experts should bring in a fresh sight on the problem and the system because they have no information about the system and no training on it. The results obtained from first new probands are in accordance with the system and with the former evaluators. This also shows that the system interface is intuitive enough such that an untrained expert can easily use it without major difficulties.

6 Conclusion

The main goal of the presented work is the development of an innovative automated knowledge-based interpretation system for predicting rockfalls and landslides with advanced measurement techniques. Towards this goal, we have carried out extensive knowledge acquisition and performed knowledge analysis with the help of several experts, leading to a rich knowledge base that builds on a number of influence factors, determined from various information sources (e.g., measurement data, expert knowledge, maps, etc.).

An experimental prototype for risk assessment we developed shows good results for the completed “*Initial Risk Factor*” part, in where it behaves like a human expert.

Currently, the *dynamic system* (the second part of the final system) is under development – besides extending the existing rule set, further rule components (e.g., for updating risk factor, including meteorological data, etc.) will be added. Future work will include testing (collecting field data is targeted for the summer of 2010) and the integration of the knowledge base component into an on-line geo-risk management system.

Acknowledgments

The research presented in this paper has been supported by the Austrian Science Foundation (FWF) Project L514 “*i-MeaS – An Intelligent Image-Based Measurement System for Geo-Hazard Monitoring*”.

References

1. Agrawal, R., Mannila, H., Srikant, R., Toivonen, H., Verkamo, I.: Fast Discovery of Association Rules. In: *Advances in Knowledge Discovery and Data Mining*, pp. 229–248. MIT Press, Cambridge (1996)
2. Clarke, K., O’Moore, R., Smeets, R., Talmon, J., Brender, J., McNair, P., Nykanen, P., Grimson, J., Barber, B.: A Methodology for Evaluation of Knowledge-Based Systems in Medicine. *Artificial Intelligence in Medicine* 6(2), 107–121 (2004)
3. Farrand, P., Hussain, F., Hennessy, E.: The Efficacy of the Mind Map Study Technique. *Medical Education* 36(5), 426–431 (2002)
4. Friedmann-Hill, E.: *JESS in Action. Rule-Based Systems in Java*. Manning (2003)
5. Fujisawa, K.: *Monitoring Technique for Rock Fall Monitoring*. Internal Report (2000)
6. GeoMoS, <http://www.leica-geosystems.com>
7. GOCA, <http://www.goca.info>
8. i-MeaS, <http://info.tuwien.ac.at/ingeo/research/imeas.htm>
9. JESS, <http://www.jessrules.com>
10. Kahmen, H., Niemeier, W.: OASYS – Integrated Optimization of Landslide Alert Systems. In: *Österreichische Zeitschrift für Vermessung und Geoinformation*, vol. 91, pp. 99–103 (2003)
11. McHugh, E.: Video Motion Detection for Real-Time Hazard Warning in Surface Mines. In: NIOSHTIC No. 20026093 (2004)
12. Negnevitsky, M.: *Artificial Intelligence – A Guide to Intelligent Systems*, 2nd edn. Addison-Wesley, Reading (2004)
13. Scaioni, M., Giussani, A., Roncoroni, F., Sgrezaroli, M., Vassena, G.: Monitoring of Geological Sites by Laser Scanning Techniques. In: *IAPRSSIS*, vol. 35, pp. 708–713 (2004)
14. Scheikl, M., Poscher, G., Grafinger, H.: Application of the New Automatic Laser Remote System (ALARM) for the Continuous Observation of the Mass Movement at the Eiblschrofen Rockfall Area – Tyrol. In: *Workshop on Advances Techniques for the Assessment of Natural Hazards in Mountain Areas, Austria* (2000)
15. Sokkia: <http://www.sokkia.com>
16. Walser, B., Braunecker, B.: Automation of Surveying Systems through Integration of Image Analysis Methods. In: Grün, Kahmen (eds.) *Optical 3-D Measurement Techniques VI*, Zurich, vol. I, pp. 191–198 (2003)

PART IX
Short Papers

Workflows Recognition through Multi Agents in Surveillance Systems

Manolis Sardis, Vasilis Anagnostopoulos, and Theodora Varvarigou

National Technical University of Athens (NTUA), Athens, Greece
{sardis, vanag, dora}@telecom.ntua.gr

Abstract. Workflow management systems exactly enact business procedures and processes described in a process description language. This strict adherence to the prescribed workflow makes it impossible for the system to adapt to unforeseen circumstances. Surveillance systems have unpredicted information especially in difficult environments like the industrial ones. In this paper, we are presenting a workflow recognition architecture through the use of a multi agent system that controls and evaluates the recognized processes from the surveillance algorithms and adaptively creates environment warnings or alarms. The related methodology is based on Java technologies which are presented and latest innovations from the multi agents and workflow processes composition.

Keywords: multi agents, workflows, surveillance.

1 Introduction

Watching a video from an online camera that is performing surveillance in a room, in a road, or in an industrial environment like the picture in Fig. 1, is very difficult even for the human eyes to recognize objects, humans or even more workflows of objects and humans. In Fig. 1, the workflow is performed by two workers that are getting from the automobile spare parts boxes a spare part and transfer it to a soldering environment, where a robot is performing the related accurate and fast soldering. The sequence of events either in place or time is critical for the start and stop of robot activity. Any delay or fault from the workers could cause a delay or even worst a stop in the whole manufacturing chain. This is an example of a workflow problem, where humans need to know online and automatically the workflow steps and any possible problems, in order to protect and prevent system and business process delays and faults. This paper presents an architecture framework for solving the above problem, using the state of the art of multi agent technologies and web technologies according on the surveillance restrictions of the algorithms that provide the recognized objects or the humans in the scenes.

Paper section two presents the problem statement and its requirements. Section three continues with the scientific tools technologies that we are using for our architecture. Section four presents the architecture details and implementation, and finally in section five the conclusions with future research activities.



Fig. 1. An industrial environment, a workflow stages, human recognition (white box)

2 System and Problem Environment

The workflow sequence in Fig. 2, presents the three steps/states (x_1 , x_2 , and x_3) that should be performed by one worker in order to fulfill the workflow number 1. The worker is available to go through a_{12} task to x_2 state, and a_{23} to x_3 state. The problem is that the worker can pause the execution from x_2 to x_3 , if he returns back to x_1 using the a_{21} route. The IT system should get all related information from the environment using the four sensors (y_1 , y_2 , y_3 , and y_4), which are installed cameras for surveying the place there.




Fig. 2. Workflow composition from different procedures

The signals from the y_i ($i=1,\dots,4$) sensors should inform the IT system that the workflow has not been completed. These dynamic changes, in a workflow tasks (coming as signals b_j in the system) execution, which cannot be described in a static workflow sequence, are the base of our problem in that paper.

3 Scientific Tools

There are several commercial implementations of workflow engines, with their own proprietary formats. None of them is yet focused on multi-agent systems, but in technologies like web services. Web services workflow can be automated using many tools such as Web Services Business Process Execution Language (WSBPEL) [1] and Microsoft’s X Language (XLANG) [2] specification. Workflow description is even more necessary, when composing a system from multiple web services, with multiple operations. Web Services Flow Language (WSFL) [3] can be used to define separately the flow model and the global composition model and its hierarchies. Web Services Conversation Language (WSCL) [4] focuses on protocols and document exchange and supports workflow concepts. WSBPEL is ideally for facilitating static Web Service composition. In case of dynamic workflows modeling we need an extra




Fig. 3. Modeler Agent interconnections

layer “*individual task layer*” where dynamic linking and manipulation is performed through WSBPEL. This layer in our case is performed by agents’ functionality as will be presented in the rest of the paper for enhancing and integrating the dynamic manipulation of web services [5]. The comparisons in [6] show that semantic web services need higher level information for agent-based automation. Modeling a workflow as an extended transaction means that the sub-transactions correspond to the tasks of the workflow and the execution structure of the extended transaction corresponds to the control flow of the workflow.

The main distinction between multi agents and web services is that web services are user-driven; multi agents act on their own [7] after user has given the necessary initial system instructions, providing adaptability on the workflow changes. Autonomous agents need more content information from the web service standards. Integrating web services and multi agent systems to a composite service requires a common method for using the system [8]; there has to be standard client software for the user, to contact agents, web services, or both. The need for this extra middleware software in our case is implemented through the usage of the *Modeler Agent*. The tasks and its position in the core architecture are presented in the following Fig. 3. The aim of this agent is to ensure that global constraints are not violated and that global efficiencies can be achieved.

According to Fig. 3, the *Modeler Agent* is responsible to perform all the intelligence of data (from sensors) manipulation. Data manipulation means that the web services data has to be evaluated and filtered in a manner readable for the end user. The system architecture is composed of multi-*Modeler Agents* that control and support a part of web services sensors. Each agent works standalone and is responsible for the online data from the web service. The distributed architecture, per sensor manipulation, benefits the system performance and response time, something very important for online systems and online characterization or identification of workflows in surveillance systems.

Workflows and Agents Development Environment (WADE) [9] enables a group of agents to cooperatively execute complex tasks defined as workflow. An Agent to perform its tasks may need to communicate with other Agents in the Platform. The Java Agent Development Framework (JADE) [10] provides the agents with the ability to communicate. WADE adds to JADE the support to the workflow execution. In our case, the use of WADE is to implement the internal behavior of each single system, and the full web services orchestration will be performed by BPEL usage.

4 Composition of Workflows Tasks and System Implementation

The workflows are synthesized by tasks and related procedures where an object/human has to perform in a sequence, and based on related time restrictions. The main types of workflows can be separated in two categories of related tasks/procedures of web services composition. The *static* and the *dynamic* web services composition, in order the system to recognize and label a workflow.

In *static* composition, the service to which the agents will be connected is determined before the workflow execution takes place (prior to run-time). Assume that we know exactly for a robot in the surveillance environment the root of its actions and the related paths during time. For internal automation of tasks, the recognition of the sub events from the *Camera Agent* (Fig. 4.) are composed by the workflow *Modeler Agent*, which links services offered by different services (from different *Camera Agents*) creating a static composition since the existence of such a service is known prior to run-time (for the case of the robot performance in the scene) and the agent can have design-time knowledge of connecting to such a service. Static composition of web services can be implemented through execution of BPEL constraints. Although the case works perfect for constant robots' based routes in a scene, not all services and procedures are known in case of humans actions, especially during the workflow design time. However, some of the services need to be discovered during workflow execution (run-time).

Dynamic composition of web services can be supported through the multi agents' environment where agents' services provide system adaptability [11]. Using a number of cameras and related services per *Camera Agent* the intelligence system is able to provide the correct answers for the recognized event and its position in the workflow sequence. This 'intelligence' is performed through the *Modeler Agents* that using WADE and according on the workflow process constraints, are manipulating and report the achievement or failure of the related surveyed workflow. The end user is informed online through the generated alarms or warnings signals from the gui. Dynamic composition and adaptability support can be ensured by applying more agents in the surveillance scene. The *Sensor Agents* can provide more detailed information of sub tasks. This information that can be from the simple *start* and *stop* of an event, into more detailed functions of the workflow allows the delays, the route changes, and the time and position drifting from the original positions of the surveyed object. The result is that the *Modeler Agents* have available extra information that limits possible dummy workflows processes that could create errors on the workflow recognition.

The system has been modeled using Java technologies on web services and on multi agents systems. The JADE platform has been used for the construction of the agents and their manipulation. The surveillance system consists during the lab tests of 3 PTZ Axis 213 cameras saving the surveyed scenes in .jpg files per frame. The network is a Gigabit Ethernet for the control of the cameras network and the agents virtual machines have been applied in 2 computers. Each *Camera Agent* is collecting the surveillance information from the object detection and tracking algorithms and then through an xml format is getting this information from a related web service. Composition of the web services information is done through the *Modeler Agent*. This JADE based agent controls the inputs and cooperates with the extra *Sensors Agents* for extracting correct workflows recognitions. Finalization of the whole process is performed through the end user GUI that integrates the whole architecture.

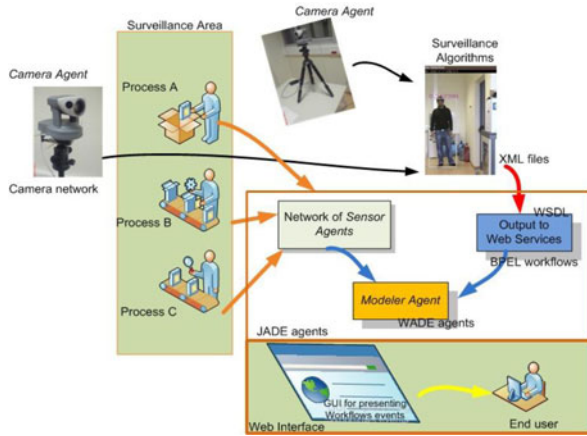


Fig. 4. Architecture framework

5 Conclusions

Through this paper the authors presented an architecture for workflows recognition using multi agent technologies on surveillance systems, under complex environments for surveillance of workflows, like the industrial environments. The paper presents the main architecture components and the innovation in manipulating dynamic composition of the processes that a workflow contains and characterizes it. Already the main building blocks of the proposed architecture have been implemented by the authors and there is an on going work especially on the Modeler agent rule engine integration and simulation.

Acknowledgments

The work presented by this paper is supported by the FP7 European Research Project SCOVIS (g. a. no. 216465) (<http://www.scovis.eu>).

References

1. WSBPEL standard, <http://docs.oasis-open.org/wsbpel/2.0/OS/wsbpel-v2.0-OS.html>
2. XLANG, <http://msdn.microsoft.com/en-us/library/aa577463%28BTS.20%29.aspx>
3. WSFL, <http://xml.coverpages.org/wsfl.html>
4. WSCL, <http://www.w3.org/TR/wscl10/>
5. Bohn, H., Golatowski, F., Timmermann, D.: Dynamic device and service discovery extensions for WS-BPEL. In: International Conf. on Service Systems and Service Management, pp. 1–6 (2008)

6. Sollazzo, T., Handschuh, S., Staab, S., Frank, M., Stojanovic, N.: Semantic Web Service Architecture—Evolving Web Service Standards toward the Semantic Web. In: Proc. of the 15th International FLAIRS Conference (2002)
7. Xie, N., Cao, C., Ma, B., Zhang, C., Si, J.: Ontology-Based Web Agents Using Concept Description Flow. In: Knowledge-Based Intelligent Information and Engineering Systems 8th International Conference, KES 2004, Proc. Part II, Wellington, New Zealand (2004)
8. Montebello, M., Abela, C.: DAML enabled Web Services and Agents in the Semantic Web, NetObject Days. In: Chaudhri, A.B., Jeckle, M., Rahm, E., Unland, R. (eds.) NODeWS 2002. LNCS, vol. 2593, pp. 46–58. Springer, Heidelberg (2003)
9. Giovanni, C.: WADE Tutorial, AAMAS 2008 (2008)
10. JADE environment, <http://jade.tilab.com/>
11. Briot, J.P., Muñoz-Meléndez, A., Cardon, A.: Adaptability and Embodiment Using Multi-Agent Systems. In: Mařík, V., Štěpánková, O., Krautwurmová, H., Luck, M. (eds.) ACAI 2001, EASSS 2001, AEMAS 2001, and HoloMAS 2001. LNCS (LNAI), vol. 2322, pp. 121–150. Springer, Heidelberg (2002)

Speaker Verification System Using a Hierarchical Adaptive Network-Based Fuzzy Inference Systems (HANFIS)

Zohreh Soozanchi-K.¹, Mohammad-R. Akbarzadeh-T.², Mahdi Yaghoobi¹,
and Saeed Rahati¹

¹ Department of Artificial Intelligence, Islamic Azad University, Mashhad Branch, Iran

² Center for Applied Research on Soft Computing and Intelligent Systems,

Ferdowsi University of Mashhad, Iran

z.soozanchi@gmail.com

Abstract. We propose the use of a hierarchical adaptive network-based fuzzy inference system (HANFIS) for automated speaker verification of Persian speakers from their English pronunciation of words. The proposed method uses three classes of sound properties consisting of linear prediction coefficients (LPC), word time- length, intensity and pitch, as well as frequency properties from FFT analysis. Actual audio data is collected from fourteen Persian speakers who spoke English. False acceptance ratio and false rejection ratio as are evaluated for various HANFIS trained with different radius. Results indicate that vowel sounds can be a good indicator for more accurate speaker verification. Finally, the hierarchical architecture is shown to considerably improve performance than ANFIS.

Keywords: ANFIS, Speaker verification, LPC, FFT, intensity and pitch coefficients, HANFIS.

1 Introduction

The computer industry controls most informational, financial and security systems. Currently, the research community has considered properties of face, sound, fingerprint, and iris for new identification schemes [1]. In addition to above approaches, sound is an important information source, is very simple for the user, and can potentially provide high speed recognition. To date, sound recognition for security purposes and with the use of ANFIS in particular, has been largely neglected. Few articles have applied neural networks and considered sound samples directly as input to the network, without any pre-processing. In this paper, we advocate the use of a hierarchical ANFIS (HANFIS) for fast clustering and identification of sound. Initially, several features are produced from sound of individuals after omission of noise and silence [2], and validation function is accomplished with these features by ANFIS. The experimental result of this paper which is practical represented by final Table at the end and it is MATLAB programme.

2 Sound Verification System

It is well known that people can be generally identified by their voices, not by the message sent by them. Basically, there are two kinds of voice based recognition or speaker recognition: speaker identification and speaker verification [3]. Speaker recognition is further divided into two categories, which are text dependent and text-independent [4]. The method that use in this paper is based on the concept of speaker verification since the objective in the access control is to accept or reject a person to enter a specific building or room. In general the speaker verification system for accessing control makes four possible decisions [5]. The accuracy of the access control system is then specified based on the rate in which the system makes decision to reject the authorized person and to accept the unauthorized person. The quantities to measure the rate of the access control accuracy to reject the authorized person is then called as false rejection rate (FRR) and that to measure the rate of access control to accept the unauthorized person is called to as false acceptance rate (FAR). Mathematically, both rates are expressed as percentage using the following simple calculations [3]:

$$FRR = \frac{NFR}{NAA} * 100\% \quad (1)$$

$$FAR = \frac{NFA}{NIA} * 100\% \quad (2)$$

NFR and NFA are the numbers of false rejections and false acceptance respectively, while NAA and NIA are the number of the authorized person attempts and the numbers of impostor person attempts. For achieving high security of the door access control system, it is expected that the proposed system will have both low FRR and low FAR. In order to give a definite answer of access acceptance or rejection, a threshold is set. When degree of similarity between a given voice and the model is greater than threshold, the system will accept the access; otherwise the system will reject the person to access the building/room.

3 Plan Executions by Hierarchical ANFIS

ANFIS, proposed by Jang [6], is an architecture which functionally integrates the interpretability of a fuzzy inference system with adaptability of a neural network. ANFIS structure is a weightless multi-layer array of five different elements [7].

In this research ANFIS-based speaker model is developed using Fuzzy Logic Toolbox of MATLAB and design of the ANFIS structure is done by determining premise parameters. Here the subtractive clustering method is used with different radius parameters. Once the premise parameters are obtained, the ANFIS model is trained by using hybrid learning algorithm.

In this case, we can use some ANFIS orders as shown in Fig. 1. We propose the use of a hierarchical adaptive network-based fuzzy inference system (HANFIS) for automated speaker verification of Persian speakers from their English pronunciation of words. In this way, we should allocate whole responses from each ANFIS as new recent data to other ANFIS input for its instruction. Here, we give answer to test process of HANFIS with different pitch and intensity features, frequency features and LPC coefficients as input to the other network, then evaluate program with new data.

As it was mentioned, we can obtain better results with the change of parameters and ANFIS's radius. Here, the percent of errors is zero with the change of radius. You can see combination of LPC features, frequency, pitch and intensity features in Table 4. The main important point in here is that the value of FAR should be close to zero, until unauthorized recognized better.




Fig. 1. HANFIS architecture

4 Data Collection

Whole used data in this article is statement of word in English language by Persian speaker. At first, for gathering data by speaker, 20 samples from considered words were stated and recorded, then the silence between pronounced words at each sentence was omitted by PRAAT software and stored at different separated files (PRAAT is a program for doing phonetic analyses and sound manipulations. It can be downloaded from www.praat.org). During arrangement, it is necessary to equalize length of whole samples because signal length in different speeches of fixed sentences is not the same. So, in this field, length of whole samples was set equal to longest sample, and empty space was loaded by zero.

4.1 Feature extraction

Feature extraction is the main important processing, which in fact is the conversions programme of raw signal to feature vector for uses classification. Features are quantities which resulted from sound pre-processing and can be used for sound representation. In proposed method, every sentence was divided to 20 sub words and different features, and related specifications of sub words are extracted and used in speech recognition.

Three groups of separated features are as follows:

LPC is a statistical method by which that we can calculate some different coefficients for forecasting signals with high precision. In this method, we sampling from each signal and different whole coefficients has been settled by mathematical methods. At first, for extraction of LPC coefficients, we should normalized each signal of sample sentence between 0 and 1, then divide them to 10 blocks and extract suitable

LPC coefficient for every block by MATLAB software which we have used length feature of word as sound features.

Frequency features of sound signals are extracted by Fourier transform. So, we divide and separate each sentence to 10 blocks for better extraction of frequency features, then extract their coefficients by Fourier transform.

Pitch and intensity are two features of sound. Sound pitch depends on its frequency. The intensity of sound measured by db, and is magnitude of frequency at air-pressure which is caused by sound waves. After extraction of these features by PRAAT software, we can store them in a file and considered them as network input.

5 Results and Discussion

In this research, we recorded the sound of 14 persons and every individual repeated one word 20 times and this function was done in environment with acoustic walls due to noise omission. Also, according to accomplished researches it was cleared that acoustic features of individuals when saying some letters containing vowel sounds would be distinguished better and it's because of pitch feature that peer better in vowels. We considered 2 individuals as authorized and the other 12 individuals as unauthorized among these 14 persons. Also, we have used 4 sounds of new individuals as unauthorized whose acoustic features were not in training data. Obtained results with ANFIS in different stages of coefficients training could be observed in Table 1, 2 and 3. We changed radius value of ANFIS then compared results in every Table. In these Tables, P1 and P2 show authorized individuals and P3- P14 show unauthorized individuals and P15-P18 are new unauthorized persons who weren't in data training.

According to obtained results in our proposed method, it was clear that represented method producing results towards the others about gathered data. Also, obtained present results have been trained and tested with limited data. We can use more data for better examination and obtain optimum results in this field.

Table 1. Comparison between the LPC results with different radius by ANFIS

	Radius=0.5 ANFIS		Radius=1 ANFIS	
	FAR	FRR	FAR	FRR
P1	-	10	-	10
P2	-	10	-	20
P3	2		4	
P4	0		0	
P5	0		0	
P6	0		0	
P7	0		0	
P8	7		3	
P9	0		0	
P10	-	0	-	0
P11	-	19	-	10
P12	0		0	
P13	0		0	
P14	4		4	
P15	0		5	
P16	0		0	
P17	0		0	
P18	1		4	

Table 2. Comparison between the FFT results with different radius by ANFIS

	Radius=0.5 ANFIS		Radius=1 ANFIS	
	FAR	FRR	FAR	FRR
P1	-	20	-	18
P2	-	40	-	40
P3	4		1	
P4	6		1	
P5	0		0	
P6	4		0	
P7	0		0	
P8	4		0	
P9	4		0	
P10	-	10	-	10
P11	-	0	-	10
P12	4		0	
P13	5		5	
P14	0		0	
P15	4		0	
P16	0		0	
P17	0		6	
P18	4		9	

Table 3. Comparison between the Pitch and intensity results with different radius by ANFIS

	Radius=0.5 ANFIS		Radius=1 ANFIS	
	FAR	FRR	FAR	FRR
P1	-	0	-	0
P2	-	0	-	0
P3	0		0	
P4	12		0	
P5	0		0	
P6	0		0	
P7	0		0	
P8	12		4	
P9	0		0	
P10	-	0	-	0
P11	-	0	-	0
P12	0		0	
P13	2		2	
P14	0		0	
P15	1		0	
P16	0		0	
P17	1		8	
P18	20		0	

Table 4. Comparison between the combination of LPC, FFT, pitch and intensity results from hierarchy ANFIS with different radius by HANFIS

	Radius=0.5 HANFIS		Radius=1 HANFIS	
	FAR	FRR	FAR	FRR
P1	-	0	-	0
P2	-	0	-	0
P3	0		1	
P4	0		0	
P5	0		0	
P6	0		0	
P7	0		0	
P8	0		0	
P9	0		0	
P10	-	0	-	0
P11	-	0	-	0
P12	0		0	
P13	0		0	
P14	0		0	
P15	3		0	
P16	0		0	
P17	4		0	
P18	0		0	

References

1. Anonymous, Door-access-control System Based on Finger-vein Authentication. Hitachi Review (2004), <http://www.hitachi.com/rev/archive/2004>
2. Xiong, Y.: Speaker Identification System using HMM and Mel Frequency Cepstral Coefficient. Pattern Recognition and Clustering (May 2006)
3. Campbell, J.P.: Speaker Recognition: a Tutorial. Proc. IEEE, 1437–1462 (1997)
4. Ratha, N.K., Senior, A., Bolle, R.M.: IBM Thomas J. Watson Research Center: Automated biometrics
5. Wahyudi, Astuti, W., Mohamed, S.: Intelligent Voice-Based Door Access Control System Using Adaptive-Network-based Fuzzy Inference Systems (ANFIS) for Building Security. Journal of Computer Science 3(5), 274–280 (2007)
6. Jang, J.S.R., Sun, C.T., Mizutani, E.: Neuro-Fuzzy and Soft Computing. Prentice Hall, Upper Saddle River (1997)
7. Jang, J.S.: ANFIS: Adaptive-network-based fuzzy inference system. IEEE Trans. on System, Man, and Cybernetics 23, 665–685 (1993)

Case-Based Decision Support in Time Dependent Medical Domains

Stefania Montani

Dipartimento di Informatica, Università del Piemonte Orientale, Alessandria, Italy

Abstract. Medical applications often require to capture the observed phenomenon dynamics, in order to properly support physicians' decision making. Case-based Reasoning (CBR), and more specifically case-based retrieval, is recently being recognized as a valuable decision support methodology in these domain. However, adopting CBR in this field is non trivial, since the need for describing the process dynamics impacts both on case representation and on the retrieval activity itself. In this work, we survey different methodologies introduced in the literature in order to implement medical CBR applications in time dependent domains, with a particular emphasis on *time series* representation and retrieval. Among the others, a novel approach, which relies on Temporal Abstractions, is analysed in depth.

1 Introduction

Several real world applications require to capture the evolution of the observed phenomenon over time, in order to describe its behaviour, and to exploit this information for future problem solving. Despite the fact that this requirement characterizes many different domains, like e.g. financial and physics ones, it appears to be particularly relevant in medical applications, where the physician typically needs to recall the clinical history that led the patient to the current condition, before prescribing a therapy. In several medical applications, (many) process features are naturally collected in the form of **time series**, either automatically generated and stored by control instruments (as e.g. in Intensive Care Unit monitoring), or obtained by listing single values extracted from temporally consecutive situations (as e.g. the series of glycated haemoglobin values, measured on a diabetic patient once every two months). Analysing long and complex time series of measurements at a screen or on paper can be tedious and prone to errors for physicians; an automated decision support strategy is therefore strongly needed in these domains.

Case-based Reasoning (CBR) [1], a reasoning paradigm that exploits the knowledge collected on previously experienced situations, known as *cases*, is recently being recognized as a valuable knowledge management and decision support methodology in time dependent applications (see e.g. [8]). Cases are composed of a problem description part, and a solution part. CBR operates by *retrieving* past cases that are similar to the current one and by *reusing* past successful solutions after, if necessary, properly *adapting* them; the current case can

then be *retained* and put into the system knowledge base, called the *case base*. *Purely retrieval* systems are also very valuable decision support tools, especially when automated adaptation strategies can hardly be identified, as in medicine. As a matter of fact, various case-based retrieval works dealing with cases with time series features have been recently published, in the medical domain (see section 2), as well as in different ones (see e.g. 7).

However, adopting case-based retrieval can be non trivial in these applications 8. In particular, pre-processing techniques are required in order to simplify the representation of cases with time series features, and to optimize the retrieval activity. Most of the approaches proposed in the literature to this end are founded on the common premise of **dimensionality reduction**, which allows to reduce memory occupancy, still capturing the most important characteristics of time series. Dimensionality is typically reduced by means of a **mathematical transform**, able to preserve the distance between two time series (or to underestimate it). On the other hand, our research group has recently proposed 8,3 to exploit a different technique for dimensionality reduction, namely **Temporal Abstractions** (TA) 11. These techniques will be described in section 2. Finally, section 3 will address conclusions and future work.

2 Methods for Dimensionality Reduction

Mathematical methods. A wide literature exists about time series dimensionality reduction (see the survey in 6). Dimensionality is typically reduced by adopting a mathematical transform that preserves the distance between two time series (or underestimates it). Widely used transforms are the Discrete Fourier Transform (DFT) 2, and the Discrete Wavelet Transform (DWT) 4.

The choice of the most cost-effective transformation to apply should be done on the basis of the application at hand. For instance, dimensionality reduction in medical CBR by resorting to DFT has been proposed in 9, where the application domain is the one of haemodialysis. In 9 only the first DFT coefficients are kept for similarity calculation, and Retrieval is implemented as a multi-step procedure. In particular, the most similar cases with respect to each one of the time series features, considered individually, are first extracted. To guarantee that retrieved cases have the required level of similarity on each time series feature, the intersection of the sets of returned cases is then computed. On the result set, global distance is finally calculated, by computing a weighted average; returned cases are ordered in terms of overall distance.

A significant application of DWT-based dimensionality reduction in medical CBR, on the other hand, can be found in 10. The tool is able to classify Respiratory Sinus Arrhythmia (RSA), i.e. the respirations affecting the heart rate: persons do sometimes have physiological or psychological (e.g. stress-related) disorders, that appear as dysfunctions in the RSA patterns. The system uses CBR as the method for classification of dysfunctional patterns within the RSA. In the system, RSA time series are first dimensionally reduced by means of DWT; then the most similar clusters of previous cases are retrieved from the case base, in

order to support disorder classification. The authors were also able to compare the use of DWT to the one of DFT, since they originally applied DFT to the same domain; in RSA classification, they have obtained an increased retrieval rate when resorting to wavelets.

Temporal Abstractions. Mathematical methods for time series retrieval are widely accepted in the scientific community, and have been experimentally tested in several domains. Nevertheless, they have a number of limitations. For instance, they can be computationally complex, and their output is often not easily interpretable by end users (e.g. physicians). Additionally, they work well with signals with relatively simple dynamics, but they can fail to characterize more complex patterns, as the ones often encountered in medicine. The study of an alternative way to deal with time series retrieval is therefore well justified.

The idea of relying on Temporal Abstractions (TA) methods for time series dimensionality reduction and retrieval support, originally introduced by our research group [8], starts to be reported in the literature, especially dealing with medical applications (see e.g. [53]). TA is an Artificial Intelligence methodology able to move from a *point-based* to an *interval-based* representation of the data, where: (i) the input points are the elements of the discretized time series; (ii) the output intervals aggregate adjacent points sharing a common behavior, persistent over time. More precisely, the method described above should be referred to as *basic* TA. Basic abstractions can be further subdivided into *state* TA and *trend* TA. *State* TA are used to extract episodes associated with *qualitative levels* of the monitored feature, e.g. low, normal, high values; *trend* TA are exploited to detect specific *patterns*, such as increase, decrease or stationary behaviour, from the time series.

In particular, we are currently working at the definition of a time series retrieval framework, which allows for *multi-level abstractions*, according to two *dimensions*, namely a taxonomy of (trend or state) TA symbols, and a variety of time granularities. TA symbols can be organized in a taxonomy, in order to provide different levels of detail in the description of episodes. On the other hand, time granularities allow one to describe episodes at increasingly more abstract levels of temporal aggregation. Our framework also takes advantage of a forest of *multi-dimensional orthogonal index structures*, allowing for early pruning and focusing during the retrieval process, each of which orthogonally spans both the time and the symbol taxonomy dimensions. The root node of each index structure is represented by a string of symbols, defined at the highest level in the symbol taxonomy and in the time granularity taxonomy. An example, taking as a root the *I* (increase) symbol, is provided in figure 1. Here, the root node *I* is refined along the time dimension from the 4 hours to the 2 hours granularity, so that the nodes *II*, *IS* and *SI* stem from it (provided that *IS* and *SI* generalize to *I* when switching to a coarser time granularity level, see figure 1; *S* stands for stationary). Orthogonally, each node can be specialized in the symbol taxonomy dimension; for instance, *SI* can be specialized into *S I_W* (stationary, weak increase), *S I_S* (stationary, strong increase).




Fig. 1. An example multi-level orthogonal index structure

Users are allowed to issue queries at any level of detail. To answer the query, we first progressively generalize it in the symbol taxonomy dimension, while keeping time granularity fixed. Then, we generalize the query in the time dimension as well. Generalization allows us to identify the root of one index in the forest. Following the generalization steps backwards, we are then able to descend along the index, until we reach the node which fits the original query time granularity. If an orthogonal index stems from this node, we can descend along it as well, always following the query generalization steps backwards. We stop when we reach the same detail level in the symbol taxonomy as in the original query. If the query detail level is not represented in the index, because the index is not complete, we stop at the most detailed possible level. We then return all the cases indexed by the selected node.

Interactive and progressive query relaxation and refinement are supported as well in our framework. Query relaxation (as well as refinement) can be repeated several times, until the user is satisfied with the obtained results.

Some first experiments on the use of our framework were executed in the field of haemodialysis. In this domain, we compared our TA-based approach with the DFT-based one adopted in [9]. The experiments proved that, when focusing on the shape of time series (rather than on values), the use of TA provides more reliable results. Moreover, our TA-based method is computationally more efficient. The interested reader may find additional details about the framework and about the experiments in [3].

3 Conclusions

Time series retrieval is a critical issue in all medical domains in which the observed phenomenon dynamics have to be dealt with. Case-based Reasoning (CBR), and more specifically case-based retrieval, is recently being recognized as a valuable decision support methodology in these domains, as testified by the growing number of works in the field. In this paper, we have analyzed the main methodologies used to implement case-based retrieval as a decision support strategy in these applications. Among the others, we have described a framework

in which time series dimensionality is reduced by means of TA. The framework supports multi-level abstractions, both along the time dimension, and along a symbol taxonomy one, thus increasing the flexibility of retrieval. Query answering is interactive and is made faster by the use of orthogonal index structures. In our opinion, flexibility and interactivity represent a relevant advantage of our approach with respect to more classical techniques, in which end users are typically unable to intervene in the retrieval process, that often operates in a black-box fashion. In the future, we plan to extensively test our framework in several different application domains, thus validating its significance, and studying ways of making it more and more efficient and usable.

References

1. Aamodt, A., Plaza, E.: Case-based reasoning: foundational issues, methodological variations and systems approaches. *AI Communications* 7, 39–59 (1994)
2. Agrawal, R., Faloutsos, C., Swami, A.N.: Efficient similarity search in sequence databases. In: Lomet, D. (ed.) *Proc. 4th Int. Conf. of Foundations of Data Organization and Algorithms*, pp. 69–84. Springer, Berlin (1993)
3. Bottrighi, A., Leonardi, G., Montani, S., Portinale, L., Terenziani, P.: Intelligent data interpretation and case base exploration through temporal abstractions. In: Bichindaritz, I., Montani, S. (eds.) *ICCBR 2010. LNCS (LNAI)*, vol. 6176, pp. 36–50. Springer, Heidelberg (2010)
4. Chan, K.P., Fu, A.W.C.: Efficient time series matching by wavelets. In: *Proc. ICDE 1999*, pp. 126–133. IEEE Computer Society Press, Washington (1999)
5. Funk, P., Xiong, N.: Extracting knowledge from sensor signals for case-based reasoning with longitudinal time series data. In: Perner, P. (ed.) *Case-Based Reasoning in Signals and Images*, pp. 247–284. Springer, Heidelberg (2008)
6. Hetland, M.L.: A survey of recent methods for efficient retrieval of similar time sequences. In: Last, M., Kandel, A., Bunke, H. (eds.) *Data Mining in Time Series Databases*. World Scientific, London (2003)
7. Jaere, M.D., Aamodt, A., Skalle, P.: Representing temporal knowledge for case-based prediction. In: Craw, S., Preece, A.D. (eds.) *ECCBR 2002. LNCS (LNAI)*, vol. 2416, pp. 174–188. Springer, Heidelberg (2002)
8. Montani, S., Portinale, L.: Accounting for the temporal dimension in case-based retrieval: a framework for medical applications. *Computational Intelligence* 22, 208–223 (2006)
9. Montani, S., Portinale, L., Leonardi, G., Bellazzi, R., Bellazzi, R.: Case-based retrieval to support the treatment of end stage renal failure patients. *Artificial Intelligence in Medicine* 37, 31–42 (2006)
10. Nilsson, M., Funk, P., Xiong, N.: Clinical decision support by time series classification using wavelets. In: Chen, C.S., Filipe, J., Seruca, I., Cordeiro, J. (eds.) *Proc. Seventh International Conference on Enterprise Information Systems (ICEIS 2005)*, pp. 169–175. INSTICC Press (2005)
11. Shahar, Y.: A framework for knowledge-based temporal abstractions. *Artificial Intelligence* 90, 79–133 (1997)

Solving the Capacitated Vehicle Routing Problem and the Split Delivery Using GRASP Metaheuristic

Joseph Gallart Suárez and Manuel Tupia Anticona

Engineering department, Pontificia Universidad Católica del Perú, Av. Universitaria 1801,
San Miguel, Lima, Perú
{jgallart,tupia.mf}@pucp.edu.pe

Abstract. This paper presents two GRASP metaheuristic algorithms for the vehicle routing problem, considering the capacity and shared demand of the customers. In this paper the solution obtained is compared with a greedy solution and two hybrid solutions (greedy and random). The results obtained show that the GRASP algorithm obtains a better quality solution for this kind of problem.

Keywords: Capacitated Vehicle Routing Problem, Split Delivery Vehicle Routing Problem, Metaheuristic, GRASP.

1 Introduction

The Vehicle Routing Problem (VRP) is probably the best known node routing problem nowadays. It is generally defined as a graph $G = (V, D)$ with a set of nodes $V = \{v_0, v_1, v_2 \dots v_n\}$ a set of edges D , where v_0 represents the depot, with K identical delivery vehicles with Q capacity, the other nodes represent the consignee customers with a demand q_i and each edge (i, j) has a symmetrical cost $c_{ij} = c_{ji}$ [1].

The problem consists in determining a group of K routes for K vehicles, where total cost is minimal, beginning and ending in a depot; such that each node is visited just one time and each delivery vehicle does not surpass its Q capacity [2].

A particular variant of the previously explained scenery is to add delivery vehicle capacity as a restriction (CVRP), besides, in this paper, split delivery is added as a second restriction (SDVRP); this last consideration applies when the consignee customer has bigger demand than the delivery vehicle capacity, thus a consignee customer can be visited by more than one delivery vehicle.

Since the introduction of the proposed problem by Dantzig and Ramser [3], the evolution of solutions, both exact and heuristic, has improved notoriously, among heuristic solutions the works of Laporte, Mercure and Nobert [4]; Fischetti, Toth and Vigo [5]; and Fisher [6] stand out, proposing Ramification and quotation algorithms. We have to mention that Fisher's method solves up to 71 destiny nodes (customers). However, this demands a big computational cost [2].

In the exact ramification and cut algorithm category the work of Cornuéjols, Fonlupt and Naddef [7]; Naddef and Rinaldi [8]; Toth and Vigo [9] stand out, however, this kind of algorithm depends a lot on the particular structure of the VRP since it can complicate the solution if cuts are applied to a structure that does not favor the algorithm [10].

Within the most outstanding heuristic constructive methods we have Clarke and Wright's saving algorithm [11], and the saving algorithm based on matching, as the works of Desrochers and Verhoog [12] and Altinkemer and Gavish[13].

In the category of algorithms with heuristic insertion methods the works by Mole and Jameson [14]; Christofides, Mingozzi and Toth [15]; Solomon [16] stand out, these algorithms start with initially empty routes and evaluate the best form to insert a node in any route iteratively, remaining with a pair (node, route) that represents the best insertion.

Nowadays, we have metaheuristic methods that obtain very good solutions for the proposed problem, such is the case of the ant colony algorithm [17], genetic algorithms [18] or the Tabu search algorithm [19].

In this document, we present two GRASP metaheuristic algorithms that allow us to solve the previously mentioned problem and 3D bin packing problem (3D-BPP) within the bins of the delivery vehicles, besides, we try to improve the solution, optimizing it through a 2-Opt local search algorithm. The combination of the two previously mentioned problems constructs a real transport and package packing problem.

This paper is organized as follows: in section 2 the algorithms we used are described, in section 3 the tests and the results obtained are shown, finally, the conclusions obtained and future work are shown in section 4.

2 Description of the Algorithms

2.1 GRASP Algorithm

The GRASP algorithm is an aleatory and adaptative short-sighted search procedure that helps us to find good quality solutions for combinatory optimization problems [20]. This algorithm is a multistart method where each iteration contains two phases: construction and improvement.

In the construction phase, a solution is constructed by relaxing the greedy criteria. In order to do so, an element is iteratively added in each step, the election of each element (S_i) is determined by three elements: the objective function (f_o), the viability range (E) and the restricted list of candidates, which is defined by three factors: the relaxation constant (α), the best value of the objective function (β) and the worst value of the objective function (τ) [20]. The restricted list of candidates (RCL) to minimize cases and to maximize the objective function respectively is detailed in equations 1 and 2.

$$\text{RCL} = \{S_i \in E: \beta \leq f_o(S_i) \leq \beta + \alpha(\tau - \beta)\}$$

Equation 1. Restricted list of candidates in case minimization of the objective function is searched

$$\text{RCL} = \{S_i \in E: \tau + \alpha(\beta - \tau) \leq f_o(S_i) \leq \beta\}$$

Equation 2. Restricted list of candidates in case maximization of the objective function is searched

An element is extracted aleatorily out of this group of candidates, thus forming a possible solution.

In the improvement phase the solution generated by the construction phase is optimized; in order to do so, an iterative search is performed replacing successively the present solution by one that is in the group of neighborhood solutions. When a solution that improves the present one is not found, it is said that the solution is locally optimal.

These two phases are repeated several times and the best solution found after all GRASP iterations is considered the best quality solution [20].

In this paper we present 2 GRASP algorithms that are detailed bellow:

2.1.1 First GRASP Algorithm

The first GRASP algorithm provides an initial solution for the capacitated vehicle routing problem, taking split delivery as second restriction. Within the construction phase a second GRASP algorithm is applied that solves the 3D bin packing problem which helps us to optimize packing in the vehicle, considering vehicle load weight and the order of delivery of the consignees.

This algorithm receives two entry parameters: a graph (conformed by the starting node and the consignee nodes) and the delivery vehicle list; the algorithm returns the delivery route to be followed by each vehicle.

For the construction of the first GRASP algorithm we use an objective function that has three entry parameters: the traveled distance of the edge, the traffic constant (a value between 1 and 5) and the cost function of the second GRASP algorithm¹. The objective function to be minimized is shown in equation 3.

$$F_{(objective\ GRASP-1)} = \frac{Distance * TrafficConstant}{F_{(cost\ GRASP-2)}}$$

Equation 3. Objective function of the first GRASP algorithm

To know the quality of the solution of the algorithm we apply a cost function that has four entry parameters: the summation of the distance traveled by all the vehicles in the route, the average of the traffic constant of all the avenues that form the delivery vehicle route, the quantity of vehicles used to perform the delivery and the average of the cost function of the second GRASP algorithm². The cost function to be minimized is shown in equation 4.

$$F_{(Cost\ GRASP-1)} = \frac{\sum Distance * TrafficConstant * Quantity\ of\ vehicles}{\overline{F_{(cost\ GRASP-2)}}$$

Equation 4. Cost function of the first GRASP algorithm

¹ The cost function of the second algorithm is detailed in section 2.1.2.

² The cost function of the second algorithm is detailed in section 2.1.2.

2.1.2 Second GRASP Algorithm

The second GRASP algorithm is applied in the construction phase of the first GRASP algorithm. This algorithm provides a solution for the 3D bin packing problem. This algorithm receives two entry parameters: the delivery vehicle and the list of packages of the consignees; the algorithm returns the vehicle with the configuration of packed packages according to the space (in three dimensions) and the load weight of such vehicle.

For the construction of the second GRASP algorithm an objective function that has six entry parameters is used: package measurements (width, length and height), the width and the length of the cut region of the vehicle to be examined (*WRVehicleCut*, *LRVehicleCut*) and, finally, vehicle height. The objective function to be maximized is shown in equation 5.

$$F_{(Objective\ GRASP-2)} = \frac{widthPackage * lengthPackage * heightPackage}{WRVehicleCut * LRVehicleCut * heightVehicle}$$

Equation 5. Objective function of the second GRASP algorithm

In this second algorithm a cost function that has 4 entry parameters is applied: the volume used to pack the packages in the vehicle (*Vol Tp*), total storage space volume of the vehicle (*Vol Veh*), the summation of the weight of all the packages packed in the vehicle (*Weight Tp*) and the load weight that the vehicle supports (*Weight Veh*), in this function the priority is to optimize the volume, that is why such factor is doubled. The cost function to be maximized is shown in equation 6.

$$F_{(Cost\ GRASP-2)} = \left(\frac{(Vol\ Tp / Vol\ Veh) * 2 + Weight\ Tp / Weight\ Veh}{3} \right) * 100\%$$

Equation 6. Cost function of the second GRASP algorithm

2.2 Local Search Algorithm: 2-OPT

The 2-OPT algorithm is a local search based improvement algorithm, the procedure to be followed is as follows: two groups of edges $\{x_1, \dots, x_k\}$ and $\{y_1, \dots, y_k\}$ are determined, such that their interchange reduces the solution cost. The arcs “x” must be part of the route and both groups must be disjunct, besides, when eliminating the arcs “x” and adding the “y” arcs they must form a closed route [21].

This algorithm receives three entry parameters: a graph (conformed by the start node and the consignee nodes) and the list of delivery vehicles with the route of the consignees; the algorithm returns a list of vehicles with the optimized routes.

3 Test Execution and Results

The algorithms were implemented in Java and the tests were executed in an Intel Core 2 Duo Computer with 2.0 GHz of processing speed, 2 GB RAM and with a Windows XP SP2 operating system.

A certain quantity of graphs and random nodes were used following the pattern described by Christofides, Mingozzi and Toth for CVRP problems [2]. 300 random graphs with 50, 100 and 200 consignee nodes with random packages and random weight and measures were generated. In Table 1, the parameters used by the two GRASP algorithms are shown.

Table 1. Parameters of the GRASP algorithms

	n=50	n= 100	n=200
#Iterations to GRASP-1	50	50	40
#Iterations to GRASP-2	100	100	80
Relaxation constant to GRASP-1 (α)	0.2	0.2	0.2
Relaxation constant to GRASP-2 (α)	0.8	0.6	0.8

The objective of the tests is to compare the two GRASP algorithms and 2-Opt optimization (acronym G-2OPT) against greedy algorithms and hybrid algorithms (greedy and random) described in Table 2.

Table 2. Acronyms of the algorithms

First algorithm	Second algorithm	Acronym
Greedy algorithm	Greedy algorithm	GG
Greedy algorithm	Random algorithm	GR
Random algorithm	Greedy algorithm	RG

In Tables 3, 4 and 5 the average results obtained by the algorithms according to their execution time (in seconds) and cost (detailed in equation 4) are shown.

Table 3. Result for 50 consignee nodes

	G-2OPT	GG	GR	RG
Average time(s)	306.30	403.20	429.84	532.36
Average cost	2175.29	2283.96	3225.82	4385.58

Table 4. Result for 100 consignee nodes

	G-2OPT	GG	GR	RG
Average time (s)	772.44	925.53	961.17	1216.25
Average cost	17583.11	18516.72	20223.74	29546.84

Table 5. Result for 200 consignee nodes

	G-2OPT	GG	GR	RG
Average time (s)	1849.56	2620	2838	3563.11
Average cost	96193.21	118110.66	144171.91	214684.65

The three tables indicate that the two GRASP algorithms with 2-OPT optimization obtain better quality solutions than the rest of algorithms, in the 3 tables the Greedy algorithm obtains better solutions than the two Hybrid algorithms. In Table 3 the proposed algorithm (G-2OPT) and the greedy algorithm obtain very close solutions (in cost and time), however, for 100 and 200 nodes the proposed algorithm renders a solution of better quality.

We can observe that the Random-Greedy algorithm (RG) presents a higher cost and more execution time than the rest of algorithms, this is because it parts from an random solution, therefore it demands higher cost and more time to find a solution.

4 Conclusions and Future Work

In this work we implemented 2 GRASP algorithms with 2-OPT optimization to solve the capacitated vehicle routing problem, considering the split delivery restriction and solving, besides, the 3D bin packing problem, the combination of both problems let us have a more realistic model of the transport and package loading problem.

Computing tests were performed to study the aplicability and the quality of the solution of the two algorithms proposed, the computational results indicate that the two GRASP algorithms with 2-OPT optimization obtain better solutions than a Greedy algorithm and against two hybrid algorithms.

As future work we have to implement other metaheuristic algorithms, as the Ant Colony and the Genetic algorithms, compare them with the solution that was proposed and to verify if it obtains the best solutions for the proposed problem.

References

1. Prins, C.: A GRASP X Evolutionary Local Search Hybrid for the Vehicle Routing Problem. In: *Bioinspired Algorithms for the Vehicle Routing Problem*, vol. 1, p. 35 (2008)
2. Universidad de Málaga (2009), <http://neo.lcc.uma.es/radi-aeb/WebVRP/>
3. Marinakis, Y., Migdalas, A.: *Operational Research* 7(1), 27–46 (2007)
4. Laporte, G., Mercure, H., Nobert, Y.: An exact algorithm for the asymmetrical capacitated vehicle routing problem. *Networks* 16, 33–46 (1986)
5. Fischetti, M., Toth, P., Vigo, D.: A branch-and-bound algorithm for the capacitated vehicle routing problem on directed graphs. *Operations Research* 42, 846–859 (1994)
6. Fisher, M.L.: Optimal solution of vehicle routing problems using minimum k-trees. *Operations Research* 42, 626–642 (1994)
7. Cornuéjols, G., Fonlupt, J., Naddef, D.: The traveling salesman problem on a graph and some related integer polyhedra. *Mathematical Programming* 33, 1–27 (1985)
8. Naddef, D., Rinaldi, G.: The graphical relaxation: A new framework for the symmetric traveling salesman polytope. *Mathematical Programming* 58, 53–88 (1993)

9. Toth, P., Vigo, D.: *The Vehicle Routing Problem*. Society of Industrial and Applied Mathematics (SIAM), Philadelphia (2002)
10. Claudio Contardo Vera: *Formulación y solución de un problema de ruteo de vehículos con demanda variable en tiempo real, trasbordos y ventanas de tiempo*. Universidad de Chile (2005)
11. Clarke, G., Wright, J.W.: Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research* 12, 568–581 (1964)
12. Desrochers, M., Verhoog, T.W.: A new heuristic for the fleet size and mix vehicle routing problem. *Computers & OR* 18, 263–274 (1991)
13. Altinkemer, K., Gavish, B.: Parallel Savings Based Heuristics for the Delivery Problem. *Operations Research* 39(3), 456–469 (1991)
14. Mole, R.H., Jameson, S.R.: A sequential route-building algorithm employing a generalized savings criterion. *Operations Research Quarterly* 27, 503–511 (1976)
15. Christofides, N., Mingozzi, A., Toth, P.: *The vehicle routing problem*. Wiley, Chichester (1979)
16. Solomon, M.M.: Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research* 35 (1987)
17. Dong, L.W., Xiang, C.T.: Ant Colony Optimization for VRP and Mail Delivery Problems. In: 2006 IEEE International Conference on Industrial Informatics, pp. 1143–1148 (2006)
18. Ombuki-Berman, B., Hanshar, F.T.: Using Genetic Algorithms for Multi-depot Vehicle Routing. *Studies in Computational Intelligence* 161, 77–99 (2009)
19. Wassan, N.A., Wassan, A.H., Nagy, G.: A reactive tabu search algorithm for the vehicle routing problem with simultaneous pickups and deliveries. *Journal of Combinatorial Optimization* 15(4) (2008)
20. Resende, M.G.C., Velarde, J.L.G.: GRASP: Greedy Randomized Adaptive Search Procedures. *Inteligencia Artificial, Revista Iberoamericana de Inteligencia Artificial* 19, 61–76 (2003)
21. Olivera, A.: *Heurísticas para Problemas de Ruteo de Vehículos*. Instituto de Computación, Facultad de Ingeniería, Universidad de la República, Montevideo, Uruguay (2004)

Author Index

- Abdullah, Noraswaliza 57
Adda, Mo 47
Akbarzadeh-T., Mohammad-R. 233
Almeida, Ana 205
Anagnostopoulos, Vasilis 227
Anticona, Manuel Tupia 243
- Barla, Michal 35
Bieliková, Mária 35
Bou Ezzeddine, Anna 187
Bramer, Max 47
- Carneiro, Davide 174
Coelho, Bruno 205
Costa, Ricardo 174
- Davey, Samuel 89
Debenham, John 3
De Giusti, Armando 143
Domagk, Steffi 25
Dwyer, Dan 67
- Egly, U. 215
Eiter, T. 215
El-Mahassni, Edwin 89, 123
- Fukuda, Shinsaku 197
- García-Martínez, Ramón 143
Geva, Shlomo 57
Guedes, Gilleanes Thorwald Araujo 133
- Hada, Ryukichi 197
Hanabata, Norihiro 197
- Ierache, Jorge 143
- Kostková, Katarína 35
- Legg, Jonathan 89
Leonardi, Giorgio 101
Looi, Mark 57
Løvlid, Rikke Amilde 79
Lucht, Martina 25
- MacMahon, Clare 67
Martins, Constantino 205
Micalizio, Roberto 13
Mikami, Tatsuya 197
Mohring, Martin 25
Montani, Stefania 101, 238
Mora, Matias 111
- Návrát, Pavol 165, 187
Neumann, Bernd 155
Neves, José 174
Novais, Paulo 174
Nuzzolo, Giancarlo 13
- Ofoghi, Bahadorreza 67
- Rahati, Saeed 233
Reiterer, A. 215
Rieke-Zapp, D. 215
- Sardis, Manolis 227
Sasaki, Yoshihiro 197
Scala, Enrico 13
Sierra, Carles 3
Solar, Mauricio 111
Soozanchi-K., Zohreh 233
Stahl, Frederic 47
Suárez, Joseph Gallart 243
Suchal, Ján 165
- Terzic, Kasim 155
Torasso, Pietro 13
- Varvarigou, Theodora 227
Vicari, Rosa Maria 133
Vicovac, T. 215
- Xu, Yue 57
- Yaghoobi, Mahdi 233
Yoshimura, Tetsuro 197
- Zeleznikow, John 67