# Reports on the AAAI 1999 Workshop Program

*18–19 July 1999*

The AAAI-99 Workshop Program (a part of the sixteenth national conference on artificial intelligence) was held in Orlando, Florida. The program included 16 workshops covering a wide range of topics in AI. Each workshop was limited to approximately 25 to 50 participants. Participation was by invitation from the workshop organizers. The workshops were

- Agent-Based Systems in the Business Context
- Agents' Conflicts
- Artificial Intelligence for Distributed Information Networking
- Artificial Intelligence for Electronic Commerce
- Computation with Neural Systems Workshop
- Configuration
- Data Mining with Evolutionary Algorithms: Research Directions (Jointly sponsored by GECCO-99)
- Environmental Decision Support Systems and Artificial Intelligence
- Exploring Synergies of Knowledge Management and Case-Based Reasoning
- Intelligent Information Systems
- Intelligent Software Engineering
- Machine Learning for Information Extraction
- Mixed-Initiative Intelligence
- Negotiation: Settling Conflicts and Identifying Opportunities
- Ontology Management
- Reasoning in Context for AI Applications

Summaries of 13 of the 16 workshops follow. A longer version of the Workshop on Artificial Intelligence for Electronic Commerce will be published in a subsequent issue.

## Agent-Based Systems in the Business Context

*Work-flow management systems* are integrated software tools for supporting the modeling, analysis, and enactment of business processes. The development of this technology has been driven by the move toward process-oriented management in the 1990s using initiatives such as Continuos Business Process Improvement and Business Process Reengineering. Its coordination benefits have generated a commercial market estimated in 1996 to be worth $2.6 billion. Despite this market size, the technology is currently limited to simple administrative-type problems, and work-flow researchers are seeking to develop the new techniques demanded by more dynamic applications. In contrast, the AI community has been involved with related research on dynamic process management for several decades. This workshop brought together researchers and practitioners from each field to identify and discuss the AI technologies that could be leveraged to meet the requirements of future work-flow systems.

Speakers with a range of backgrounds presented 12 papers, which led to wide-ranging debates on the issues raised. It became clear that although AI provides a wealth of relevant technologies, and the work-flow field offers opportunity for their widespread deployment in commercial applications, the AI community has much to learn from exploring work flow as a test bed for its ideas. Domains that involve the active control of computational entities and physical devices have traditionally motivated AI. The list of relevant technologies includes agent-based architectures, plan synthesis, scheduling and constraint management algorithms, and reactive control architectures. However, the physical and computational agents to which they have traditionally been applied are more crisply defined and more willing to accept instructions without explanation than their human counterparts in work-flow applications. People want to see a range of options that they can explore to tailor a solution to their preferences, and it is more difficult to crisply define their roles and capabilities in the context of a process. There is much in common here with the recent mixed-initiative push in AI research. Work flow offers opportunities for AI to really add value to commercial applications, also offering an interesting test bed to motivate the development of the field. The workshop proceedings contain excellent papers that outline the real-world requirements of work flow and the embryonic efforts of AI researchers in addressing them. It provides an excellent starting point if you are interested in investigating how this exciting area can benefit from your research.

*Brian Drabble*
University of Oregon

## Agents' Conflicts

This workshop followed the ECAI'98 workshop "Conflicts among Agents: Avoid or Use Them?" and the IJCAI-97 workshop entitled "Collaboration, Cooperation, and Conflict in Dialogue Systems." The aim this time was to focus on definitions of agents' conflicts and their roles within a multiagent system, that is, how this system might evolve thanks to, despite of, or because of conflicts. The workshop included three invited papers and six submitted papers.

The following original ideas were proposed: (1) within a team of agents, conflict handling is a team goal; (2) search for an agreement is better than averaging (numeric) criteria; (3) spatial conflicts among a team of robots can generate a derived behavior; (4) conflicts are not necessarily disturbing for the set of agents; (5) conflict handling might not be necessary; and (6) decisions can be made on the basis of conflicts.

The main conclusions of the discussion were the following: First, it is easier not to be in conflict than to be in conflict: the former can mean that the agents are not even interacting, and the latter supposes that the agents are within the same context. Second, conflict is often symmetric, but it might happen that it is not. Third, are conflicts useful? The answer depends on

the problem. To be useful, a conflict must be observed. Fourth, what we learn from a conflict depends on the situation. Learning is possible if agents are aware of the conflict. Fifth, conflicts are positive in certain cases; for example, they can create specific behaviors, create competition, or stimulate inference. Please visit www.cert.fr/fr/dcsd/PUB/AAAI99/conflicts.html for more details!

*Laurent Chaudron*
*Catherine Tessier*
Onera-Cert-DCSD

## Artificial Intelligence for Distributed Information Networking (AIDIN'99)

This third workshop follows the IJCAI-95 and IJCAI-97 workshops on the use of AI in wired and wireless distributed information networking (DIN).

The workshop again focused on complex DIN problems. Twelve papers were presented on these topics: security management, fault management, service management, user applications, routing, quality of service, and congestion management. Four invited speakers gave excellent presentations on intelligent agents in wired and wireless networking, electronic marketplaces, network routing with reinforcement learning, and network routing with collective intelligence. See liawww.epfl.ch/AiDIN99/ for full details on the workshop.

Two panel discussions were held. The first was entitled "An Agent Future for DIN Management?" and was led by Sue Abu-Hakima. Intelligent agents have quickly become a key technology in information networking, leaving us with many open issues, especially standardization now tackled by FIPA (www.fipa.org) and MASIF (www.omg.org). The second panel was lead by Beat Liver and was entitled "The Future of Resource Network Management?" The panel discussed the main goal in resource management: maximizing service provider profits while maximizing the use of available resources. This goal requires the management of complex value chains and service agreements as well as the maximizing of economic efficiencies. All these areas provide rich application domains for AI technology.

The attendees agreed that the need for AI technology in DIN is essential in the complex world of converging networks integrated with seamless information provision from the World Wide Web. We agreed that we might hold another workshop at AAAI-2000 or organize a spring symposium on this continuously evolving hot topic.

*Sue Abu-Hakima*
AmikaNow! Corporation

*Steven Willmott*
Laboratoire d'Intelligence Artificielle

## Computation with Neural Systems

This workshop focused on the interrelationship between artificial and real neural networks and computation and formed a part of a series of workshops to be held on this topic (the second workshop took place at the International Conference on Artificial Neural Networks in 1999). The aim of the workshops was to transfer findings on how the brain might operate with computer science with the possibility of inspiring novel research in computer science. The workshop highlighted the issues of how the brain deals with synchronization, processing speed, timing, robustness, modular construction, information representation, and transmission because all these areas are important problems in computer systems. The workshop attracted nine papers, which managed to capture all these issues. Approximately half the papers covered results arising from models of neural systems that were balanced by papers presenting artificial systems that were inspired by biological systems. The workshop concluded that there are many lessons that can be learned from both disciplines, but a major hurdle is the use of a common language between the areas. It was thought that the future workshops would clearly allow this to develop and might then provide a platform for developing novel collaborative research with novel computer architectures.

The workshop organizers were Jim Austin, Stefan Wermeter, Vasant Honavar, and Victoria Hodge. Details of future workshops can be found at www.his.sunderland.ac.uk/emernet/.

*Jim Austin*
University of York

## Configuration

*Configurators* are a cornerstone for successful applications of the mass customization paradigm. Combined with e-business solutions, these techniques have a high market impact. Configuring can be defined, in a simplifying way, as the task of identifying the set of parts or components that, assembled together, will satisfy a customer's request. In more sophisticated applications, configuring also includes the configuration of hardware and software as well as the determination of correct parameter settings of the components.

Configurators are generating new business opportunities in many industries for new products and new ways to interact with customers. However, efficient development and maintenance of configurators require sophisticated software development techniques. AI methods, more than ever, are central to the development of powerful configuration tools.

This workshop succeeded in drawing together researchers and practitioners. From a total of 45 attendees, 20 participants came from industry. The affiliations of the other 25 participants were universities or national research laboratories.

The accepted papers were grouped around common themes: knowledge acquisition and representation, reconfiguration and diagnosis of knowledge bases, reasoning, and tools and applications. Ample time was left for discussion.

An old, important but still unsolved challenge is the generation of helpful explanations. This issue was one of the central discussion points in the session entitled "Knowledge Acquisition and Representation."

The session entitled "Reconfiguration and Diagnosis of Knowledge Bases" showed the importance of supporting the maintenance of existing product configurations. In business

domains where products are continuously changed to fit the customer's needs, efficient reconfiguration is a necessity. The discussion proved that enhancements of the current techniques are desirable.

The main focus in the session on reasoning dealt with improvements of configuration algorithms based on constraint satisfaction. An interesting set of questions posed to industry was, "Is it worth improving our configuration algorithms, or are they already good enough? As a consequence, should we invest our person-power in other research topics such as knowledge acquisition and maintenance?" The uncontradicted answer was that improving reasoning is still an important issue for various reasons, for example, optimization or integration of configurators in the production process.

This need for integrating configurators at various points from design through sales and maintenance was pointed out in the session entitled "Tools and Applications." The integration of configuration and production planning and scheduling is promising from an economic point of view because it contributes to an optimization of the manufacturing process.

In addition to the panel discussions, system demonstrations were presented by Baan, GMD FIRST, Tacton Systems, University Bremen, University Hamburg, and Lenze and showed the current state of the art in configuration tools. Additional information can be found at www.ifi.uni-klu.ac.at/Conferences/aaai99_ws_configuration.

*Boi Faltings*
*Eugene C. Freuder*
*Gerhard Friedrich*
Universitaet Klagenfurt

## Data Mining with Evolutionary Algorithms: Research Directions

The Workshop entitled "Data Mining with Evolutionary Algorithms: Research Directions" was jointly sponsored by the American Association for Artificial Intelligence and GECCO-99. The general goal of the workshop was to discuss research issues concerning the integration of two areas: (1) data mining and (2) evolutionary algorithms. The workshop brought together people from these two research communities, from both academia and industry. This half-day workshop was attended by approximately 40 participants and consisted of 5 paper presentations, as follows:

First, James Thomas and Katia Sycara (both of Carnegie Mellon University) presented a genetic programming–based system for trading rule discovery, whose performance was evaluated over real-world exchange rate data in the dollar-yen and dollar-dm markets. They focused on the issues of rule complexity and how to fight overfitting. Gary Weiss (Rutgers University) addressed the problem of predicting rare events from a sequence of events. The author has developed a genetic algorithm–based system, TIMEWEAVER, that, given a prespecified target event, learns to identify patterns in the data that successfully predict the future occurrence of this event. In essence, TIMEWEAVER uses a Michigan-style genetic algorithm to evolve a set of prediction rules. It also uses a niching strategy to ensure that a diverse set of rules is achieved.

Second, Cesar Guerra-Salcedo and Darrell Whitley (both of Colorado State University) proposed the use of genetic algorithms to select features for an ensemble of classifiers. They have experimented with two classifiers—(1) C4.5 and (2) Euclidean decision tables—and several ensemble-construction methods, including bagging and boosting.

Third, Simon Thompson (BT Labs, United Kingdom) proposed using a genetic algorithm to prune an ensemble of classifiers. The genetic algorithm works with a real-valued encoding, and it is used to optimize the weightings of the classifiers in the ensemble. He also discussed some research directions in the use of genetic algorithms for postprocessing of the discovered knowledge, including the supplementation of this knowledge, for example, the discovery of new rules covering data that are not correctly classified by the current rules.

Fourth, Nicolas Monmarche (University of Tours, France) presented a hybrid method for clustering, combining the stochastic principles of the artificial ants paradigm with a well-known statistical method for clustering, the K-MEANS algorithm.

Fifth, in addition to these projects, the workshop proceedings also contains a paper that was not presented at the workshop, entitled "Genetic Algorithms for Selection and Partitioning of Attributes in Large-Scale Data-Mining Problems," by William Hsu, William Pottenger, Michael Welge, Jie Wu, and Ting-Hao Yang (National Center for Supercomputing Applications).

The proceedings of the workshop were published by AAAI Press in 1999 as a technical report, WS-99-06, ISBN 1-57735-090-1.

*Alex A. Freitas*
Pontifical Catholic
University of Parana (PUC-PR),
Curitiba, Brazil

## Environmental Decision Support Systems and Artificial Intelligence: New Issues

This workshop was the perfect opportunity to bring together researchers from both environmental sciences and AI. The workshop had 25 submissions, and 13 papers were accepted mainly from European research centers. This workshop is the second that AAAI has hosted to show the ongoing and growing collaboration of AI researchers and environmental scientists.

Accepted papers touched a wide variety of environmental problems, such as environmental and technological risk management, forest resource management, ozone concentration prediction, trumpeter swan management, and waste water treatment plants, showing the importance and strengths of interaction between the two communities. All the contributors were asked to organize the presentation of their environmental decision support system (EDDS) around the following questions: What is the environmental problem we are trying to solve? What are the AI methods used? How are they facing the problem? How are we testing the methods?

The workshop was organized in five blocks. The first block was "From classic DSS to EDSS." The second block was "Soft Computing and the Environment." The third block was "Distributed and Integrated AI and Environmental Issues." The fourth block was "Ontologies and KBS in Environmental Issues." The last block was a general discussion about the future of the interaction between AI and environmental sciences.

During the final panel of the workshop, we tried to identify those issues that are central in the development of an EDDS and that require more in-depth work in future years.

*U. Cortés*
*M. Sànchez-Marrè*
Technical University of Catalonia

## Exploring Synergies of Knowledge Management and Case-Based Reasoning

This workshop, which attracted approximately 40 participants, concentrated on intersecting issues of these two fields. Previous workshops focused on these two topics, but none on their intersection. This workshop was well timed; both case-based–reasoning industrialists and researchers (for example, studying e-commerce and related topics) recently explored issues related to applying case-based reasoning to knowledge management processes.

This workshop included seven invited talks, mostly from knowledge management experts and case-based–reasoning industrialists (that is, representing tecInno and The Haley Enterprise) to balance the novel contributions, which were mostly from case-based–reasoning researchers. Invited talks included discussions on enterprise resource planning systems, case studies, process-centered knowledge management, and the use of ontologies to support knowledge management processes. Contributed papers consisted mainly of descriptions of interactive applications and knowledge management and case-based reasoning frameworks; they often highlighted knowledge management processes (for example, embedding the case-authoring process in a problem-solving context). Several presenters argued that knowledge management techniques are not limited to business tasks and that the case-based reasoning task-decomposition cycle does not address all the behaviors needed to support knowledge management processes.

The final panel identified several lessons learned: Case-based–reasoning applications to knowledge management should closely integrate case-based–reasoning processes into the knowledge management tasks and models that they target; knowledge management tasks demand more advanced case-authoring tools; and organizational dynamics must be considered when applying case-based reasoning to knowledge management processes. We anticipate that future work relating knowledge management with case-based reasoning will address these issues.

*David W. Aha*
National Research Laboratory

*Irma Becerra-Fernandez*
FIU

*Hector Munoz-Avila*
University of Maryland

## Intelligent Software Engineering

There is an emerging consensus that the areas of AI and software engineering need to enter into a new, mutually symbiotic, relationship. AI offers software engineering crucial pointers for building tools to support and automate, at least partially, the key human-centric tasks involved in software development and maintenance. Software engineering provides AI researchers perhaps the best test bed to evaluate the efficacy of new tools and techniques. The intent of this workshop was to provide a forum for discussing ideas at the interface of these two areas.

The response to the idea of such a forum was encouraging. The workshop included 10 contributed papers, an invited talk, and a panel discussion. The contributed papers provided a representative sample of research in this area, covering topics such as inconsistency handling in requirements engineering, updating of formal specifications, automated component retrieval and synthesis, the application of abductive techniques in software testing and debugging of hardware designs, the interplay between software engineering and knowledge engineering, and methods for building lightweight domain-specific tools that exploit commercial off-the-shelf products. The invited talk challenged some common assumptions on the role of metaknowledge in knowledge maintenance. The panel discussion explored the application of AI techniques in software maintenance.

The vast majority of participants were AI researchers and practitioners sharing a conviction in the utility of AI principles and techniques to software engineering, both at a foundational and at a practical level. How this message could be convincingly conveyed to the software engineering community was a concern shared by many. A final (unscheduled) discussion, involving a large number of workshop participants, indicated a real need for a forum such as this becoming a regular feature. Plans are under way to make this workshop a continuing series.

*Aditya Ghose*
University of Wollongong

*Tim Menzies*
NASA Software Independent Verification and Validation Facility

*Ken Satoh*
Hokkaido University

## Machine Learning for Information Extraction

The dramatic growth in the number and size of online textual information sources has fueled increasing research interest in the information-extraction problem. Given a set of text documents from some domain, an information-extraction system automatically populates a predefined database by extracting relevant fragments from the documents. Because manually constructed information-extraction systems must be adapted to each new problem domain, various machine-learning techniques have recently been applied to the information-

extraction problem. The main purpose of the workshop was to deepen the information-extraction community's understanding of the state of the art by bringing together researchers who address the information-extraction problem from different perspectives.

The workshop had two invited presentations. "New Directions in Information Extraction," by Claire Cardie, emphasized that learning extraction rules is just one of many parts of the information-extraction problem to which machine learning can be applied. "Information Extraction from the Web," by Tom Mitchell, integrated the information-extraction problem into the larger context of text learning. Paper presentations explored the application of propositional rule learning, inductive logic programming, and hidden Markov models to the problem of learning extraction patterns. Two overview papers were also presented: a review of the state of the art in learning extraction rules and a survey of the roles that machine learning plays within different components of information-extraction systems. For more details, see the online proceedings at www.isi.edu/~muslea/RISE/ML4IE/.

*Mary Elaine Califf*
Illinois State University

## Mixed-Initiative Intelligence

*Mixed-initiative intelligence* represents an amalgam of human and machine cognition that together produces intelligent behavior. Mixed-initiative systems integrate human and automated reasoning to take advantage of their respective reasoning styles and computational strengths. The benefit is the potential to combine the resources available to both; the challenge is to manage the interaction and responsibilities encountered in joint decision making. Such were the assumptions that motivated the workshop, although many objected to some of the finer points expressed.

This workshop was the first gathering to bring together various groups performing research on mixed-initiative systems. The participants included two major camps that have seldom worked together closely: (1) the mixed-initiative planning community and (2) the mixed-initiative dialogue communities. The mixed-initiative planning community emphasizes the "mixed" aspect, and the mixed-initiative dialogue community emphasizes the "initiative" aspect. In addition to these members, many other representatives from various disciplines added to the interaction. Also, the workshop included persons presenting research based on cognitive engineering, decision-theoretic, and experimental psychology approaches. The result was a pleasant mix of 33 individual points of view.

One of the highlights of the workshop was an invited presentation by Paul Cohen (University of Mas-

sachusetts). In the discussion, he declared that in all the approaches to this topic, a common issue arises, that is, to get computers to effectively understand and work intimately with people, with a kind of mind reading. For example, when a user makes a mistake, the program ideally should understand what the human really meant. The conclusion is that for mixed-initiative systems to be realized, a shared context and shared semantics must be present.

Various flavors of mixed-initiative computing share many human-computer–interaction issues because they must address the complexity of communication between human and machine. However, mixed-initiative computing transcends the fundamental concerns of other areas because the design of such systems implies problems of heterogeneous decision making. The issues, therefore, include the division of processing and knowledge access, the coordination of natural language dialogue among agents (both human and machine), and the shift of initiative and control during problem-solving and planning.

Although the participants held much in common, each had his or her own emphasis. Some concentrated on learning, some on agent-human interactions. Others presented fundamental results, and others examined implemented systems and system-building approaches. Numerous persons discussed planning and decision making, dialogue, sociological and sensory contexts, and many other features that differentiate their results. The workshop goal was to examine mixed-initiative computing from a wide perspective. In this regard, the workshop was a success.

*Michael Cox*
Wright State University

## Negotiation: Settling Conflicts and Identifying Opportunities

As the internet becomes an increasingly viable medium for the deployment of autonomous agents acting on behalf of their users, it becomes evident that such agents should be endowed with the capability of negotiating settlements with other agents to further user interests. Typical negotiation instances in this context include settling prices for services and goods, forming dynamic alliances with other agents to effectively exploit windows of opportunities, and settling goal or resource conflicts.

This workshop was organized to investigate the critical need for negotiation mechanisms and reasoning procedures by which two or more agents can quickly reach consensus about arranging their actions to avoid conflicts and enable new collaborations. Research papers presented discussed the following topics: using Bayesian networks to represent influences on the parties to a negotiation and creating a favorable negotiation context based on such models; eliminating possible conflicts between agents by exchanging abstract high-level plans; designing an agent architecture and a multistage negotiation protocol to resolve a multi-issue dispute; repeatedly negotiating among subgroups to form larger and larger coalitions leading to the formation of a grand coalition involving all agents; switching from minimal communication to more verbose dialogue to resolve conflicts when the breakdown of consensus is detected through probing; learning from reinforcements to bid competitively in electronic markets; prescribing optimal contracting decisions to promote expected social welfare under different decommitment procedures; establishing negotiation procedures to allocate continuously divisible goods that improve participant satisfaction and guarantee "envy-freeness"; resolving intra-agent and interagent conflict at different objective levels; and detecting settlement of all negotiation issues in a distributed setting. The workshop also hosted a lively panel entitled "Negotiation Possibilities in E-Commerce," with presentations by panelists Tuomas Sandholm, Michael Wellman, and Sandip Sen. More information is available at euler.mcs.utulsa.edu/~sandip/wshop/aaai99/.

*Sandip Sen*
University of Tulsa

## Reasoning in Context for AI Applications

Twenty researchers from various scientific areas participated in this one-day workshop. This workshop followed several months of electronic discussion in which 50 participants discussed the topic, sometimes vigorously.

The workshop began with a short session in which the participants introduced themselves and their positions. Some participants took a cognitive approach to context and were mainly interested in how context is represented and what the nature of reasoning with context is. Others took a software-engineering approach and used context to help manage knowledge and decision making in their systems. Within their applications, participants focused on how context is represented, how agents exploit context in reasoning, and how context affects interagent interaction.

The rest of the workshop was devoted to group discussion. Results included identifying the several distinct entities commonly referred to as "context" or "situation" and recognizing that a large number of items have to be considered in an agent's context; the items do not always have direct links to each other, they are not always at the same level of detail, and they can belong to the past or the future of the current context. In addition, shared (or at least compatible) context was seen to be important between agents, although having the same view of the context is not enough. Each agent must know that other agents possess these views.

*Patrick Brezillon*
*Jean-Charles Pomerol*
University of Paris VI

*Roy Turner*
*Elise Turner*
University of Maine