

Preface

The European Society for Artificial Intelligence in Medicine in Europe (AIME) was established in 1986 following a highly successful workshop held in Pavia, Italy, the year before. The aims of AIME are to foster fundamental and applied research in the application of Artificial Intelligence (AI) techniques to medical care and medical research, and to provide a forum for reporting significant results achieved at biennial conferences. In accordance with the latter aim, this volume contains the proceedings of AIME 2001, the eighth conference on Artificial Intelligence in Medicine in Europe, held in Cascais, Portugal, July 1–4, 2001. Previous conferences were held in Marseille (1987), London (1989), Maastricht (1991), Munich (1993), Pavia (1995), Grenoble (1997), and Aalborg (1999). This latter was a joint conference of AIME and ESMDM, the European Society for Medical Decision Making.

The call for papers of AIME 2001 required original contributions regarding the development of theory, techniques, and applications of AI in medicine. Contributions to theory included presentation or analysis of the properties of novel AI methodologies potentially useful in solving medical problems. Papers on techniques described the development or the extension of AI methods and their implementation, and discussed the assumptions and limitations of the proposed methods. Application papers described the implementation of AI systems to solve significant medical problems, and most of them presented an evaluation of the practical benefits of the system proposed.

The call resulted in 79 submissions, covering the areas of knowledge management, machine learning, data mining, decision support systems, temporal reasoning, case based reasoning, planning and scheduling, natural language processing, computer vision, image and signal interpretation, intelligent agents, telemedicine, careflow systems, and cognitive modeling.

All papers were carefully reviewed by at least two independent referees (77% by three referees), belonging to the Program Committee, supported by some additional reviewers. The review form addressed relevance of the paper content to AIME, originality and quality of the research, completeness, and organization of the paper. Eventually, 31 contributions were accepted for oral presentation, and 30 for poster presentation, with a “full paper” acceptance rate of about 39%. Thus, this volume contains 31 full papers and 30 short papers. In addition, the volume contains two keynote lectures written by the invited conference speakers. This year, keynote areas were the communication between agents within healthcare organizations and the sociotechnical approach to the design, implementation, and evaluation of knowledge-based systems. The choice of these areas stems from the recent debate within the medical community about the consequences of lack or default of co-operation among health care professionals. This is one of the main causes of poor care delivery. We think that AIME has the potentiality to take an active role in this debate, devoting efforts to the development

of systems that take into account this medical community need. On the other hand, the 30 years history of AI in medicine shows that effective and efficient implementation of AI systems and, more generally, decision support systems in medicine, is often impaired by poor consideration of the real-world environment where such systems are intended to work.

We finish by thanking all those who contributed to the success of AIME 2001: the authors, the program committee members together with the additional reviewers, the local organizing committee members, the invited speakers Enrico Coiera from Australia and Jos Aarts from The Netherlands, the satellite workshops organizers, Peter Lucas (Bayesian Models in Medicine) and Stephen Rees (Computers in Anaesthesia and Intensive Care: Knowledge-Based Information Management), the tutorials' presenters Christoph Schommer (Application of Data Mining in Medicine), Jeremy Wyatt (Knowledge Management and AI in Medicine: What's the Link?), Gabriela Guimaraes (Unsupervised Neural Networks for Knowledge Discovery in Medicine), and Dan Steinberg (Multivariate Adaptive Regression Splines).

Last but not least we thank all the Institutions that sponsored the conference, namely IPE, Investimentos e Participações Empresariais, SA, Fundação para a Ciência e Tecnologia, and Fundação Calouste Gulbenkian.

April 2001

Silvana Quaglini
Pedro Barahona
Steen Andreassen

Organization

Scientific Program Chair:	Silvana Quaglini (Italy)
Organizing Chair:	Pedro Barahona (Portugal)
Tutorials and Workshops Chair:	Steen Andreassen (Denmark)

Program Committee

Wil Van der Aalst (The Netherlands)	Jim Hunter (United Kingdom)
Stig K. Andersen (Denmark)	Elpida Keravnou (Cyprus)
Steen Andreassen (Denmark)	Cristiana Larizza (Italy)
Pedro Barahona (Portugal)	Nada Lavrac, Ljubljana (Slovenia)
Robert Baud (Switzerland)	Johan van der Lei (The Netherlands)
Riccardo Bellazzi (Italy)	Leonard Leibovici (Israel)
Jan van Bommel (The Netherlands)	Silvia Miksch (Austria)
Marc Berg (The Netherlands)	Marc Musen (USA)
Enrico Coiera (Australia)	Stelios Orfanoudakis (Greece)
Carlo Combi (Italy)	Silvana Quaglini (Italy)
Michel Dojat (France)	Alan Rector (United Kingdom)
Rolf Engelbrecht (Germany)	Yuval Shahar (Israel)
John Fox (United Kingdom)	Mario Stefanelli (Italy)
Douglas Fridsma (USA)	Mário Veloso (Portugal)
Catherine Garbay (France)	Jeremy Wyatt (United Kingdom)
Barbara Heller (Germany)	Blaz Zupan (Slovenia)
Werner Horn (Austria)	

Additional Reviewers

Ivano Azzini	Giovanni Magenes	Andreas Seyfang
Robert Kosara	Laura Maruster	Ton Weijters
Matjaž Kukar	Marko Robnik-Šikonja	Xenophon Zabulis
Peter Lucas	Patrick Ruch	

Organizing Committee

Pedro Barahona	Joaquim Aparício	Gabriela Guimarães
Mário Veloso	João Moura Pires	

Workshops

Bayesian Models in Medicine

Co-chairs: Peter Lucas, University of Aberdeen, United Kingdom

Linda van der Gaag, Utrecht University, The Netherlands

Ameen Abu-Hanna, University of Amsterdam, The Netherlands

Computers in Anaesthesia and Intensive Care: Knowledge-Based Information Management

Chair: Stephen Rees, Aalborg University, Denmark

Co-chairs: Silvia Miksch, Vienna University of Technology, Austria

Michel Dojat, INSERM, France

Jim Hunter, University of Aberdeen, United Kingdom

Tutorials

Application of Data Mining in Medicine

Christoph Schommer, IBM German Development Laboratory, Germany

Knowledge Management and AI in Medicine: What's the Link ?

Jeremy Wyatt, University College, London, United Kingdom

Unsupervised Neural Networks for Knowledge Discovery in Medicine

Gabriela Guimarães, Universidade Nova de Lisboa, Portugal

Multivariate Adaptive Regression Splines

Dan Steinberg, Salford Systems, San Diego, CA, USA

Balancing Reactivity and Social Deliberation in Multi-Agent Systems – A Short Guide to the Contributions

Markus Hannebauer¹, Jan Wendler², and Enrico Pagello^{3,4}

¹ GMD – German National Research Center for Information Technology,
Research Institute for Computer Architecture and Software Technology (FIRST)
Planning and Optimization Laboratory, Kekuléstr. 7, D-12489, Germany

markus.hannebauer@gmd.de

² Humboldt University Berlin, Department of Computer Science,
AI Laboratory, Unter den Linden 6, D-10099 Berlin, Germany

wendler@informatik.hu-berlin.de

³ University of Padua, Department of Electronics and Informatics,
Intelligent Autonomous Systems Laboratory,
Via Gradenigo 6/a, I-35131 Padova, Italy

epv@dei.unipd.it

⁴ LADSEB-CNR, Padova, Italy

Abstract. The focus of “Balancing Reactivity and Social Deliberation in Multi-agent Systems” is the right balance between the two extremes of pure reactivity and in-depth social deliberation in the context of collaborative work in multi-agent systems. This article briefly motivates this problem and provides a short guide to the contributions contained in this volume.

1 Covered Topics

Today’s envisioned applications of intelligent systems in general and multi-agent systems in particular confront researchers and developers with the difficulty of finding the right balance between reactive and socially deliberative behavior. Reactive systems are capable of adapting very quickly to unforeseen changes in the environment and are hence said to be more robust and efficient. On the other hand, they usually lack the necessary overview to produce behavior that can compete with the results of in-depth reasoning techniques. In contrary to that, socially deliberative systems allow for exploiting environmental information and coordination mechanisms to build up through-thought individual and even team-oriented strategies. Though their problem solving results are usually much better than that of reactive systems, deliberative systems are much more susceptible to dynamic environments and often lack the potential for real-time computation.

This book focuses on theoretical, technical and practical work on balancing between these two extremes in the context of collaborative work in multi-agent systems. The call for contributions for this book directed the attention to the following topics of interest.

- Extension of reactive systems by cooperation
- Teaching deliberative systems reactivity and real-time
- Efficient coordination, cooperation and organization approaches
- Anytime approaches and algorithms
- Design and evaluation of hybrid multi-level agent architectures
- Short-term, medium-term and long-term intentionality
- Enriching group behavior by environmental, opponent and social models
- Individual and social adaptivity

As application areas we had encouraged practical contributions to be directed at the following.

Multi-agent problem solving Real-world problems often require distributed solving strategies, because of natural distribution, social competence and efficiency matters. Multi-agent approaches to such problems are said to be more robust than monolithic systems, but usually entail worse solutions. How can this be overcome by better balancing between reactive and deliberative behavior?

RoboCup RoboCup has proven to be a great and challenging benchmarking scenario both to Robotics and Artificial Intelligence. In RoboCup reactivity and real-time are a must, but social deliberation gets more and more important to match the world's leading teams.

2 The Contributions

2.1 General Observations

From the beginning on, the difficulty of defining the notions of “reactivity” and “social deliberation” in a commonly acceptable way was evident. Though all authors had a more or less precise personal understanding of these categories it showed up in a lengthy discussion that a common understanding was quite far away. A first step towards creating a sound taxonomy for this kind of area is made by Iocchi, Nardi and Salerno in their introductory chapter on reactivity and deliberation in multi-robot systems (page 9). The taxonomy is mainly based on the system's cooperative capabilities. It consists of four levels: cooperation, knowledge, coordination and organization which characterize the major features of multi-robot systems. The concepts of reactivity and social deliberation are then assigned to the nodes of the taxonomy. Finally several deployment fields for multi-robot systems are described, examined, and characterized according to the proposed taxonomy. Apart from this introduction, all other authors of this book clarify their understanding of reactivity and social deliberation in the very beginning of their contributions.

As a second important point, even though most of the contributions do not mainly focus on applications, in almost all of them the authors show the relevance of the described techniques or theories by applying them to application areas from multi-agent problem solving. These areas are represented by

testbeds, such as coordinated foraging (Carpin, Ferrari, Pagello, and Patuelli, page 35), tile-world chessboards (Mavromichalis and Vouros, page 53), “Capture the Flag” (Atkin, Westbrook, and Cohen, page 92), “Robot Sheepdogs” (Sigaud and Gérard, page 150) and RoboCup soccer (Bredendfeld and Kobialka, page 111; Behnke and Rojas, page 125; Reis, Lau, and Oliveira, page 175), as well as by *real-world case studies*, such as driver support systems (Malec, page 76), electrical power grid control (Riedmiller, Moore, and Schneider, page 137) and task allocation for cooperative rovers (Bouazid, Hanna and Mouaddib, page 198).

We have compiled the contributions in three parts according to the special focus they put on reactivity, social deliberation and the transition between them. These parts and the contributions belonging to them are briefly described in the next subsections.

2.2 Architectures and Frameworks

The first part deals with architectures and frameworks that comprise reactive components as well as deliberative components in a unique system. In the work of Carpin et al. and Malec reactive and deliberative components are clearly distinguishable, whereas Mavromichalis and Vouros or Atkin et al. do not represent reactivity and deliberation by distinct parts of their frameworks but by hierarchical planning concepts that allow for a fluent transition between reactive and deliberative behavior.

Carpin, Ferrari, Pagello and Patuelli (page 35) analyze the problem of balancing reactivity and social deliberation in the case of cooperative multi-robot systems. They outline the issues which need to be coped with to solve this problem and introduce a balancing method based on the “map focus” concept. This concept couples the reactive and the deliberative module of their architecture. The basic idea of the “map focuser” is to compute a simplified and localized version of the deliberative module state that can be used by the reactive module. The architecture proposal is supported by a case study of a simulated coordinated multi-robot system that performs a foraging task.

The contribution of Mavromichalis and Vouros (page 53) discusses a framework called ICAGENT that allows an agent designer to specify behaviors that combine reactivity with deliberation. Some internal mental states of the agent are explicit in this framework to allow the agent to decide when to react and when to deliberate. The transition between reactive and deliberative behavior is realized by the amount of external and internal information that is incorporated into a hierarchical planning process. The article addresses balancing only for single agent systems, but the authors state that it can be generalized to multi-agent systems. The framework is evaluated with a tile-world example as testbed.

In his contribution Malec (page 76) discusses the controlled augmentation of predictable reactivity with limited deliberation to preserve hard real-time requirements. The author illustrates his thoughts in terms of the “Generic Layered Architecture”, which in its current status can be used for creating reactive agents

that act in dynamic environments. The behavior of the agents is subject to worst-case guarantees in the sense of temporal predictability. The problem of guaranteeing similar worst-case bounds for deliberation processes is presented and the factors that influence the problem are discussed. The paper also surveys existing approaches to the problem and a number of other layered agent architectures.

Whereas many other well-known architectures conceptually distinguish between reactive and deliberative behavior, Atkin, Westbrook and Cohen (page 92) introduce an agent architecture with a uniform representation of both. To achieve this, the behaviors/actions of the agents are arranged into a hierarchy. All the behaviors are represented in the same manner and have the same interface to a generic planner. Additionally, this agent architecture, called “Hierarchical Agent Control Architecture”, does not distinguish between single-agent and multi-agent behaviors. The architecture and the planner have been applied to the competitive, real time scenario “Capture the Flag”.

2.3 Enhanced Reactivity

This part collects contributions that concentrate on techniques to enhance existing reactive approaches in particular from robotics by social behaviors. All these contributions argue that socially deliberative behavior observed by an external observer can emerge from the interaction of complex reactive behaviors. But they also show that the extended reactive systems do not only “show” deliberative behavior but also dominate their purely reactive ancestors in efficiency and problem solving capability. This is proven by various testbed simulations and case studies.

Bredenfeld and Kobińska (page 111) present a team coordination approach, which is an extension of the behavior-based Dual Dynamics scheme. They specify team behaviors to realize team coordination in multi-agent systems by introducing special blackboard-like “team variables” that are exported and read by the local reactive decision makers. They also show that the set of design, simulation and monitoring tools developed for the Dual Dynamics approach is also suitable for allowing a smooth integration of team behavior with non-team behavior. The approach has been successfully applied to the coordination of soccer playing robots.

The work of Behnke and Rojas (page 125) builds upon the Dual Dynamics scheme. It introduces a temporal hierarchy of behaviors, fast and simple at the bottom of the hierarchy and getting slower and more complex to the top. Each layer in this hierarchy consists of a sensor module, an activation module and an actuator module. All of these modules are subject to the temporal hierarchy, i.e. not only the activation module (the decision maker) but also the sensors and actuators are classified according to their temporal characteristics. The generic approach is substantiated by a precise description of its application to robots in the small size league of RoboCup.

Riedmiller, Moore and Schneider (page 137) extend purely reactive agents by the ability to learn cooperative patterns via reinforcement. Driven by a global optimization goal the agents are forced to establish communication mechanisms.

The authors investigate two different settings of distributed power grid control to assess their approach. The learning agents fail more seldomly in solving the stated problems and produce globally better problem solutions.

Following a nice discussion of the notion of reactivity in the behaviors research community and in the Markov decision process research community, Sigaud and Gérard (page 150) enhance reactive controllers, which are based on the Classifier System formalism, by social roles. Though the performance of the system with handcrafted roles dominates the performance of a set of purely reactive agent without social roles, the latter approach turns out to be more robust to the size of the problem. To cope with this, the authors propose “social reactivity” enabling the reactive agents to dynamically change the social roles initially assigned to them. This approach turns out to be both more efficient and more robust than the system without roles.

2.4 Controlled Social Deliberation

The last part of the book is devoted to contributions that stem from the opposite direction of the spectrum compared to the contributions in the preceding part. Contributions in this part start from deliberative methods and propose techniques to constrain the resources needed by the deliberative reasoning process. Hence, they are trying to “teach deliberative systems reactivity”.

Reis, Lau and Oliveira (page 175) facilitate a classification of the environment observed by an agent into active and strategic situations to balance between reactivity and social deliberation. In active situations that demand a very responsive behavior the agent uses high level and low level skills to directly manipulate the environment. As soon as the situation is classified as strategic more thorough reasoning techniques are used including social ones such as “situation based strategic positioning” and “dynamic role and position exchange”. These concepts are evaluated by their impact on RoboCup soccer matches.

Bouzid, Hanna and Mouaddib (page 198) compare two quite different architectures for task allocation in collaborative problem solving and investigate their responsiveness properties. The first architecture consists of a completely deliberative central task allocator and reactive worker agents. The second architecture comprises deliberative worker agents making local task allocation decisions that are coordinated by a central instance. Decisions in both architectures are made based on Markov decision processes. Evaluation criteria include the overall quality of the task allocation process as well as the communication overhead in the second architecture.

The book closes with a contribution of a speculative yet visionary character. Weiss (page 217) argues that existing agent architectures often tend to inherently limit an agent’s flexibility because they imply a discrete cognitive and social behavior space. He proposes a constraint-centered architectural framework that restricts the deliberative and social behavior by constraint handling and such forms a continuous behavior space. This behavior space is characterized by the two dimensions cognition (reactive vs. deliberative) and sociability

(isolated vs. interactive). Four special conditions are highlighted under which constraint handling seems suitable to guide an agent with controlled social deliberation capabilities. The work additionally provides an extensive review of the current literature.

Artificial Intelligence in Medicine

8th Conference on Artificial Intelligence in Medicine in
Europe, AIME 2001 Cascais, Portugal, July 1-4, 2001,
Proceedings

Quaglini, S.; Barahona, P.; Andreassen, S. (Eds.)

2001, XIV, 469 p., Softcover

ISBN: 978-3-540-42294-5