

Preface

Intelligent agents are one of the most important developments in computer science of the past decade. Agents are of interest in many important application areas, ranging from human-computer interaction to industrial process control. The ATAL workshop series aims to bring together researchers interested in the core/micro aspects of agent technology. Specifically, ATAL addresses issues such as theories of agency, software architectures for intelligent agents, methodologies and programming languages for realizing agents, and software tools for applying and evaluating agent systems. One of the strengths of the ATAL workshop series is its emphasis on the synergies between theories, languages, architectures, infrastructures, methodologies, and formal methods.

This year's workshop continued the ATAL trend of attracting a large number of high quality submissions. In more detail, 71 papers were submitted to the ATAL 2000 workshop, from 21 countries. After stringent reviewing, 22 papers were accepted for publication and appear in these proceedings.

As with previous workshops in the series, we chose to emphasize what we perceive as important new themes in agent research. This year's themes were both associated with the fact that the technology of intelligent agents and multi-agent systems is beginning to migrate from research labs to software engineering centers. As agents are deployed in applications such as electronic commerce, and start to take over responsibilities for their human users, techniques for controlling their autonomy become crucial. As well, the availability of tools that facilitate the design and implementation of agent systems becomes an important factor in how rapidly the technology will achieve widespread use. Consequently, the ATAL 2000 program included two special tracks on *Autonomy — Theory, Dimensions, and Regulation* and *Agent Development Tools*. Besides papers in each of these special tracks, the program also featured two associated panels (organized by Cristiano Castelfranchi and Keith Decker respectively). Another highlight of this year's program was the invited talks by leading exponents of agent research:

ARCHITECTURES	Craig Boutilier	Structured Online and Offline Solution of Decision-Theoretic Planning Problems
THEORIES	Phil Cohen	On Group Communication

It is both our hope and our expectation that this volume will be as useful to the agent research and development community as its six predecessors have proved to be. We believe that ATAL and the *Intelligent Agents* series, of which these proceedings form a part, play a crucial role in a rapidly developing field, by focusing specifically on the relationships between the theory and practice of agents. Only through understanding these relationships can agent-based computing mature and achieve its widely predicted potential.

March 2001

Cristiano Castelfranchi
Yves Lespérance

Workshop Organization

Organizing Committee

Cristiano Castelfranchi University of Siena, Italy
Yves Lespérance York University, Canada

Steering Committee

Nicholas R. Jennings University of Southampton, UK
Yves Lespérance York University, Canada
Jörg P. Müller Siemens Research, Germany
Munindar P. Singh North Carolina State University, USA
Michael Wooldridge University of Liverpool, UK

Program Committee

Chitta Baral	(USA)	Suzanne Barber	(USA)
Michael Beetz	(Germany)	Stefan Bussmann	(Germany)
Lawrence Cavedon	(Australia)	Paolo Ciancarini	(Italy)
Phil Cohen	(USA)	Rosaria Conte	(Italy)
Giuseppe De Giacomo	(Italy)	Keith Decker	(USA)
Frank Dignum	(The Netherlands)	Alexis Drogoul	(France)
Jacques Ferber	(France)	Klaus Fischer	(Germany)
Michael Fisher	(UK)	Stan Franklin	(USA)
Fausto Giunchiglia	(Italy)	Piotr Gmytrasiewicz	(USA)
Keith Golden	(USA)	Barbara Grosz	(USA)
Henry Hexmoor	(USA)	Wiebe van der Hoek	(The Netherlands)
Marc Huber	(USA)	Mark d’Inverno	(UK)
Nick Jennings	(UK)	David Kinny	(Australia)
Sarit Kraus	(Israel)	Michael Luck	(UK)
John-Jules Meyer	(The Netherlands)	Jörg Müller	(Germany)
Anand Rao	(UK)	Murray Shanahan	(UK)
Onn Shehory	(Israel)	Carles Sierra	(Spain)
Munindar Singh	(USA)	Liz Sonenberg	(Australia)
Katia Sycara	(USA)	Milind Tambe	(USA)
Jan Treur	(The Netherlands)	Manuela Veloso	(USA)
Tom Wagner	(USA)	Wayne Wobcke	(UK)
Mike Wooldridge	(UK)	Eric Yu	(Canada)

Additional Reviewers

Jaelson Castro	Rino Falcone	Anuj Goel
Claudia V. Goldman	John Graham	Daniel Gross
Pu Huang	Luke Hunsberger	Manuel Kolp
Dung Lam	Foster McGeary	Wieke de Vries
Wei Yang		

Introduction

Like its predecessors [5,6,3,4,2,1], this volume of the Intelligent Agents series focuses on the relationships between the theory and the practice of intelligent autonomous agents. The volume is divided into eight sections. Sections I and III present work on agent theories. Section II discusses work on development tools and platforms for agents and multiagent systems. Section IV presents work on models of agent communication and coordination. Section V deals with the issue of autonomy and presents additional work on agent coordination. Section VI deals with programming and specification languages for agents. Section VII focuses on issues in agent architectures, such as how to support planning, decision making, and learning. Finally, Sections VIII and IX provide summaries of the discussions held at the two ATAL panels, one on Agent Development Tools and the other on Autonomy — Theory, Dimensions, and Regulation.

Section I: Agent Theories I

Wooldridge and Dunne develop a formal model of the problem of designing an agent that achieves and/or maintains given goals and prove some new results about the complexity of the agent design problem under various assumptions. They look at optimistic agent design, where the agent is only required to achieve/maintain its goals for *some execution* of the specified environment. They also look at disjunctive design, where the agent is allowed to achieve/maintain only one among a set of alternative goals.

Dragoni et al. develop a formal framework for modeling how agents' mental states are updated as a consequence of communication acts. Mental states are modeled as multi-context propositional theories. An abductive account of mental state change is developed for the framework. The use of the framework is illustrated by formalizing the "inform" speech act and showing how its effects on the hearer's mental state can be captured.

Zimmerbaum and Scherl propose a generalized version of a framework for reasoning about how knowledge is affected by action, to deal with concurrent actions and time. The framework, which is based on the situation calculus, captures the knowledge effects of both sensing and non-sensing actions. Indexical knowledge of time and absolute knowledge of time are both handled.

Section II: Agent Development Tools and Platforms

DeLoach and Wood present their MaSE methodology for multiagent system development and a tool that supports its use. The methodology is comprehensive, covering both the analysis and design phases. Diagram-based notations are provided for developing detailed models of the behaviors associated with roles and conversations.

Riley et al. present their *layered disclosure* approach to supporting the debugging and fine tuning of complex agent systems. In layered disclosure, the system designer defines an information hierarchy, which is then used to control the level of detail of the information

provided in response to queries about the agent's internal state and why it acted in a particular way. The implementation and use of the approach in a simulated robot soccer playing system is discussed.

Bryson and Stein propose a meta-methodology for reducing the cost of the proliferation of architectures and methodologies. They propose the use of "idioms" to compare architectures and methodologies. An idiom is a compact, regularized way of expressing a frequently used set of ideas or functionality. As an example, they describe the "Basic Reactive Plan" idiom and compare various agent architectures in terms of how they can express this idiom.

Bellifemine et al. present JADE, a software framework that facilitates the development of agent applications in compliance with the FIPA specifications for multiagent systems. They discuss the advantages of this agent platform, for instance, flexible and efficient messaging, and its limitations, such as the primitive character of the basic agent model used.

Section III: Agent Theories II

Shanahan and Witkowski present a logic-based agent programming framework that allows the interleaving of planning, plan execution, and perception, and preserves reactivity. The framework is based on the event calculus. Planning and observation assimilation are viewed as abductive theorem proving tasks and are implemented in a logic programming meta-interpreter. Simple application programs for robot navigation and map building are provided as examples.

Terán et al. present an approach for analyzing the behavior of multiagent systems through constraint-based search of possible system trajectories. The approach can be viewed as doing a model-based proof that the under the given constraints, the system's behavior satisfies certain properties. They describe a method for compiling a system specification into one that is more efficiently searchable.

Section IV: Models of Agent Communication and Coordination

Norman and Reed study delegation and responsibility by analyzing the delegation of a task from one autonomous agent to another in terms of the performance of some imperative communication act, and the creation of a normative state of affairs. They distinguish between an agent doing something, being responsible for getting something done, and being responsible for bringing about a state of affairs. In their framework, it is possible to evaluate the coherence of an agent's action with respect to the agent's commitments.

Dignum et al. present an architecture for agents that can discuss the formation of a team, construct a collective plan by dialogue, and work as a team member until the collective goal has been fulfilled. They explain how fixed protocols like contract net do not suffice for this, and propose the use of *structured dialogues*, with an emphasis on persuasion.

The dialogues in which agents engage are described formally using modal logic and speech acts.

Nodine et al. explore issues related to the coordination of complex and long running tasks in systems involving unreliable agents. Their analysis of coordination issues uses three crucial categories: tasks, roles, and conversations. They claim that current agent communication languages make certain assumptions about the agent system, such as stability or intelligence, that are not always justified in information-centric applications. They discuss what additional requirements agent communication languages should satisfy for this, and propose changes to fulfill them.

Section V: Autonomy and Models of Agent Coordination

Luck and d’Inverno introduce the notion of *sociological agents*, agents who can model their social environment and in particular the social relationships in which they participate, for instance, autonomy, dependence, and power relations. They explain how by making use of this information, the agents can enhance their performance by exploiting social opportunities and exerting power over others. The models of social relationships used are not limited to characterizing problem situations, such as dependencies, or the notion of “owned resource”, but also provide the agents with solutions, for example, through self-sufficient plans.

In their paper, *Sierra et al.* present a new approach to landmark-based robot navigation based on the use of bidding mechanisms to coordinate the actions requested by a group of agents in charge of guiding a robot towards a specified target in an unknown environment. The approach uses a qualitative (fuzzy) representation of landmark information.

Clement and Durfee, working on the complexity of the coordination problem, show how coordinating CHiPs, i.e. individually formed concurrent hierarchical plans, at a higher level is still exponentially cheaper than at lower levels, in spite of the possibly exponential growth of the associated summary information. The paper also proposes heuristics that take advantage of summary information to smartly direct the search for a global plan. Experiments that confirm the effectiveness of the proposed heuristics are described.

Section VI: Agent Languages

Hindriks et al. address a fundamental deficiency in existing agent programming languages, namely that their notions of goal are more like procedure calls than true declarative goals. They propose a simple agent programming language that supports declarative goals. The language incorporates simple mechanisms for selecting actions based on the agent’s beliefs and goals and for goal revision. A programming logic is also proposed for reasoning about programs in this language.

Shapiro and Lespérance present the latest version of their CASL language for specifying multiagent systems. CASL supports the specification of complex agent behaviors

and how they depend on agents' knowledge and goals. The new version has a simpler model of goals and encrypted communication actions. Use of the language is illustrated by modeling a fairly complex multiagent system that resolves feature interactions in telecommunications.

van Eijk et al. present enhancements to their multiagent programming framework inspired from work on the semantics of concurrent programming. To better support the programming of open and complex agent systems, they incorporate notions drawn from object-oriented programming, such as agent classes, and agent communication interfaces defined in terms of question templates.

Marini et al. propose a metalanguage for specifying agent architectures and multiagent system configurations. The metalanguage allows specifications to be given in a procedural format. Concrete architectures can be specified as instances of abstract architecture classes. A formal semantics is provided for the specification metalanguage.

Section VII: Planning, Decision Making, and Learning

Antunes et al. present a novel approach for decision making, an alternative to classical utilitarian theories. Their BVG agent architecture relies on the use of *values*, i.e. multiple dimensions against which to evaluate a situation, to perform choice among a set of candidate goals. In the BVG model, agents adjust their scale of values by feeding back evaluative information about the consequences of their decisions. In the paper, new experiments are presented where the agents have to make decisions without having all relevant information.

Koenig addresses the problem of synthesizing optimal plans for agents, e.g. delivery robots, in the presence of uncertainty and immediate soft deadlines. He does this by introducing the additive and multiplicative planning-task transformations: fast representation changes that transform planning tasks with convex exponential utility functions to planning tasks that can be solved with variants of standard deterministic or probabilistic AI planners. Advantages of such representations include that they are context insensitive, fast, scale well, allow for optimal or near-optimal planning, and are grounded in utility theory.

In his paper, *Weiß* presents M-Dyna-Q, a multiagent coordination framework that integrates joint learning, planning, and reacting. M-Dyna-Q is an extension of Dyna-Q, which is an instantiation of Dyna, a single-agent architectural framework that integrates learning, planning, and reacting. M-Dyna-Q bridges the gap between plan-based and reactive coordination in multiagent systems. After summarizing the key Dyna features, M-Dyna-Q is described in detail, and experimental results on its use are provided. Benefits and limitation of the framework are discussed.

Panel Summaries

Bryson et al. describes the questions that the panel on *Agent Development Tools* attempted to address, the essence of the panelists' statements, and the general conclusions that were drawn.

Castelfranchi, who chaired the other panel on *Autonomy — Theory, Dimensions, and Regulation*, introduces the topics that were discussed at this panel. This is followed by statements by the panelists.

March 2001

Cristiano Castelfranchi (Siena, Italy) and
Yves Lespérance (Toronto, Canada)

References

1. N.R. Jennings and Y. Lespérance, editors. *Intelligent Agents VI — Agent Theories, Architectures, and Languages, 6th International Workshop, ATAL'99, Orlando, Florida, USA, July 15–17, 1999, Proceedings*, volume 1757 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, 2000.
2. J.P. Müller, M.P. Singh, and A.S. Rao, editors. *Intelligent Agents V — Agent Theories, Architectures, and Languages, 5th International Workshop, ATAL'98, Paris, France, July 4–7, 1998, Proceedings*, volume 1555 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, 1999.
3. J.P. Müller, M.J. Wooldridge, and N.R. Jennings, editors. *Intelligent Agents III — Agent Theories, Architectures, and Languages, ECAI'96 Workshop (ATAL), Budapest, Hungary, August 12–13, 1996, Proceedings*, volume 1193 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, 1997.
4. M.P. Singh, A. Rao, and M.J. Wooldridge, editors. *Intelligent Agents IV — Agent Theories, Architectures, and Languages, 4th International Workshop, ATAL'97, Providence, Rhode Island, USA, July 24–26, 1997, Proceedings*, volume 1365 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, 1998.
5. M. Wooldridge and N.R. Jennings, editors. *Intelligent Agents — ECAI-94 Workshop on Agent Theories, Architectures, and Languages, Amsterdam, The Netherlands, August 8–9, 1994. Proceedings*, volume 890 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, 1995.
6. M. Wooldridge, J. Müller, and M. Tambe, editors. *Intelligent Agents II — Agent Theories, Architectures, and Languages, IJCAI'95-ATAL Workshop, Montreal, Canada, August 19–20, 1995 Proceedings*, volume 1037 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, 1996.