

# AAAI 2000 Workshop Reports

*Yves Lesperance, Gerd Wagner, William Birmingham, Kurt Bollacker, Alexander Nareyek, J. Paul Walser, David Aha, Tim Finin, Benjamin Groszof, Nathalie Japkowicz, Robert Holte, Lise Getoor, Carla P. Gomes, Holger H. Hoos, Alan C. Schultz, Miroslav Kubat, Tom Mitchell, Joerg Denzinger, Yolanda Gil, Karen Myers, Claudio Bettini, and Angelo Montanari*

The AAAI-2000 Workshop Program was held Sunday and Monday, 30–31 July 2000 at the Hyatt Regency Austin and the Austin Convention Center in Austin, Texas. The 15 workshops held were (1) “Agent-Oriented Information Systems,” (2) “Artificial Intelligence and Music,” (3) “Artificial Intelligence and Web Search,” (4) “Constraints and AI Planning,” (5) “Integration of AI and OR: Techniques for Combinatorial Optimization,” (6) “Intelligent Lessons Learned Systems,” (7) “Knowledge-Based Electronic Markets,” (8) “Learning from Imbalanced Data Sets,” (9) “Learning Statistical Models from Relational Data,” (10) “Leveraging Probability and Uncertainty in Computation,” (11) “Mobile Robotic Competition and Exhibition,” (12) “New Research Problems for Machine Learning,” (13) “Parallel and Distributed Search for Reasoning,” (14) “Representational Issues for Real-World Planning Systems,” and (15) “Spatial and Temporal Granularity.”

## Agent-Oriented Information Systems

Information systems are, and continue to be, the predominant application of computing technologies. The development, maintenance, and evolution of information systems are the primary preoccupation of many computing professionals. In spite of ongoing advances in many areas of informa-

tion technology, computing professionals continue to face numerous challenges. Information systems are now placed in ever-more demanding roles, in distributed and networked environments, and with ever-higher expectations from users—in functional as well as nonfunctional qualities,

---

The AAAI-2000 Workshop Program was held Sunday and Monday, 30–31 July 2000 at the Hyatt Regency Austin and the Austin Convention Center in Austin, Texas. The 15 workshops held were (1) “Agent-Oriented Information Systems,” (2) “Artificial Intelligence and Music,” (3) “Artificial Intelligence and Web Search,” (4) “Constraints and AI Planning,” (5) “Integration of AI and OR: Techniques for Combinatorial Optimization,” (6) “Intelligent Lessons Learned Systems,” (7) “Knowledge-Based Electronic Markets,” (8) “Learning from Imbalanced Data Sets,” (9) “Learning Statistical Models from Relational Data,” (10) “Leveraging Probability and Uncertainty in Computation,” (11) “Mobile Robotic Competition and Exhibition,” (12) “New Research Problems for Machine Learning,” (13) “Parallel and Distributed Search for Reasoning,” (14) “Representational Issues for Real-World Planning Systems,” and (15) “Spatial and Temporal Granularity.”

---

such as robustness, responsiveness, and flexibility.

Agent orientation has a lot to offer information systems, both in terms of

enabling technologies and in terms of conceptualizing what information systems are and can be and how they can be developed, sustained, and evolved.

Despite this potential, agent concepts have mostly played only a niche role in information systems, for example, in offering specialized capabilities in information searching and brokering. The Agent-Oriented Information Systems (AOIS) Workshop series aims to promote and stimulate interest and discussion in an agent-oriented conception of information systems, covering all aspects of information systems and all stages of development—requirements analysis, design, and implementation.

To promote cross-fertilization of ideas between the AI agents community and the information systems community, the AOIS workshop was held at two locations this year, at AAAI in Austin, Texas, and at the Conference on Advanced Information Systems Engineering (CAISE) in Stockholm, Sweden. This year’s workshop continues the success of AOIS’99, which was held at the Autonomous Agents’99 conference in Barcelona and at CAISE’99 in Heidelberg, Germany.

Approximately forty-five participants attended AOIS-2000, with around 25 coming to the AAAI/Austin location. The program in Austin included invited presentations by James Odell and Bill McCarthy, and eight paper presentations dealing with a wide range of issues in the area.

The first invited talk, by well-known object-oriented software development author Odell, discussed efforts within FIPA and the OMG Agents Working Group to develop and standardize extensions to the UML modeling language to support the design of agent-based systems. He described several aspects of the extended language agent UML that is being developed, with particular attention to modeling agent-interaction protocols. This work is important for furthering the use of agent technology in industry by relating agent-oriented models to existing techniques.

In the other invited talk, McCarthy,

a well-known business and information systems researcher, presented his work on the resource-event-agent (REA) model of economic transactions and how it has been applied to represent enterprise value chains and work flows. He discussed how the model might be used to develop agent-based systems where the agents act as full-fledged economic participants.

Among the paper presentations, Cernuzzi and Giret described their work on comparing different agent-oriented software-engineering methodologies and discussed a case study that was done using Kinny, Georgeff, and Rao's (KGR) methodology. The KGR methodology was criticized for its lack of support for specifying fuzzy goals and behavior evolution and for its translation of a design into an implementation.

Vercouter, Beaune, and Sayettat presented a decentralized approach to the integration of new agents in an open multiagent system. Agents learn stereotypes to classify their acquaintances. When they encounter a new agent, they interact and form a description of it, classify it, and then can recommend acquaintances to it.

Camacho, Molina, and Borrajo described their experience in building a multiagent system for electronic travel planning using web information resources. The architecture involves planning agents based on the PRODIGY planner and web agents that access information sources and interpret the results.

Filipe argued that socially aware agents need to be able to cope with different norm systems. He presented a model inspired from the INFORMATION FIELD paradigm for how agents might resolve normative conflicts based on their values and beliefs.

Zhu, Greenwood, Huo, and Zhang discussed how quality management in information systems was being affected by agent orientation, the web, and other recent changes in information technology. They presented a model of information system evolution through "growing up" and sketched how agent-based quality management tools might be designed to cope with this.

Vinaja, Slinkman, and Mykytyn presented the results of an experimental study of the effects of internet information delivery and agent facilitation on a decision-making task. The study found that agent facilitation leads to fewer information sources being examined and higher satisfaction with the process. Workshop participants were interested in the approach but questioned whether the results could be generalized beyond the system studied.

Kaminka, Pynadath, and Tambe presented a nonintrusive approach to the monitoring of complex multi-agent systems based on eavesdropping on the messages exchanged by the agents. The approach performs plan recognition on the messages and uses a model of team coherence to infer the system state more accurately.

Mukherjee, Dutta, and Sen discussed the use of agents to support query reformulation, a key problem in information retrieval. Such agents will need rich domain ontologies to perform their task. Characteristics of domain ontologies that facilitate the reformulation task are identified, in particular the presence of suitable ordering relations.

Among the wealth of impressions and questions that were left in mind from the workshop, there are two that perhaps stand out: One is the enormous breadth of the AOIS domain. Our goal is to build systems modeled after social organizations that will be deployed in the real world, with all the engineering skill that this task requires. We will need the contributions of people from many disciplines. The second is the importance of reaching out to practitioners and presenting agent-based concepts and methods in ways that ease the transition from traditional approaches.

The number of submissions and participants, and the interactions at the workshop, suggest that AOIS will develop into an important area both from the AI agents' standpoint and from the information systems standpoint. The organizers are planning to continue holding the workshop to facilitate and promote the develop-

ment of this area. Further information about this workshop and future ones can be found at [aois.org](http://aois.org).

*Yves Lesperance*

York University

*Gerd Wagner*

Eindhoven University of Technology

*Eric Yu*

University of Toronto

## Artificial Intelligence and Music: Toward Formal Models for Composition, Performance, and Analysis

The workshop on Artificial Intelligence and Music presented recent research organized into three main topics. The first, harmony, included papers on harmonic analysis, prediction of harmonic progressions, educational software, and harmonization of melodies. The second topic, transcription, included work on transcription from MIDI (essentially, piano-roll data) to notation, beat tracking, benchmarking tools, and composer identification. The third topic, performance, included working on real-time improvisation, generating expressive performance, learning about expressive performance, and learning to classify performance styles.

A general feature of this research is that there are few well-defined standard problems. Instead, music research tends to be task oriented, often with artistic goals rather than utilitarian ones, making objective evaluation particularly difficult. The need for more scientific evaluation was discussed. In addition, there was a call for more formal analysis of problems and algorithms, so that researchers can better understand and communicate research problems and results. Both top-down and bottom-up approaches to music processing were represented.

The workshop featured two invited speakers. Greg Wakefield gave an overview of music transcription, from the earliest efforts to the most recent results. Unfortunately, David Temperly was unable to attend because of a flight cancellation, but his invited paper is an excellent presentation of his recent work on understanding counterpoint. Overall, the workshop

provided a friendly and interactive forum for researchers to meet and discuss current research results and new directions. The chairs are grateful to the organizing committee and attendees for making this workshop so rewarding.

*William Birmingham*  
University of Michigan

## Artificial Intelligence for Web Search

The World Wide Web offers an immense diversity of knowledge; however, the web's size, rapid growth, high diversity, and erratic organization often make it difficult to find information. This workshop was designed to allow for exploration of AI techniques as applied to the problems of organizing, searching, and classifying information on the web. Although there was a diversity of interesting work, there were some clusters of related research: (1) building ontologies, rich indexes, and other sophisticated contexts for improving the web search process; (2) exploring similarity measures between documents to find more relevant, useful ones; and (3) attempting to improve query result ranking, including using user feedback.

### Search in Knowledge-Rich Contexts

When a "context-free" search on the web is not sufficient, there was a strong case made for the creation of knowledge-rich, structured contexts in which searches can be made. One approach suggested that web documents could be built to include indexable semantic tag metadata to describe the context of each document. This context would be prebuilt and would help to better judge the document's relevance. Another researcher built a framework for extracting keyword "concepts" from documents that was more powerful and precise than simple text indexing, and a third actually designed a system for categorizing web documents in the Library of Congress classification hierarchy to help overcome the vocabulary mismatch problem of finding relevant documents.

## Similarity Measures

If a user can identify a relevant document, asking a search engine to "find more like this one" might be a powerful approach to web search. There was general agreement that no single similarity measure was always superior and that this was an open problem. A comparative study of general web document similarity measures was presented, and a system using the powerful WORDNET knowledge base was demonstrated to be highly effective in some applications. Finally, an interesting exploration into finding not just relevant but relevant and useful documents showed that active research in this area will be needed for some time to come.

### Query Result Ranking and User Feedback

Often, search engines will return many results, but most will be uninteresting to the user. If a search tool can learn preferences or contexts from the user, this information can be used to hone the search or document-ranking process. Because it is often difficult to get a user to explicitly rank pages, the challenge becomes building a user profile from implicit preference data. Techniques presented included using TFIDF similarity with web documents previously bookmarked by the user or using documents that exist in the user's "community" of web searchers to help define the user's interests.

The organizing committee for this workshop consisted of Kurt Bollacker (Internet Archive), Justin Boyan (NASA Ames), Lee Giles (NEC Research), Haym Hirsh (Rutgers University), and Steve Lawrence (NEC Research).

*Kurt Bollacker*  
Internet Archive

## Constraints and AI Planning

Interest in the use of constraint techniques for AI planning problems has grown considerably in the recent years. The goal of the AAAI-2000 Workshop on Constraints and AI Planning was to foster exchange between researchers from the AI planning and

the constraint programming communities. The workshop was attended by approximately 30 scientists, a large number of them researchers from the National Aeronautics and Space Administration (NASA). Besides the technical presentations, the program included two invited talks, two panel discussions, and a small poster session.

The workshop began with an invited talk by Robert Fourer (Northwestern University) on AI planning from an operations research perspective. Fourer compared and highlighted the differences between operations research, constraint programming, and local search approaches.

Kanna Rajan (NASA AMES) gave the second invited talk on experience gained with the planning system that operated NASA's *Deep Space 1* spacecraft during two experiments. It focused on the use of constraint techniques.

The technical presentations were of high quality and covered a broad range of approaches that apply constraint technology for planning, for example, constraint propagation within hierarchical planning, constraint-based planning that is not bound to any kind of maximal planning graph, and techniques for improved resource handling. Many approaches combined planning and scheduling.

The workshop's panel discussions proved interesting too. The first panel was on different frameworks for the use of constraints, which included a discussion on constraint programming versus propositional satisfiability versus integer linear programming. The second panel on future directions addressed topics such as the integration with scheduling as well as the market situation for planning applications. A forthcoming article is planned that will present and discuss the panel topics in more detail.

*Alexander Nareyek*  
GMD First

## Integration of AI and OR: Techniques for Combinatorial Optimization

The Workshop on Integration of AI and OR: Techniques for Combinatorial Optimization was held on 30 July

2000 in Austin, Texas, in conjunction with the Seventeenth National Conference on Artificial Intelligence (AAAI-2000). It was sponsored by the American Association for Artificial Intelligence.

The workshop, organized by James Crawford and J. Paul Walser, was the first of its kind held in conjunction with AAAI. It aimed at fostering research that crosses the traditional boundaries between AI and operations research approaches to combinatorial optimization and placed particular emphasis on practical applications.

The workshop program included an invited talk by William R. Pulleyblank, director of the Deep Computing Institute and Mathematical Sciences at IBM T.J. Watson Research ("Integer and Linear Programming for Logic Programmers"). In addition to the select technical presentations, two panel discussions were organized: (1) Applications of Integrated AI-OR Approaches, chaired by Martin Savelsbergh, and "Future Directions of Integrated AI-OR Research," chaired by John Hooker.

The number of registered participants was over 40, which demonstrates the sustained interest in the topic within the AI community. The organizers were particularly pleased about the high-profile participation from researchers of the operations research community as well as from practitioners from several of the leading optimization companies.

The panel discussions were valuable in identifying successful application areas of hybrid technology between constraint and integer-linear programming, in particular, in production scheduling and planning. Interestingly, there was broad agreement among researchers that the AI-OR community could benefit from both open CP/LP solver platforms and from a set of standard benchmarks.

Another topic in focus by panelists was how a university curriculum could support education in the AI-OR interface. This topic is becoming relevant given the consensus among participants that to make a living in the field of combinatorial optimization,

researchers and practitioners need a solid background in both operations research and AI.

The workshop's program committee consisted of Alexander Bockmayr (Universite de Nancy, LORIA), Crawford (i2 Technologies), Hooker (Carnegie Mellon University), Dana Nau (University of Maryland), Martin Savelsbergh (Georgia Institute of Technology), Bart Selman (Cornell University), Peter Van Beek (University of Alberta), and J. Paul Walser (i2 Technologies).

*J. Paul Walser*  
i2 Technologies

## Intelligent Lessons Learned Systems

The AAAI-2000 Workshop on Intelligent Lessons Learned Systems, which attracted 37 participants, was the first meeting of subject matter experts on lessons learned systems and AI researchers. Lessons learned systems concern the collection, validation, and dissemination of lessons among an organization's employees. Although lessons learned processes require knowledge management solutions that address both organizational dynamics and technological issues, this workshop focused primarily on the latter.

Currently, most of the (hundreds of) deployed lessons learned (and related) systems focus on dissemination using stand-alone, passive approaches. These are underused because they are divorced from the decision processes that the lessons are intended to support. Current relevant research issues include developing embedded approaches for eliciting, extracting, indexing, validating, and disseminating lessons. Maintaining existing lesson repositories is also of keen interest.

In this workshop, subject-matter experts presented overviews of their organizations' lessons learned processes-systems and identified existing problems, and AI researchers presented potential (technological) solutions to these problems. The audience was surprisingly well balanced: In addition to the 6 attendees from academic institutions, industry (15), govern-

ment laboratories (8), and active military (5) were well represented, showing a broad interest in this topic. The invited talks included John Bickford's description of the lessons learned process used at the Department of Energy (DoE), Kevin Ashley's survey on textual case-based reasoning (CBR), Peter Foltz's introduction to latent semantic analysis, and Commander John Moorman's description of the NAVY LESSONS LEARNED SYSTEM (NLLS). (The workshop WWW site at [www.aic.nrl.navy.mil/AAAI00-ILLS-Workshop](http://www.aic.nrl.navy.mil/AAAI00-ILLS-Workshop) contains slides from several of the workshop's presentations.) Although several of the other contributed papers discussed CBR approaches for lessons learned systems, rule-based, model-based, and other types of approaches were also presented. Panels, on the relation of CBR to lessons learned systems and a workshop-ending summary, and targeted discussion periods with preassigned leaders assisted with enhancing communication among the attendees.

What was learned or shared? First, many organizations face similar problems with operating lessons learned systems (for example, how to encourage lesson use, how to evaluate system utility). Second, several organizations (for example, Xerox, NLLS, DoE) have developed organization-wide lessons learned processes that can serve as prototypes for creating lessons learned processes in other organizations. Third, although AI approaches have rarely been incorporated into deployed lessons learned systems, practitioners are strongly encouraging the AI research community to team with them to enhance existing systems. Fourth, there is general agreement on the utility of integrating lessons learned systems with decision support systems that can benefit from the stored lessons. Finally, AI researchers have an excellent opportunity to contribute solutions on lesson elicitation and extraction; intelligent approaches to these tasks are in high demand but have not yet been deployed.

*David Aha*  
NRL

## TARK VIII

# Eighth Conference on Theoretical Aspects of Rationality and Knowledge

July 8–10, 2001

The Certosa di Pontignano, University of Siena, Italy

<http://www.tark.org>

The biannual TARK conferences bring together researchers from a wide variety of fields, including Artificial Intelligence, Cryptography, Distributed Computing, Economics, Game Theory, Linguistics, Logic, Philosophy, and Psychology, to further our understanding of interdisciplinary issues involving reasoning about rationality and knowledge. Topics of interest include semantic models for knowledge, belief and uncertainty, bounded rationality and resource-bounded reasoning, commonsense epistemic reasoning, logics of knowledge and action, formal analysis of games, applications of reasoning about knowledge and other mental states, belief revision, and the role of knowledge in general information flow. TARK has been influential in building interfaces between disciplines and accelerating new research trends, both through the actual event and the circulation of its proceedings. The 2001 conference, which will be held in Siena, Italy, will again serve this important purpose, with an exciting mix of contributed and invited talks. The conference is open to all; early pre-registration is encouraged. Check <http://www.tark.org> for further details.

### Knowledge-Based Electronic Markets

The AAAI-2000 Workshop on Knowledge-Based Electronic Markets assembled about 40 researchers who are focused on exploring the challenges, opportunities, and practical applications of knowledge-based electronic markets, or “e-markets.” This workshop followed up on the successful workshop entitled “Artificial Intelligence for Electronic Commerce,” held at AAAI in 1999 (see [agents.umbc.edu/aiec/](http://agents.umbc.edu/aiec/) for a record of this workshop).

E-markets include internet- or web-based markets where buyers interact and transact with sellers and are characterized by infrastructural-support and intermediary services and players, for example, yellow pages, catalogs, shopping search, advertising, sales assistants, brokers-aggregators, intermediaries, reputation-trust, authentication, and payments. The kinds of knowledge-based techniques that are being applied include knowledge rep-

resentation and reasoning services; machine learning; and communication-oriented services, for example, in systems of intelligent agents.

The day consisted of 12 presentations organized into 4 sessions. The first session included three papers on various aspects of collaborative filtering and recommendation systems applied to markets. The second session included papers on problems in managing supply-chain problems. The third session included papers and discussion of automated negotiation in electronic markets. A final session addressed the software infrastructure requirements for environments to support knowledge-based electronic markets.

A record of the workshop, including draft copies of the working papers, is available at [igec.umbc.edu/kbem/](http://igec.umbc.edu/kbem/).

*Tim Finin*

University of Maryland Baltimore  
County

*Benjamin Grosz*

IBM T.J. Watson Research Center

### Learning from Imbalanced Data Sets

The AAAI-2000 Workshop on Learning from Imbalanced Data Sets provided a venue for researchers to discuss fundamental questions pertaining to machine learning and challenge some of the field’s institutional practices.

Several observations were made, and certain issues were explored in particular depth. First, it was observed that a large number of applications suffer from the class imbalance problem. A distinction, nonetheless, was drawn between the small sample versus the imbalance problem, and it was remarked that although smart sampling can, sometimes, help, it is not always possible. Among the issues that received a lot of attention was the problem of evaluating learning algorithms in the case of class imbalances. It was emphasized that the use of common evaluation measures can yield misleading conclusions. More accurate measures include ROC curves and cost

curves. An evaluation measure was also proposed for the case where only data from one class are available. The other issues concerned the design of learning algorithms. It was shown that concept-learning methods can use a one-sided approach, focusing on either the majority or the minority class. If both classes are used, however, avoiding fragmentation in the minority class is useful. Another important issue concerned the close connection between the class-imbalance problem and cost-sensitive learning. Finally, the goal of creating a classifier that performs well across a range of costs-priors was declared to be an important one.

*Nathalie Japkowicz*  
Dalhousie University  
*Robert Holte*  
University of Ottawa

## Learning Statistical Models from Relational Data

The AAAI-2000 Workshop on Learning Statistical Models from Relational Data brought together researchers from diverse research areas, including machine learning, inductive logic programming, statistics, and databases. The workshop included nine paper presentations and two invited talks, "Relational Learning: Lessons Learned from the Web," by Tom Mitchell and Sean Slattery, and "Learning Probabilistic Relational Models," by Daphne Koller. The workshop closed with a roundtable discussion of potential application domains. Roundtable participants included Ted Senator and Henry Goldberg discussing structure discovery in financial data, Chris Manning discussing statistical natural language processing, Stephen Muggleton discussing structural molecular biology, and Lyle Ungar discussing recommender systems and DNA microarrays.

The bulk of the research presented at the workshop shared a common motivation: to uncover patterns and make predictions from structured data. However, there are multiple paths toward the common goal of statistical relational learning. One path begins with machine learning and statistical methods for "flat" or attribute-value representations and expands these approaches to incorporate relational

structure. However, a key assumption of many existing learning techniques—*independent and identically distributed instances*—might no longer hold; so, the naive approach of flattening structured data might introduce important statistical errors. A second path extends techniques for relational learning in nonprobabilistic domains, especially inductive logic programming, to incorporate stochastic models. This research area is active, and several new languages and learning algorithms have been proposed.

There was general consensus that a longer workshop should be held in the near future, allowing more time for discussion and synthesis of the many different approaches and applications. The workshop web page, [robotics.stanford.edu/srl](http://robotics.stanford.edu/srl), includes pointers to the workshop papers, related relevant papers, software, and, eventually, data sets and will have information about a mailing list and future events.

*Lise Getoor*  
Stanford University

## Leveraging Probability and Uncertainty in Computation

There is an increasing interest in computational approaches based on randomization, probability, and uncertainty. Such approaches can help to computationally model and manage resources more realistically and achieve significantly increased speed and robustness in problem solving.

To achieve this goal, probability and uncertainty can be used in various ways: First, randomized search methods, including stochastic local search algorithms as well as randomized systematic search techniques, have been successful in solving hard combinatorial problems more efficiently. Second, advanced probabilistic analysis techniques facilitate the empirical characterization of algorithm behavior, providing the basis for improved algorithm design, in particular, allowing the design of anytime strategies as well as more efficient mechanisms for allocation and management of computational resources. Third, methods and models for reasoning and problem solving under uncertainty facili-

tate representing and solving more realistic problems in which information can be incomplete or unreliable.

The workshop aimed to bring together researchers from different areas of AI and operations research to discuss various topics in randomization, stochastic search techniques, probability analysis of algorithms, flexible computation, and uncertainty.

The presentations spanned a broad range of topics, ranging from Monte Carlo methods and their application in AI to logic-based high-level robot control. The workshop contributions fell into five major categories: (1) randomization and adaptive strategies, (2) portfolios and cooperative methods, (3) approximation and probabilistic methods, (4) local search, and (5) applications. Additionally, two invited talks (by Eric Horvitz and Stuart Russell) provided higher-level views of the workshop topic and showed connections between the various ways in which probability and uncertainty are used in AI problem solving.

In a concluding discussion, a number of research problems that are common to various current approaches to uncertainty and probability in computation were identified and addressed. One of these research problems concerns the development of effective and reasonably efficient methods for predicting or estimating the computational cost for solving a given problem with a given algorithm. Such methods could provide the basis for flexible anytime computation and, more specifically, for the control of algorithm portfolios such that an improvement in the efficiency and robustness of solving certain problem types, such as combinatorial search problems, is achieved. Current research in this direction investigates how the performance of search algorithms depends on certain characteristics of the spaces searched. However, most of the measures used today are either quite unreliable or very expensive to compute. Because much of the research in this area is based on empirical studies, another important issue is the design and use of benchmark libraries that provide problem instances that reflect important characteristics of real-world instances in a controllable way and,

thus, provide a solid basis for further empirical research.

*Carla P. Gomes*  
Cornell University  
*Holger H. Hoos*  
University of British Columbia

## Mobile Robot Competition and Exhibition

Traditionally, the workshop for the AAAI Robot Competition and Exhibition is held the last day of the AAAI-2000 conference, after the robot competition events and exhibitions have concluded. Thus, the participants are able to discuss their actual entries in the robot events and talk about the results and lessons learned.

The events of the Ninth AAAI Robot Competition and Exhibition, held 30 July to 3 August 2000, included the popular Hors d'oeuvres, Anyone? and challenge competition events as well as a new competition event, Urban Search and Rescue. The exhibition included groups that wanted to demonstrate work outside the robot competitions. Students and faculty from University of Arkansas, Northwestern University, Universite de Sherbrooke, Swarthmore College, University of South Florida, and Kansas State University presented research related to contest events. Topics included architectural issues such as integration of high-level cognition and low-level behaviors, implementation of arbitrators for subsumptionlike architectures, innovative vision approaches for face tracking and recognition, and human detection using fused sensor data. Human-robot interaction issues and interface implementations were discussed as well.

Several participants of the exhibition—University of Oklahoma, University of Minnesota, Georgia Institute of Technology, and Utah State University—also shared their research with fellow conference participants. This year's exhibition entries represented research in vision and machine learning as well as mobility for autonomous agents.

Even though participants focused on the research and work done prior to the competition and exhibition, they also discussed the performance of

the systems during the contest; instances of both failure and success were heard. The future direction of each team's research was briefly mentioned as well. At the conclusion of the workshop, participants discussed various future improvements to the Mobile Robot Competition and Exhibition. This workshop gave the participants an opportunity to share their research with fellow researchers and get valuable feedback for future work. Each of the participants was invited to submit a paper to be published this fall in the workshop proceedings.

*Alan C. Schultz*  
NRL

## New Research Problems in Machine Learning

In the last 10 years, research in machine learning has been predominantly addressing a relatively narrow range of issues. This focus has helped the field to explore in depth some of the basic methods and approaches and has produced results that have attracted wide attention. Recently, however, a growing number of researchers came to believe that time might be ripe to start exploring issues that have not received due attention in the past. To strengthen these tendencies and act as a catalyst, this workshop brought together scientists willing to present their ideas and recent efforts investigating new directions in the field and participate in a brainstorming session exploring desirable topics for future research in machine learning and new challenges in response to the needs of the industry.

About 30 researchers participated in the workshop. The presented papers can roughly be divided into three major groups: The first group concentrated on attempts to contribute to a general theory of learning, including exploratory theory formation, induction of predictive compositional hierarchies, the problems of teaching agents by computer users, and other problems underlying the logic of learning. A talk by Giordana discussed a "phase transition" in relational learning and indicated severe theoretical difficulties in applying an inductive logic programming paradigm to this problem.



Another group of papers discussed the problem of mapping between representations, including mechanisms for metadata induction, perception-based abstraction for concept representation, and some open questions from the field of combining heterogeneous sets of classifiers. The last group targeted the challenges entailed by major real-world applications, for example, the problems related to automated discovery of structural patterns.

Michalski and his collaborators have reported results from a new research direction concerned with integrating learning and evolutionary computation and have shown that a machine learning-based guidance can significantly speed up evolutionary processes (in terms of the number of evolutions or births). He also indicated a need for research on what he calls *natural induction*, a form of inductive

learning that produces knowledge in the forms natural to people. Finally, Langley suggested several new research topics: induction of constraints, induction of configuration-related knowledge, learning of strategies for automated dialogs, and model construction. These last two speakers in particular raised questions that led to extensive discussions.

The workshop participants expressed the position that the field is at a crossroad and needs to undertake new topics to sustain its dynamic growth. Many new research problems have recently been raised in connection with the technologies spawned by the web. Further, techniques developed by machine learning have been applied to real-world problems, even though probably not to the extent that the community envisioned a few years ago. However, the growing understanding of empirical learning algorithms and the accumulation of experience in their effective implementation provide strong encouragement for researchers to undertake new ambitious projects that could lead to significant advances in the field of machine learning and beyond.

*Miroslav Kubat*

University of Louisiana at Lafayette

*Tom Mitchell*

Carnegie Mellon University

## Parallel and Distributed Search for Reasoning

Many AI systems, such as learners, planners, deduction systems, or expert systems, base their reasoning on intelligently searching in very large search spaces. *Intelligently searching* means that a system combines various pieces of general knowledge with different information found so far during its search and then uses this combined knowledge to decide where to search next. Because each search step can produce new and even surprising information, only very limited predictions about the search are possible.

For search in general, many concepts for systems have been developed to speed up searching by using several processing units. These concepts must deal with problems such as balancing the processor loads or avoiding idle

times, bottlenecks, and redundancy. Very often, these problems are solved by using a simple search control that leads to easily predictable search steps.

Thus, it is difficult for systems to use intelligent search while they take advantage of parallel or distributed search. Naturally, communicating all new information between the working units overcomes the problems, but too much communication drastically reduces the gains with the use of several processing units. Nevertheless, in the last years, several systems for different application areas have been developed that achieve the combined benefits of using several computing nodes and doing intelligent search. By cooperation, several search systems were often able to solve harder problem instances than each single system could when working alone. Although the systems used, in most cases, much domain-specific knowledge, certain types of knowledge and certain parallelization and distribution concepts can be identified that are usable for several application domains.

The goal of the AAI-2000 Workshop on Parallel and Distributed Search for Reasoning was to bring together researchers from various application areas that are interested in intelligent search using several processing units and the problems such a search has to face. We were happy that with Tad Hogg, Xerox PARC, and Sesh Murthy, IBM Watson Research Center, we had two invited speakers who have worked on general problems and general solution concepts for intelligent parallel and distributed search and who also have insight into using such search systems in real industrial applications.

In the talks, a wide variety of application domains were covered, ranging from different constraint-satisfaction problems over planning problems to electronic-commerce problems. The contributors addressed problems such as designing appropriate configurations of processing units, selecting and representing the information to be passed between the units, estimating the difficulty of tasks assigned to units, or selecting the most appropriate sequential search systems for a team using portfolio-management techniques.

In the final discussion, the need for

a forum for parallel and distributed search was expressed to give newcomers a starting point, facilitate discussions about the stated problems, and continue the task of bringing together researchers from the different application domains. As a starting point, the web site for the workshop ([www.cpsc.ucalgary.ca/~denzinger/aaai-ws.html](http://www.cpsc.ucalgary.ca/~denzinger/aaai-ws.html)) will be updated with more links and information.

*Joerg Denzinger*

University of Calgary

## Representational Issues for Real-World Planning Systems

The workshop brought together a diverse set of researchers from academia, research institutes, and the commercial world to discuss representational challenges for practical planning technology. Most of the participants reported on research that stemmed from efforts to develop planning systems for real-world problems; application domains included military planning, space, games, aircraft control, manufacturing, and logistics. A smaller number of attendees were drawn to the workshop as an opportunity to relate more abstract work to realistic problems.

Several attendees presented work on increased representational expressivity and reasoning techniques to support richer models of planning. Concepts covered by this work included preferences and advice, uncertainty, time-critical tasks, and nonsymbolic information. Others discussed specialized representations developed to support work on specific domains. User-centric planning systems, control for embedded systems, and plan ontologies also played prominent roles among the presentations.

Open discussions gravitated toward broader issues related to building and deploying applications of planning technology, during which attendees shared many interesting experiences and insights. Key challenges for the field were identified, among them scalability, knowledge validation, knowledge acquisition, and maintenance. Attendees also debated the viability of an off-the-shelf planning tool

kit (planner in a box) that would enable users to develop new planning applications with little assistance from technology experts. In discussing the issue of technology transition, several attendees remarked on the need to bring the user into the development process early on to increase the likelihood of buy in and acceptance. It was noted that early user feedback is essential to ensure that the technology solves the right problem and matches the user's conceptualization of the planning process.

Although central to the work of many of the participants, issues related to building and transitioning planning systems are not often emphasized in publications and meetings of the planning community. As such, the workshop provided a valuable forum for the exchange of practical experience and the discussion of challenges related to the development and deployment of real-world planning systems.

*Yolanda Gil*

USC/Information Sciences Institute

*Karen Myers*

SRI International

## Spatial and Temporal Granularity

This was the first international workshop set up with the goal of exploring key issues and major trends in spatial and temporal granularity. Being able to provide and relate spatial and temporal granularity representations at different grain levels of the same reality is an important research theme in AI, databases, and geographic information science (GIS) and a major requirement for many applications. The inherent abstraction power of the granularity concept has been exploited successfully in several application domains, including temporal and spatial reasoning, hierarchical planning, natural language understanding, temporal database design, database interoperability, data mining, medical informatics, image processing, hierarchical reasoning in GIS, real-time system specification, and verification.

The workshop was conducted as a combination of paper presentations, an invited talk, and a panel session. We counted about 20 attendees who

*The workshop provided a valuable forum for the exchange of practical experience and the discussion of challenges related to the development and deployment of real-world planning systems.*

actively participated with questions and remarks leading to stimulating discussions. The invited talk by Jerry Hobbs proposed an original perspective about the conceptual process of abstraction and discussed its impact on knowledge representation and reasoning. Paper presentations ranged from traditional themes, such as the representation of calendars and the description of topological relationships at different scales, to emerging trends, such as the treatment of granularity in spatial cognition and quantum mereotopology. Other interesting topics addressed during the workshop were the representation of time vagueness and uncertainty, the application of fuzzy theory in both the temporal and spatial domains, and techniques to refine-coarsen granularities in spatiotemporal representations. Finally, interesting applications in the context of temporal annotation of videos and temporal tagging of texts were also presented.

The workshop was concluded by a panel discussion led by Claudio Bettini, Anthony Cohn, Christian Freksa, Kathleen Hornsby, X. Sean Wang, and Jef Wijsen, with contributions from the attendees. Among the issues that arose from the discussion, we mention the following:

Dealing with granularity in the spatial domain seems to be inherently more involved than in the temporal domain. On the one hand, this probably depends on the inherently higher complexity of the spatial context, and on the other hand, this also reflects the fact that a more extensive and systematic investigation of granularity issues has been performed in the temporal context. In particular, the exis-

tence of a glossary for temporal granularity concepts seems to help in several respects, for example, to relate different approaches. A similar effort should be undertaken for the spatial and spatiotemporal domains.

It was argued that current approaches to temporal granularity are too general with respect to the actual needs of practical applications. Although different perspectives on this issue were proposed, it emerged that there are application domains, such as the medical one, that seem to require quite expressive models of temporal granularity.

There does not seem to be any major difference in requirements between the GIS and spatial reasoning communities. The same holds for AI and databases, with the possible exception of AI cognitive approaches that mainly deal with ontological issues.

Finally, it was discussed how reasoning complexity changes when changing the granularity of a representation and how it changes when reasoning with multiple granularities.

Overall, the participation at the workshop and the discussion suggest that the subject is relevant from both a theoretical and a practical point of view and that it deserves further investigation. It was also quite clear that research on spatial and temporal granularity would benefit from a joint effort from the AI, database, and GIS communities.

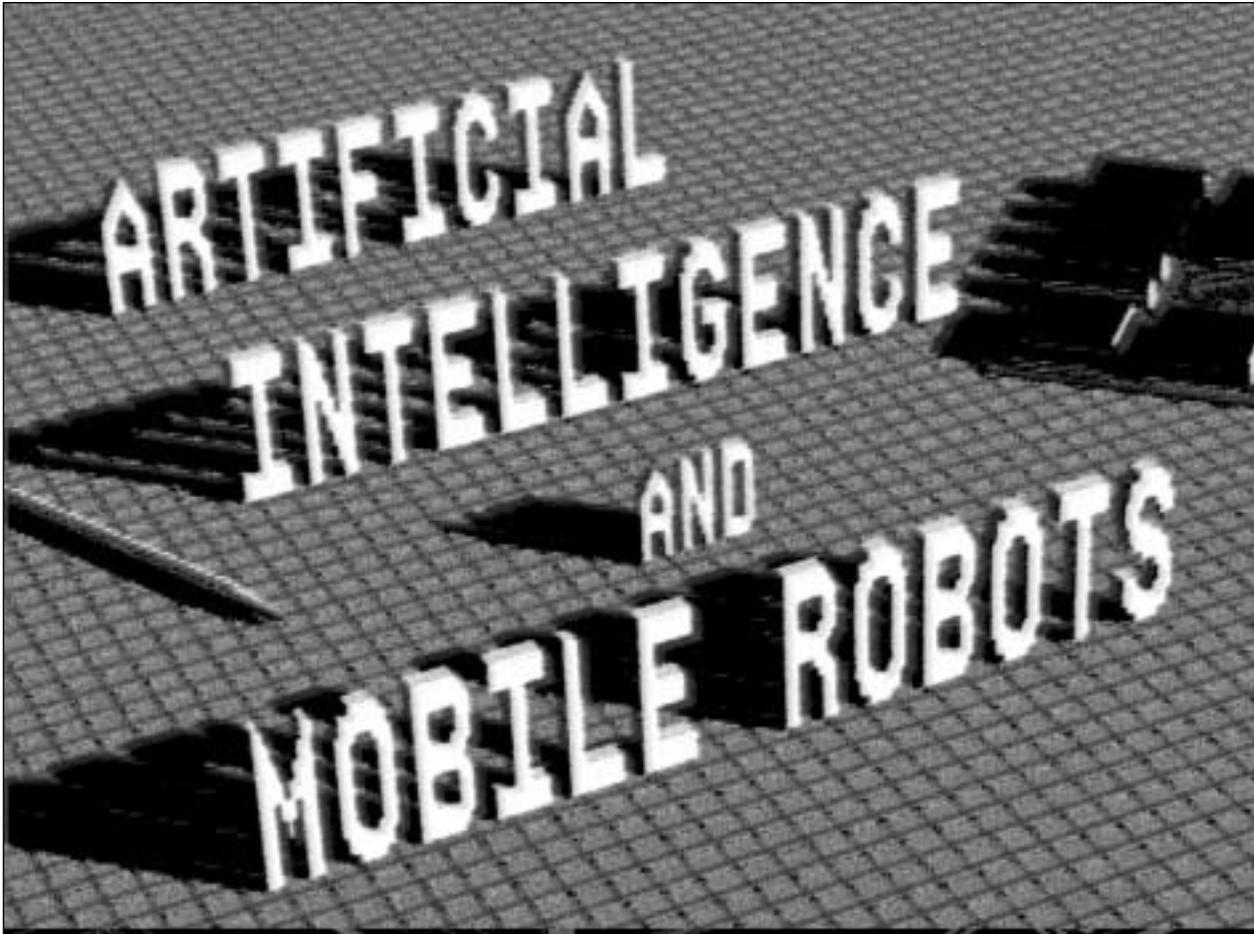
The workshop papers can be found in the AAI Technical Report WS-00-08.

*Claudio Bettini*

Università di Milano

*Angelo Montanari*

Università di Udine



## **Artificial Intelligence and Mobile Robots**

### **Case Studies of Successful Robot Systems**

*Edited by David Kortenkamp, R. Peter Bonasso, and Robin Murphy*

**T**he mobile robot systems described in this book were selected from among the best available implementations by leading universities and research laboratories. These are robots that have left the lab and been tested in natural and unknown environments. They perform many different tasks, from giving tours to collecting trash. Many have distinguished themselves (usually with first or second-place finishes at various indoor and outdoor mobile robot competitions).

Each case study is self-contained and includes detailed descriptions of important algorithms, including pseudo-code. Thus this volume serves as a recipe book for the design of successful mobile robot applications. Common themes include navigation and mapping, computer vision, and architecture.

6 x 9, 400 pp., ISBN 0-262-61137-6

**To order, call 800-356-0343 (US and Canada) or (617) 625-8569.  
Distributed by The MIT Press, 55 Hayward, Cambridge, MA 02142**