



AAAI 2002 Workshops

Brian Blake, Karen Haigh, Henry Hexmoor, Rino Falcone, Leen-Kiat Soh, Chitta Baral, Sheila McIlraith, Piotr Gmytrasiewicz, Simon Parsons, Rainer Malaka, Antonio Krueger, Paolo Bouquet, Bill Smart, Koichi Kurumantani, Adam Pease, Michael Brenner, Marie desJardins, Ulrich Junker, Jim Delgrande, Jon Doyle, Francesca Rossi, Torsten Schaub, Carla Gomes, Toby Walsh, Haipeng Guo, Eric Horvitz, Nancy Ide, Chris Welty, Frank D. Anger, Hans W. Guegen, and Gerald Ligozat

The American Association for Artificial Intelligence (AAAI) presented the AAAI-02 Workshop Program on Sunday and Monday, 28–29 July 2002 at the Shaw Convention Center in Edmonton, Alberta, Canada. The AAAI-02 workshop program included 18 workshops covering a wide range of topics in AI. The workshops were

- Agent-Based Technologies for B2B Electronic Commerce
- Automation as a Caregiver: The Role of Intelligent Technology in Elder Care
- Autonomy, Delegation, and Control: From Interagent to Groups
- Coalition Formation in Dynamic Multiagent Environments
- Cognitive Robotics
- Game-Theoretic and Decision-Theoretic Agents
- Intelligent Service Integration
- Intelligent Situation-Aware Media and Presentations
- Meaning Negotiation
- Mobile Robot Competition and Exhibition
- Multiagent Modeling and Simulation of Economic Systems
- Ontologies and the Semantic Web
- Planning with and for Multiagent Systems

-
- The American Association for Artificial Intelligence (AAAI) presented the AAAI-02 Workshop Program on Sunday and Monday, 28–29 July 2002 at the Shaw Convention Center in Edmonton, Alberta, Canada. The AAAI-02 workshop program included 18 workshops covering a wide range of topics in AI. The workshops were Agent-Based Technologies for B2B Electronic-Commerce; Automation as a Caregiver: The Role of Intelligent Technology in Elder Care; Autonomy, Delegation, and Control: From Interagent to Groups; Coalition Formation in Dynamic Multiagent Environments; Cognitive Robotics; Game-Theoretic and Decision-Theoretic Agents; Intelligent Service Integration; Intelligent Situation-Aware Media and Presentations; Meaning Negotiation; Multiagent Modeling and Simulation of Economic Systems; Ontologies and the Semantic Web; Planning with and for Multiagent Systems; Preferences in AI and CP: Symbolic Approaches; Probabilistic Approaches in Search; Real-Time Decision Support and Diagnosis Systems; Semantic Web Meets Language Resources; and Spatial and Temporal Reasoning.
-

- Preferences in AI and CP: Symbolic Approaches
- Probabilistic Approaches in Search
- Real-Time Decision Support and Diagnosis Systems
- Semantic Web Meets Language Resources
- Spatial and Temporal Reasoning

Technical reports containing the papers of these workshops are available from AAAI Press.¹

Agent-Based Technologies for B2B Electronic Commerce

With the inception of the internet, huge opportunities have arisen for intelligent automation of business-based transactions. Intelligent agent architectures have gained acceptance in implementing this domain and supporting electronic commerce and other internet-based transactions. This use of agents in electronic commerce has mostly been directed toward applications that support transactions between businesses and consumers (B2C). Another important domain is the automation of transactions that occur between businesses (B2B).

Historically, legacy B2B systems have operated over private networks that are both inflexible and expensive. The success and acceptance of B2C applications over the internet, coupled with the expense and inflexibility of legacy B2B systems, has led to the call for more research in B2B technologies that also occur over the internet. Moreover, agent researchers must rise to the challenge to offer agent-based solutions for the intelligent automation for B2B interoperability.

The Agents and B2B Electronic Commerce (AgentB2B) Workshop series aims to identify research progress in the B2B domain where agent researchers have investigated and initiated the development of agent-based B2B technologies and systems. This year's workshop continues the success of the AgentB2B 2001 Workshop held at the Autonomous Agents 2000 Conference in Montreal. AgentB2B '02 fostered strong collaboration among researchers in the area of agents and business-based internet transactions.

The workshop consisted of five scheduled papers, two invited talks, and an extremely active working session that aimed to categorize the state of agents in the B2B domain.

In the first invited talk, Norman Sadeh of Carnegie Mellon University's E-Commerce Institute discussed the institute's work in the area of automated electronic supply chain. He described their architecture where intelligent agents assist in the decision support for real-world supply-chain management. This work is important in determining the timeliness and quality when automating the processes of procurement, production, and distribution of goods.

The second invited speaker, who also presented a paper, was Mathias Salle of Hewlett Packard Laboratories. In both presentations, Salle discussed the need for agent architectures and framework support for electronic contracting. These agent-directed contract fulfillment protocols will be instrumental in electronic, business-oriented negotiations.

Other presentations reinforced the aforementioned areas. Blake presented survey work that suggested current concentrations of agent work that might be relevant to the progress of agents for B2B technologies. Aknine presented several agent-based protocols for internet-based business negotiations. Sandholm presented work where agents can moderate the clearing of the supply-and-demand curve when managing the connection of internet-based buyers and sellers. Finally, Wagner's work emphasizes the use of agents in dynamic supply-chain coordination and management.

The centerpiece of the event was the working session where all participants discussed issues surrounding the domain of agents and B2B electronic commerce. It was agreed that semantic web service interoperability would be important in the next generation of B2B coordination. Agents would be necessary to extend the areas of discovery, access, and composition. The use of agents would be a necessity in establishing long-term ontology translations among businesses. There was an interesting debate on the use of centralized versus

peer-to-peer negotiation models stemming from issues such as distributions, performance, and access to information. A final concern of all participants was the lack of real-world supply-chain and work-flow data. It was agreed that a major need in this domain is agent architectures (created by domain experts) that can produce large quantities of simulation data.²

— Brian Blake,
Georgetown University

Automation as Caregiver: The Role of Intelligent Technology in Elder Care

This symposium focused on the issues and design of intelligent systems for elder care. In previous years, much of this technology has occupied traditional assistive roles, including robotic walkers, robotic arms, and communication tools. This workshop, however, attracted researchers who are focusing on solutions that aid in monitoring, diagnosis, situation awareness, decision aiding, and training tools for either the elderly themselves or their caregivers.

Many of the papers described passive monitoring systems that use sensing in the environment to assess the elder's condition and generate appropriate responses. Several systems use more proactive devices to interact or support elders, including a mobile robot and a smart phone.

Significant focus was placed on supporting elders' cognitive limitations by providing reminders or prompting through tasks. User interfaces were also considered an important research focus, ranging from discussions of better hardware design to the modeling of emotion.

The first panel session of the day asked panelists to name the most important issue for researchers to address. Issues raised included handling caregiver burnout; handling the elder's sensory and cognitive decline over time; and responding to issues of loneliness, depression, and motivation. It was also pointed out that none of the papers had addressed significant medical problems such as arthritis or circulatory systems.

The second panel, about issues and ethics of research in this field, asked panelists to generate a "provocative" statement, whether or not the panelist agreed with it. These statements were (1) all data generated by monitoring systems should belong to the client, to give out as he/she pleases to doctors or monitoring agencies who should be required to sign non-disclosure agreements; (2) every decision made by the system should have a human in the loop who is responsible (liable) for the system's behavior; (3) system design must have professional ethicists and the target community involved so that issues of not only safety and privacy are addressed but also dignity; (4) for at least the next five years or so, there should be a certification board (not unlike the Federal Aviation Administration) that regulates system capabilities. Not surprisingly, this panel generated lively discussion.

Only a few of the presentations described systems that had been fielded or tested with the target community, suggesting that research in this area is very preliminary. Many opportunities exist for future work, both within the field of AI and without. AI topics include robotic assistive technologies, cognitive aids, passive monitoring, decision making, human factors, and adaptation to the changing elder and changing environment.

— Karen Haigh,
Honeywell

Autonomy, Delegation, and Control: From Interagent to Groups

This workshop followed the 2001 International Joint Conference on Artificial Intelligence workshop with the same title. The workshop contributed to understanding social agent interactions. This year, our focus was not only continuing to understand salient social notions in agent interaction but scaling issues of social networks on intergroup interactions. We explored theories that synthesized the interagent interaction into unified models as well as derived and implied attitudes that are beyond the immediate and direct in-

teragent attitudes. The latter plays a big role in balance of attitudes among agents in a group. As in 2001, we had researchers in multiagency as well as human-agent and human-robot interactions. In agent-agent interaction, the agents are designed to change their interaction to optimize local qualities such as cost or system qualities, for example, coherence. In organized groups, agents are designed to model the organizational structure, and the concerns are mostly with deontic concepts.

It is important for agents to understand human preferences and guidance in complex systems, which involves many issues, from abilities to comprehend issues related to delegation in natural language to the understanding of human emotions. Presentations covered several space systems and a large naval application. Discussions of applied research motivated the need for agents to explicitly reason about autonomy and delegation. Although there is a need for increased autonomy on the part of agents, there are times when autonomy is harmful, for example, when agents are fully autonomous, and their actions are uninterruptible. Humans in the loop of such agents can come to harm if they cannot interrupt actions of agents. Also, if agents have negative influence on one another, their independence can detract from their harmonious interaction. When tasks among agents are coupled, cooperating agents need to take one another's actions into account and not be self-interested.

– Henry Hexmoor,
University of Arkansas

– Rino Falcone,
National Research Council (Italy)

Coalition Formation in Dynamic Multiagent Environments

The workshop addressed a particular and ever increasingly important research issue in multiagent systems, that is, coalition formation in dynamic environments. Dynamic real-world environments pose difficult challenges for rational agents who might otherwise be capable of form-

ing optimal coalitions for problem solving under conditions of perfect information and unlimited computational resources. In dynamic environments where events are changing rapidly, information cannot be relayed among the agents frequently enough, or where centralized updates and polling are expensive, agents can be forced to form suboptimal coalitions.

The workshop had a collection of 10 papers, covering issues in emergence, real-time satisficing algorithms, hierarchical organizations, team reformation, long-term coalition, opportunistic coalition formation, and distributed optimization. Applications addressed included multisensor target tracking, robotic soccer, virtual space of mobile agents, helicopter transports, customer and vendor personal agents, and other simulated domains.

In our panel discussion, we talked about the following issues: (1) the restructuring and management of coalitions: Can coalitions be restructured in time? What motivates the restructuring? How should a coalition respond to changes? What role changes should be allowed by the agents in a coalition? (2) communication and coordination: What if message passing is noisy? Would handling noise be counterproductive? (3) the evaluation metric: What is a good coalition in real-time dynamic environments? How do you balance between the cost of forming a coalition and the solution quality of a coalition? What about intercoalition and intracoalition issues?

– Leen-Kiat Soh,
University of Nebraska

Cognitive Robotics

The Third International Workshop on Cognitive Robotics was held July 28, 2002, as part of the AAAI-02 workshop series. The workshop, cochaired by Chitta Baral (Arizona State University) and Sheila McIlraith (Stanford University), followed the very successful 1998 AAAI Fall Symposium on Cognitive Robotics held in Orlando and the Second International Cognitive Robotics Workshop,

held in conjunction with ECAI-2000.

Research in robotics has traditionally emphasized low-level sensing and control tasks, including sensory processing, path planning, and manipulator design and control. In contrast, research in cognitive robotics is concerned with endowing robots and software agents with higher-level cognitive functions that enable them to reason, act, and perceive in changing, incompletely known, and unpredictable environments. Such robots must, for example, be able to reason about goals, actions, when to perceive and what to look for, the cognitive states of other agents, time, collaborative task execution, and so on. In short, cognitive robotics is concerned with integrating reasoning, perception, and action within a uniform theoretical and implementation framework.

This workshop aimed at bringing together researchers involved in all aspects of the theory and implementation of cognitive robots to discuss current work and future directions. Although all aspects of cognitive robotics were of interest to the workshop, discussions and demonstrations of implemented systems were particularly encouraged.

The workshop was attended by approximately 35 researchers from a diversity of research communities related to cognitive robotics. Sixteen papers were accepted for poster presentation. Of these, a subset, deemed to be of broad or thematic interest, was selected for oral presentation. The workshop commenced with an invited talk by Ray Reiter (presented by Fiora Pirri) entitled "Stochastic Actions, Probabilities, and Markov Decision Processes in the Situation Calculus." This talk was followed by a presentation of INDILOG. The rest of the workshop was divided into thematic sessions entitled Programming and Controlling Cognitive Robots (Chair: Michael Thielscher), Object Recognition (Chair: Fiora Pirri), Diagnosis and Explaining Observations (Chair: Gerhard Lakemeyer), and Challenge Problems in Cognitive Robotics (Chair: Chitta Baral).

Each session comprised presenta-

tions by two or three participants with accepted papers. Sessions commenced with an overview of the challenges and state of the art in cognitive robotics posed by the session theme. These were presented by the session chair. Session participants were encouraged to go beyond the content of their workshop papers and place their work and research programs in a broader context. The workshop closed with a two-hour poster session in which participants were given an opportunity to discuss their work at length.

The workshop generated a lot of fruitful discussion. New to the 2002 workshop was the focus on object recognition. It generated a rousing debate on the role of object recognition in cognitive robotics. In the end, most participants agreed that it is an important component.³

- Chitta Baral,
Arizona State University
- Sheila McIlraith,
Stanford University

Game-Theoretic and Decision-Theoretic Agents

Over the last few years, decision and game theories have proved to be powerful tools with which to design autonomous agents and understand interactions in systems composed of many such agents.

Decision theory provides a general paradigm for designing agents that can operate in complex uncertain environments and can act rationally to maximize their preferences. Decision-theoretic models use a precise mathematic formalism to define the properties of the agent's environment, the agent's sensory capabilities, the ways the agent's actions change the state of the environment, and the agent's goals and preferences. The agent's rationality is defined as behavior that maximizes the expectation of the degree to which the preferences are achieved, and the planning problem is identified as a search for the rational plan.

Game theory adds to the decision-theoretic framework the idea of multiple agents interacting within a com-

mon environment. It provides ways to specify how agents, separately or jointly, can change the environment and how the resulting changes impact their individual preferences. Building on the assumption that agents are rational and self-interested, game theory uses the notion of Nash equilibrium to design mechanisms and protocols for various forms of interaction and communication that result in the overall system behaving in a stable, efficient, and fair manner.

The workshop brought together researchers interested in theoretical aspects of these theories as well as those interested in applications. On the theoretical side, a new formalism called COM-MTDP, was proposed. An extension of Markov decision processes to communicating agents operating as a team, it offers a new view on formally modeling communication and presents new challenges to designers of efficient solution algorithms. Other papers brought up some of the fundamental issues of limitations of individual rationality, approaches to bounded rationality, logical approaches to decision making, and learning in stochastic games.

Many papers reported on work that uses game theory techniques to design multiagent protocols and mechanisms, including pricing in combinatorial economies, auctioning contracts, and false-name proof auction protocols. Two other related papers presented the use of heuristics to analyze complex interactions, such as a double auction, and the role computational limitations play as agents try to determine their valuations for goods during auctions.

Further, some promising applications were presented. They addressed solving complex resource-allocation problems using decision-theoretic planning, air traffic control, and management using Bayesian games and designing agents capable of efficient automated trading in realistic marketplaces.

- Piotr Gmytrasiewicz,
University of Illinois
- Simon Parsons,
University of Liverpool

ISAMP: Intelligent Situation-Aware Media and Presentation

The design of intelligent situation-aware media and presentations becomes increasingly important for software designers and also for the design of consumer electronic devices. This process results from the fact that software becomes more complex, and systems become integrated into a ubiquitous information infrastructure. The paradigm of "one device-one function" is over, and today we can access mutually any service through any device. A television set, for example, is today a control center for various applications, for example, video programming, electronic commerce, and web browsing. A mobile personal digital assistant can be used as a telephone, tourist guide, remote control, calendar, or web-access device. In these scenarios, new intelligent methods are necessary that decouple the application logic from the actual device and that adapt the media and presentation to the actual situation, that is, the user with their needs and preferences, the devices that are available, and the context. In the last few years, a number of approaches addressed parts of these problems involving various AI techniques. In fact, AI techniques seem to be the key methods for intelligent and scalable solutions. However, there is currently no unified framework visible.

The presentations at the workshop covered a broad range of these topics. One prominent issue was the generation of multimodal presentations for different situations not only looking at technical configurations (for example, small screens, different modalities) but also at the cognitive abilities of the users (for example, working memory load and reading skills). The presentations addressed the problem of context and dialogue modeling as well as questions on how to design situational user interfaces (for example, for human-robot interaction). Another block of presentations looked at smart artifacts and their ability to adapt to the user's situation. Smart chairs, rings, and caps were

presented that (1) monitor and classify the user's context and (2) react to the user's location and goals. The excellent presentation of the invited speaker Hideyuki Nakashima (Cyber Assist Research Center, Japan) gave relevant insights into the new research area of intelligent artifacts and intelligent environments.

Overall, the workshop was a first step in looking at intelligent information presentation systems that adapt their dialogue behavior to the particular situation of users in instrumented environments. The workshop demonstrated that the areas of AI and ubiquitous computing have significant overlap with similar questions, thus encouraging us to plan similar events for the future.

- Rainer Malaka,
European Media Lab
- Antonio Krüger,
Saarland University

Meaning Negotiation

The development of distributed applications over large networks of computers raises the issue of semantic interoperability between autonomously developed sources of information or service providers. The problem seems especially critical in applications such as the semantic web, knowledge management, web services, marketplaces, ambient intelligence, and mobile business.

A common strategy for dealing with semantic interoperability is to create large, shared conceptual schemas (often called ontologies) that are then used as a common reference for agents using autonomous (and, typically, heterogeneous) conceptualizations. However, this approach shows many limitations, both from a theoretical and a practical point of view.

Meaning negotiation is an alternative approach to the problem of semantic heterogeneity based on the idea that any real-world approach to semantic interoperability between autonomous entities (namely, entities that cannot assess semantic problems by "looking into each other's head," such as humans or software agents) should involve (among other things)

a social process of negotiating an agreement on the content (semantics) and the speaker's intention (pragmatics) of a communication.

The workshop was a first attempt at bringing together people from different communities to discuss research issues, models, approaches, and techniques involved in this area of meaning negotiation. The main outcomes of the workshop were from a theoretical perspective: The participants "negotiated" some broad distinctions between various forms of the problem, for example, between the problem of meaning coordination (namely, reaching an agreement on the meaning of an expression when agents have a common interest in achieving such an agreement, but there are many possible solutions to the problem; thus, the selection of one of these solutions can be problematic) and the problem of meaning negotiation (reaching an agreement on the meaning of an expression when an agreement is valuable for all agents, but agents have conflicting preferences over which solution should be selected, so that every agreement implies that at least someone has to concede to some extent to other agents); from an organizational perspective, the participants agreed on the project of designing a general test bed to be used as a benchmark for any proposed protocol or algorithm of meaning negotiation. It was also agreed that it would be important to organize a follow-up event in the near future, possibly as an 2003 International Joint Conference on Artificial Intelligence workshop.⁴

One of the most valuable aspects of the workshop was the variety of disciplines (AI, philosophy, psychology, linguistics) and research areas (for example, logic, human-computer interaction, natural language processing, agents, epistemology) represented at the workshop. The discussion was very lively, and the participants expressed their satisfaction for attending the workshop.

- Paolo Bouquet,
University of Trento

Mobile Robot Competition and Exhibition Workshop

The Mobile Robot Competition and Exhibition Workshop is traditionally held on the last day of the conference and is aimed at the participants of the Mobile Robot Program events. The main purpose of the workshop is to advance the state of practice in mobile robotics by sharing techniques, ideas, and experiences in the context of the Mobile Robot Program events. There were four events this year: (1) Robot Challenge, (2) Robot Host, (3) Robot Rescue, and (4) the exhibition.

Approximately 20 teams presented papers at the full-day workshop. The papers covered all aspects of the systems fielded in the competition and exhibition events, covering a range of issues from low-level algorithm details to robot-human-interaction design.

This year's workshop also included an invited talk from Chief Ron Rogers of the Hillsborough County Fire Rescue. Rogers was involved in the rescue attempts at the World Trade Center and gave his reactions to the Robot Rescue event and suggestions for making it more relevant to real-world search-and-rescue efforts.

- Bill Smart,
Washington University

Multiagent Modeling and Simulation of Economic Systems

The AAAI-02 Workshop on Multiagent Modeling and Simulation of Economic Systems (MAMSES-02) was held on 29 July in Edmonton, Alberta. It focused on multiagent modeling framework and simulation method of economic systems. Throughout the workshop, the term *economic system* is used as a weakly connected multiagent system where chain of exchange exists; that is, goods, currency, credit, and information are continuously exchanged in the system. The target system under analysis in the workshop was not limited to financial market or whole

economy systems. It covered systems where the chain of exchange exists, such as market or auction mechanism for scheduling or planning, supply chain, game-theoretic analysis of exchange, and social coordination in ubiquitous computing.

The workshop program included an invited talk by Michael P. Wellman from the University of Michigan ("Explorations in Trading Strategy Spaces") and presentations of 15 selected technical papers, which referred to the various aspects of this research area, for example, world trade and supply chain as multimarket structure, social coordination, agent strategies in trading, theoretical analysis of markets, simulation of financial systems, and computational framework and tools. The workshop revealed the importance and effectiveness of AI for economy and vice versa. In addition, potential applicability of economic problem solving in wider areas was revealed.

The workshop was attended by approximately 25 participants, half of whom were interested in introducing market or auction mechanism in AI context, and the rest were interested in modeling and simulating economic system as a multiagent to investigate the essence of target systems.

– Koichi Kurumatani,
*National Institute of Advanced
Industrial Science and Technology*

Ontologies and the Semantic Web

This workshop addressed ontologies for the semantic web. The notion of the semantic web, as promoted by Tim Berners-Lee and Jim Hendler as the focus of a major Defense Advanced Research Projects Agency-funded research effort, is to transform the current World Wide Web so that the information and services are understandable and useable by computers as well as humans. The semantic web will create an environment where software agents can readily perform sophisticated tasks and help humans find, understand, integrate, and use information. The key distinguishing feature of the semantic web will be ontologies that

will enable software agents to find the meaning of the information on web pages by following hyperlinks to definitions of key terms and rules for reasoning about them logically. Presenters discussed two large ontologies, the DAML language, several tools for search, tools and methodologies for improving ontologies, and applications of ontologies to areas as diverse as education and autonomous vehicles. The workshop was held with 10 papers presented to approximately 45 participants. Richard Fikes (Stanford University) presented an overview of the semantic web in his keynote talk, and David Martin (SRI) presented an invited talk on progress on the DAML services effort. There was a lively discussion following the presentations. One issue was the search for one or more applications that will drive exponential growth and adoption of the semantic web. These applications have to be powerful enough to make clear that the current web is not good enough for people's needs and that the semantic web offers compelling advantages. Prior shifts in technology approaches have shown that the drivers for adoption of new technology seem to come from areas very different than the technology adopters expected. One suggestion was that the online game industry, with its rich characters and environments, and emerging open architectures, might soon need the powerful expression languages and extensive ontologies that are envisioned for the semantic web. Another discussion topic results from the tension between the knowledge representation community and the mass-market concerns of the web. Many knowledge representation researchers advocate an expressive representation language is needed with relations of arbitrary rules, quantification, and possibly modal or non-monotonic features. Others argue that getting the market to adopt something even as simple as RKF is a major challenge, apart from the technical issues of allowing expressive knowledge representation languages in a distributed environment that is far less controlled than the re-

search environments that such languages have typically been employed and studied. Ultimately, this issue is one of sociology that will be market driven.

– Adam Pease, *Teknowledge*

Planning with and for Multiagent Systems

Multiagent systems have become an important subfield of AI, and several classical AI topics are extensively studied in their distributed variants now. The aim of the AAAI 2002 workshop entitled "Planning with and for Multiagent Systems" was to discuss the challenges that the classical AI topic of planning has to face in domains where multiple agents can act concurrently. Being of interest to both multiagent systems and planning researchers, the workshop was held at AAAI to attract participants from more than a single research community.

Although there has been considerable interest in multiagent planning for a long time, existing approaches differ largely, for example, in their emphasis on either the distributed planning or the distributed plan execution process, in the ways communication and perception are used, and in whether a global plan for all agents or a local plan for each agent is produced. Thus, the workshop was also intended to present and discuss the common and differing research goals of the participants.

As expected, a broad range of work was presented on topics as varied as centralized planning for complex concurrent execution models, distributed merging of hierarchical multiagent system plans, planning team formation, dynamic negotiation, or multiagent partially observable Markov decision processes (POMDPs). In addition to the 11 papers selected for presentation from the submission, 4 invited talks gave further insight into general teamwork concepts, the role of communication in multiagent system planning, the integration of planning and coordination, and the related field of distributed constraint satisfaction.

Although the presentations

showed the diversity of the field, three open discussion sessions centered on the search for common ground. It became apparent that no single definition of the multiagent planning problem or the solutions searched for exists or is desirable at the moment. Still, some basic dimensions for multiagent planning variants could be identified that might be helpful for comparing and learning from differing approaches.

As a whole, the workshop showed a fair cross-section of current work in a highly dispersed area. A lively discussion at the end of the workshop showed that the view on planning in multiagent systems is rather different in the multiagent systems and the planning communities. Still, interesting developments in both fields, many of which were presented in this workshop, should attract and deepen mutual interest in the future.

- Michael Brenner,
Albert-Ludwigs-University
- Marie desJardins,
University of Maryland

Preferences in AI and CP: Symbolic Approaches

Preferences have been gaining more and more interest in different subfields of AI such as nonmonotonic reasoning, constraint satisfaction, planning, and configuration. We therefore organized a dedicated workshop on preferences at AAAI-02 in Edmonton to discuss questions of common interest and get an overview on different symbolic and qualitative methods for treating preferences.⁵ This workshop received excellent response, confirming the high interest in preferences. Thanks to the help of an excellent program committee, we could offer an interesting workshop program consisting of 16 presentations on topics such as preferences and soft constraints, logical approaches, preferences and agents, decision theory, and preference elicitation. We also thank AAAI for allowing us to extend the workshop duration to one and a half days.

The workshop started with an invited talk by Jon Doyle. He gave a survey of his broad work on prefer-

ences covering many challenging questions and promising perspectives of preferences. He stressed that preferences can have different roles such as representing desiderata, controlling reasoning and action, or establishing belief states and preferred solutions. In addition, preferences can themselves be objects of reasoning. Because different roles require different representations, it might become clear why there is such a variety of preference formalisms, ranging from utility functions and graphic representations to logical approaches. Doyle also discussed questions of the structure of preferences and changes of preferences. He finished his talk with an intriguing mechanical perspective about understanding preference change.

The technical presentations focused on questions such as how to acquire preferences without requiring too much user input, how to translate preference representations, and how to use preferences efficiently together with constraint-solving algorithms or logic programming. Furthermore, a first classification scheme for different preference-handling approaches was presented. Even if it is limited to nonmonotonic logics, this kind of work is very helpful to compare approaches and build bridges between them.

The talks applied preferences to problems such as travel configuration, scheduling of spacecraft activities, and multiagent meeting scheduling. Interestingly, preferences find a large application field in the area of multiagent systems where they can reduce the amount of information that is exchanged among agents. Thus, the workshop brought together many interesting papers about different aspects of preferences and provided rich material for study. It provided many challenging questions, such as how to implement an effective and user-friendly preference elicitation mechanism or how to elaborate classification schemes or criteria for comparing different approaches for dealing with preferences.

The participants of the workshop were all very satisfied with the event, and many stated that they have

learned many new things. We thus have reasons to believe that the workshop might stimulate further work in this area.⁶

- Ulrich Junker, *Ilog*
- Jim Delgrande,
Simon Fraser University
- Jon Doyle,
North Carolina State University
- Francesca Rossi,
University of Padova
- Torsten Schaub,
University of Potsdam

Probabilistic Approaches in Search

There has been considerable interest in approaches based on randomization, probability, and uncertainty to speed up computation and model resources more realistically. For example, the performance and robustness of search procedures can often be improved by adding an element of randomness to the search procedures, as in randomized backtrack search and local search. Furthermore, probabilistic methods such as Markov decision processes, Monte Carlo sampling, and Bayesian learning are now being used to study and improve the behavior of search procedures.

This AAAI-02 workshop (which followed up on a related AAAI-00 conference workshop and a AAAI-01 Fall Symposium) explored the use of probabilistic methods in the modeling and understanding of search procedures as well as their role in improving search procedures. The workshop was cosponsored by the Intelligent Information Systems Institute (IISI) at Cornell University, and its generous support allowed us to bring in Joao Marques-Silva to give an excellent invited talk on the different uses of randomization in backtrack search SAT algorithms.

The rest of the workshop was divided into four themes, each of which contained a number of different presentations. In the first theme, “extensions to SAT,” both multivalued and stochastic extensions to SAT were described as well as dynamic restart policies for randomization and restart methods. In the second theme, “parallel search,” stochastic issues in dis-

tributed and parallel applications were addressed. The third theme, "heuristic and hybrid search," covered stochastic methods in combinatorial problems such as constraint satisfaction and graph coloring. Finally, the last theme, "learning," considered methods both to learn probabilities and speed up stochastic search. A special issue on this topic is planned for the near future, and announcements will be made about this soon.⁷

- Carla Gomes,
Cornell University
- Toby Walsh,
University of York

Real-Time Decision Support and Diagnosis Systems

The purpose of this full-day workshop, jointly hosted by AAAI, ACM SIGKDD, and AUAI, was to address common issues in resource-bounded decision making, especially uncertain reasoning with deadlines. The 11 technical presentations (six 30-minute, five 10-minute) covered topics ranging from anytime, any-space embedded systems to new approaches to robust search and constraint satisfaction. The workshop emphasis was on probabilistic techniques and representation by graphic models, both for inference in diagnosis and monitoring and reinforcement learning in control and planning.

Three invited addresses were given: Bruce D'ambrosio surveyed a methodological spectrum between reactive and introspective approaches to decision making. A second invited talk by Michael Horsch dealt with fine-grained control of reasoning in constraint-satisfaction problems, and a third by Sven Koenig presented new real-time heuristic search techniques for fast dynamic replanning.

The first track focused on monitoring and diagnosis of failures in complex systems such as distributed systems (Irina Rish) and assembly lines (Cem Baydar). Cozman presented his work on space-time trade-offs for Bayesian network inference techniques in real-time embedded systems. The second track focused on

control in new real-time reinforcement learning architectures (Anita Raja), scheduling (Lloyd Greenwald), and sensor integration (Qiuming Zhu). The third track included one paper on localization in real time by particle filtering (Kwok) and one on real-time airport scheduling (Andres Gomez de Silva Garza). The fourth and final track covered a novel formulation of image analysis planning in time-constrained domains as a reinforcement learning problem (Vadim Bulitko) and a prediction-based method for online inference in graphic models—CP-nets (Carmel Domshlak).

The workshop concluded with a brief discussion on common interests and themes in real-time decision support and diagnosis systems (RTDSDS) and remarks by the cochair (Hsu) on representations useful in interactive RTDSDS.

Of 18 original technical publications submitted for review, 10 were published as full-length papers and the remainder as abstracts.

Workshop minutes, slides, supplementary notes, and continuing discussions can be accessed on the workshop web site.

- Haipeng Guo,
Kansas State University
- Eric Horvitz, *Microsoft*
- William H. Hsu,
Kansas State University

Semantic Web Meets Language Resources

The AAAI-2002 Workshop on Semantic Web Meets Language Resources met on 28 July in Edmonton, Alberta. The aim of the workshop was to bring together researchers in natural language processing (NLP) who are interested in language resources such as large corpora, lexical databases, terminologies, and formats for linguistic annotations with researchers in the so-called semantic web because the NLP researchers seem natural consumers of the work of the semantic web researchers. The two organizers, Nancy Ide and Chris Welty, represented each of these respective disciplines. Unfortunately, although the workshop was well attended by

representatives of the language resources community, very few representatives of the semantic web community were present.

The workshop began with two keynote presentations intended to provide an overview of each area and give some context for the presentations and discussions to follow. The keynote presentation on the semantic web was given by Deborah McGuinness of Stanford University's Knowledge Systems Lab, who has been instrumental in designing the Defense Advanced Research Projects Agency agent markup language (DAML+OIL), which is the major influence on the W3C standard web ontology language (OWL). McGuinness discussed some of the history of OWL and its major features as well as several use cases from medical informatics and thesauri. Nancy Ide of Vassar College, who has been involved in the development, representation, and standardization of language resources for 15 years, then gave a keynote presentation on new directions in linguistic resources, discussing the history and providing examples of linguistic resources and outlining a new ISO effort to develop standards for representing linguistic annotations.

The workshop was organized to maximize discussion time (each 20-minute presentation was slotted 45 minutes), and the organizers allowed the schedule to shift rather than stifle interesting interaction between these communities that rarely meet. Several interesting things arose from the discussions, such as the need to represent "endangered" languages, the need for more expressiveness than OWL provides, the need for ontologies that can capture the kinds of subtleties and distinctions that arise in language processing, the need for standardized semantics and annotations, and a better understanding of just what exactly constitutes the semantic web.

It was clear from the discussion that the two communities do not understand each other very well, primarily because they are approaching the issue of ontology development and use from very different perspec-

tives. In general, it appeared that semantic web researchers, on the one hand, are concerned primarily with the mechanics of semantic web development, ontology construction, and means to exploit of ontological information. Those developing language resources, on the other hand, are typically engaged in this activity to support language processing applications, which to be effective require precise and, to some extent, flexible semantics to capture meaning. Their concerns, therefore, are primarily with developing adequate semantic content models and representing them in ways that are best suited for use in language-processing applications such as information retrieval, extraction, and summarization and machine translation. At times during the discussion, it was obvious that the two communities will need to continue to talk to ensure that the mechanisms under development within the semantic web community can support the kinds of semantic adequacy required for language processing.

One thing was clear: Language resources must effectively cover all linguistic knowledge, and the complexity of this information poses a fundamental problem for knowledge representation. Furthermore, for the purposes of language processing, ontological information must be represented in such a way that it can be adapted according to context. Therefore, the needs of the language resources community pose a very large and interesting challenge for representation in the semantic web. At the same time, they are a natural fit for what the semantic web aims to support and can therefore drive further requirements as well as establish clear use cases.

Linguistic researchers learned from the workshop that the semantic web is not an exclusive community, and the framework that is being built is needed to provide the kind of common infrastructure that they need. Hopefully, the two communities can work together in the future to achieve both their goals.

- Nancy Ide, *Vassar College*
- Chris Welty, *IBM*

Spatial and Temporal Reasoning

The AAAI-02 Workshop on Spatial and Temporal Reasoning continued a series of such activities over the last 10 years in spanning related communities of researchers that study representing and reasoning about either space or time—or both.

Various aspects of representational problems in space (direction, location, proximity, geometry, intersection) and time (coincidence, order, concurrency, overlap, granularity) have attracted repeated attention because of their fundamental and difficult nature. Likewise, common reasoning problems have threaded their way through many papers on space (path finding, orientation, relative position) and time (constraint satisfaction, schedule optimization, precedence). The richness of different ontologies, different applications, and objectives has proved a persistent challenge to the community.

The papers were equally divided between spatial and temporal problems, with each paper taking on a different aspect of the challenge. Selecting a route, reasoning about angular directions, and determining if one object contains another occupied the spatial research, dealing with mixed granularity, solving changing constraints efficiently, and general reasoning using a temporal logic summed up the temporal research. On the spatial side, new tools were proposed (such as action vectors for understanding route descriptions in natural language), as well as new formalisms generalizing previous ones (such as qualitative reasoning about arbitrary angular directions), and new problems were addressed (such as representing various notions of “containment” for reasoning about object access). On the temporal side, systematic and general algorithms for reasoning about granularity were presented, and new results and tableaux algorithms for reasoning in tractable fragments of Joseph Halpern and Yoav Shoham’s logic was introduced and proved sound and complete.

Besides exhibiting new applications of known techniques, as well as

new approaches to well-studied applications, the presentations, questions, and discussions evidenced a growing interest in developing integrated methods that could result in general-purpose tools and theories subsuming several aspects of space and time.

- Frank D. Anger,
National Science Foundation
- Hans W. Guesgen,
University of Auckland
- Gerard Ligozat,
Paris-Sud University

Notes

1. Information about all AAAI workshop technical reports is available at www.aaai.org/Press/Reports/Workshops/workshop-reports.html.
2. Further information on this workshop and others can be accessed at www.cs.georgetown.edu/~blakeb/-AgentB2B. A future workshop is planned for WETICE 2003.
3. For further information about CogRob2002 or information about the venue of the Fourth International Workshop on Cognitive Robotics, see www.ksl.stanford.edu/cogrob2002.
4. The workshop web site at boogie.cs.unitn.it/AAAI-02-MN will be kept alive to coordinate these future efforts.
5. See www.cs.sfu.ca/Conf/prefs02/.
6. The proceedings are available as AAAI Workshop Technical Report WS-02-13, published by AAAI Press.
7. The workshop web site (with links to these previous meetings) is at www-users.cs.york.ac.uk/~tw/aaai02.