

The Workshop Program at the Twentieth National Conference on Artificial Intelligence

*Diego Molla Aliod, Eduardo Alonso,
Srinivas Bangalore, Joseph E. Beck, Bir Bhanu,
Jim Blythe, Mark Boddy,
Amedeo Cesta, Marko Grobelink,
Dilek Hakkani-Tür, Sanda M. Harabagiu,
Alain Lege, Deborah McGuinness,
Stacy Marsella, Natasha Milic-Frayling,
Dunja Mladenic, Dan Oblinger, Paul E. Rybski,
Pavel Shvaiko, Stephen Smith, Biplav Srivastava,
Sheila Tejada, Hannes Vilhjálmsón,
Kristinn R. Thórisson, Gokhan Tur,
Jose Luis Vicedo, and Holger Wache*

■ The AAAI-05 workshops were held on Saturday and Sunday, July 9–10, in Pittsburgh, Pennsylvania. The thirteen workshops were Contexts and Ontologies: Theory, Practice and Applications, Educational Data Mining, Exploring Planning and Scheduling for Web Services, Grid and Autonomic Computing, Human Comprehensible Machine Learning, Inference for Textual Question Answering, Integrating Planning into Scheduling, Learning in Computer Vision, Link Analysis, Mobile Robot Workshop, Modular Construction of Human-like Intelligence, Multiagent Learning, Question Answering in Restricted Domains, and Spoken Language Understanding.

The AAAI-05 workshops were held on Saturday and Sunday, July 9–10, in Pittsburgh, Pennsylvania. The co-chairs of the AAAI-05 Workshop Program were Adele Howe, Colorado State University and Peter Stone, The University of Texas at Austin. The fourteen workshops were Contexts and Ontologies: Theory, Practice and Applications (held Saturday, July 9; Pavel Shvaiko and Deborah McGuinness, co-chairs); Educational Data Mining (held Sunday, July 10; Joseph E. Beck, chair); Exploring Planning and Scheduling for Web Services, Grid and Autonomic Computing (held Saturday,

July 9; Biplav Srivastava and Jim Blythe, co-chairs); Human Comprehensible Machine Learning (held Saturday, July 9; Dan Oblinger, chair); Inference for Textual Question Answering (held Saturday, July 9; Sanda M. Harabagiu, chair); Integrating Planning into Scheduling (held Sunday, July 10; Mark Boddy, chair); Learning in Computer Vision (held Sunday, July 10; Bir Bhanu, chair); Link Analysis (held Sunday, July 10; Dunja Mladenic, Natasha Milic-Frayling, and Marko Grobelink, co-chairs); Mobile Robot Workshop (held Wednesday, July 13, Sheila Tejada and Paul E. Rybski, co-chairs); Modular Construction of Human-Like Intelligence (held Sunday, July 10; Kristinn R. Thórisson, chair); Multiagent Learning (held Sunday, July 10; Eduardo Alonso, chair); Question Answering in Restricted Domains (held Sunday, July 10, Diego Molla Aliod, chair); Spoken Language Understanding (held Saturday, July 9; Gokhan Tur (chair).

Contexts and Ontologies: Theory, Practice, and Applications

*Pavel Shvaiko, Deborah McGuinness,
Holger Wache, and Alain Leger*

During the last decade, there was a series of successful workshops and conferences on the development and application of contexts and ontologies. Early workshops focused mostly on identifying what contexts and ontologies are and how they can be formalized and exploited. More recently, with the emergence of distributed systems (such as P2P systems and the semantic web), the focus has shifted toward issues of practical applications, such as semantic integration, coordination, and meaning negotiation among information sources, where both contexts and ontologies were applied as promising solutions. However, few, if any, of these meetings focused on combining the themes of ontologies and contexts and discussing them as complementary disciplines.

This contexts and ontologies workshop aimed to bring together people from the context and ontology communities and facilitate discussion

about research and approaches to information integration, thereby highlighting different perspectives and making the meeting of these communities mutually beneficial. The workshop pushed the cross-fertilization and exchange of ideas (such as what are the commonalities and differences in the methods, which of the methods from the ontology community can be successfully adopted in the context community, and vice versa, and what working definitions of terms enhance research progress). For example, one perspective was that ontology can be viewed as an explicit encoding of a domain model that may be shared and reused. Another perspective is that a context can be viewed as an explicit encoding of a domain model that is expected to be local and may contain one party's subjective view of the domain.

Some technical themes discussed in the workshop include (1) approaches to the semantic heterogeneity problem using combinations of multiple contexts and ontologies; (2) technical problems related to integration of contexts and ontologies from theoretical, practical, and application perspectives.

The workshop consisted of two invited talks, four technical sessions, two poster sessions, and a discussion and wrap-up session. We received 30 submissions: 11 were selected for technical sessions and 16 were selected for poster sessions.

In the first invited talk, Fausto Giunchiglia from the University of Trento, discussed how ontologies can be contextualized, thereby yielding contextual ontologies, which have the advantages of both ontologies and contexts. In the second invited talk, Chris Welty of IBM discussed why ontologies need contexts and why contexts need ontologies. He also considered some outbriefing from the Advanced Research and Development Activity (ARDA) Interoperable Knowledge Representation for Intelligence Support (IKRIS) program on contexts and what may be useful to include concerning them in knowledge representation languages for interoperability.

Technical sessions addressed various combinations of contexts and ontologies from theoretical, practical, and

application perspectives. The foundations session covered some bridges between contexts and ontologies in information integration scenarios (such as airfare and e-government). The language and reasoning session concentrated on ambiguities of natural language, encoding natural language into a logical language, and the trade-off between expressivity of logical languages and reasoning. The information retrieval session focused on the issues of semantic annotation, relevance, and scoping of information depending on the application context. The ontology matching session introduced a few new approaches to the semantic heterogeneity problem using the match operation.

The poster and discussion and wrap-up sessions generated many fruitful discussions on the workshop themes. In particular, participants agreed that the main themes of convergence among contexts and ontologies are information interoperability and reuse. They also agreed that the workshop was productive and demonstrated a robust interest in a contexts and ontologies workshop next year.

The workshop papers were published as an AAAI technical report and posted in AAAI's digital library.

Educational Data Mining

Joseph E. Beck

The field of educational data mining focuses on improving our knowledge of learning and teaching by extracting patterns from the data collected as part of the educational process. Computers, especially computer tutors, enable data collection over long periods of time, for many students, and at a fine time scale. These advantages provide a novel source of data for understanding how students learn. Although there have been similar workshops, those workshops were limited to specialized conferences. Therefore, one goal of this workshop was to bring together a broader group of AI researchers.

Two big ideas cut across several workshop talks: Educational data can be very fine-grained, and we need some means of "stepping back" to view an aggregation of the data. Al-

though computers enable collection of keystroke level data, this level is probably not the best one for classifying students or examining the data. One proposal was to encode special-case detectors, such as an expert in the domain being able to realize that a particular characteristic must necessarily be exhibited, while a novice must perform several tests to confirm its existence. Preprocessing log files with such detectors enables researchers to better classify students and understand how they are learning. An alternate approach was to provide a generic tool for browsing and summarizing student interactions with computer tutors. This tool allows researchers to select the level of detail they want to see and avoids the problem of being unable to "see the forest for the trees."

Using students' performance data (how they solve problems) to construct a model of the domain produces a very different result than asking domain experts to construct such models. This work focused on examining student performance data, finding questions with correlated performance, and then extracting factors that describe the domain. Constructing domain models in this way seems to result in many fewer factors to describe a domain compared to expert beliefs. While hand-crafted domain models are a useful theoretical description, students do not seem to perform or recognize differences at such a subtle level.

Although this workshop was the first one at AAAI, the quality of papers was strong, and we look forward to having a second such workshop at AAAI 2006 in Boston.

The workshop papers were published as an AAAI technical report and posted in AAAI's digital library.

Exploring Planning and Scheduling for Web Services, Grid, and Autonomic Computing

Biplav Srivastava and Jim Blythe

Planning and scheduling, long active areas of research, have recently received increasing attention for their role in managing workflows on the

Web and on computational grids, encompassing workflow generation, storage and retrieval, analysis, composition, allocation of resources, execution, and repair. The appropriate use of these technologies will have enormous significance for scientific applications and business process integration. Two successful workshops were held at ICAPS 2003 and 2004 on planning and scheduling for Web and grid services. This year, we recognized the need to bring the discussions to a wider AI audience and also extended the scope to self-management of resources (that is, autonomic computing).

The workshop was held in three sessions corresponding to the three application areas—web services, grid, and autonomic computing. In the web services session, Marco Pistore from the University of Trento gave an invited talk on automatic composition of executable web services with a case study on its use for integrating complex business processes. We also had papers on qualitatively evaluating Web services composition approaches, learning to induce source descriptions and matching of services, and interesting ideas on composition and enactment approaches. In the ensuing discussions, it emerged that a key challenge in using planning and scheduling for web services is the formulation of the realistic problem and this often leads to interesting new techniques for advancing the field.

Presentations in the grid session described a formal framework to model the scheduling problem for grids, the use of techniques developed for distributed agents, and an architecture for developing distributed data mining applications. As in the previous session, authors focused more on managing realistic systems and exposing their underlying assumptions, and less on specific planning or scheduling algorithms. The papers also shared the idea that the tasks of workflow allocation and repair are likely to be distributed across different hosts rather than centralized. Both these trends mark interesting departures from the previous workshops held at ICAPS.

Management of resources, including software and hardware that may

be located centrally or distributed across a network, has emerged as the biggest challenge in reducing the cost of information technology while increasing business productivity. Self-management of systems for configuration, protection, recovery, and optimization comes under the ambit of autonomic computing and similar initiatives in industry. Here, workflows have been adopted as the underlying representation to connect interrelated tasks required for self-management. Gerry Tesauro from IBM Research gave an invited talk on autonomic computing in which he presented an overview of the system self-management problem. He related that his group has found that learning-based (nonmodel) approaches can have comparable performance to model-based approaches for resource allocation in distributed systems, enabling good performance without requiring in-depth system model information.

Overall, the workshop was successful in generating thoughtful discussions and building a broad understanding of the challenges ahead.

The workshop papers were published as an AAAI technical report and posted in AAAI's digital library.

Human Comprehensible Machine Learning

Dan Oblinger

Humans need to trust that intelligent systems are behaving correctly, and one way to achieve such trust is to enable people to understand the inputs, outputs, and algorithms used, as well as any new knowledge acquired through learning. As the use of machine learning increases in critical operations, it is being applied increasingly in domains where the learning system's inputs and outputs must be understood or even modified by human operators. For instance, e-mail classification systems may need to gain the user's trust by explaining their predictions in a language the user can understand. Intelligent office assistants learn from a user's preferences and behavior, but if an agent is to be useful, its users must trust that it will make the same decisions that they

would under the same conditions. Machine learning has also been widely used to support credit approval decisions, yet banks are becoming increasingly responsible for explaining the reasons behind a denial of credit. Autonomic systems are beginning to employ machine learning to support common administrative policies, yet system administrators are reluctant to trust automated technology they do not understand.

The workshop's ten presentations addressed comprehensibility from three perspectives.

First, several of the more senior learning researchers offered their personal overview of the last 20 years of learning research on comprehensibility. Pat Langley and Michael Pazzani described a range of their work and provided good overviews of how it related to other's work in comprehensibility.

Second, several researchers presented frameworks that improved comprehensibility for known learning algorithms in some way. For example, John Burge provided a pair of metrics for dynamic Bayes nets and showed how each provided learned nets with a specific meaning for the user. Flavian Vasile provided an extension to RIPPER that produces more comprehensible rules. Pedro Domingos provided a Markov logic network as a basis for comprehensible machine learning.

A third type of presentation focused on comprehensibility in specific domains or for specific applications. For example, Rich Caruana looked at how understanding or misunderstanding of induced models in the medical field has its own particular nuances.

Some attendees and presenters were from allied fields. For example, Simone Stumpf looked at comprehensibility from a human-computer interaction (HCI) perspective, and Noboru Matsuda discussed learning as an authoring tool for cognitive tutoring systems.

A panel of five speakers guided a lively discussion on open research issues in comprehensibility. Topics for the discussion were solicited before and during the workshop. There was a general sense of surprise among the participants that the issues being

raised were of broad interest to the community.

The degree of interaction and interest from both audience and presenters suggests that this is a fertile area for follow-up forums.

The workshop papers were published as an AAAI technical report and posted in AAAI's digital library.

Inference for Textual Question Answering

Sanda M. Harabagiu

The workshop on inference for textual question answering provided a forum for researchers involved in studying various forms of reasoning used in question-answering systems. When reasoning about a question and a candidate answer, an inference engine uses two forms of knowledge: (1) knowledge derived only from the question and the candidate answer, and (2) world knowledge extracted from various ontologies and knowledge databases. Several examples of inference methods used for solving complex questions were presented: abductive reasoning, default reasoning, and inference based on epistemic logic or description logic. Language interpretation also requires its own forms of inference, such as conversational implicatures, processing of metonymies, and metaphors. Inferring the answer to a question is often constrained by temporal and spatial reasoning.

The challenge issued by this workshop was to find a suitable knowledge representation and a robust inference mechanism that handles a majority of the ambiguities generated by natural texts. This problem is of interest for both the natural language processing (NLP) community and the knowledge representation and reasoning (KRR) community. The workshop was an excellent opportunity for researchers from both areas to meet and discuss this problem. The talks highlighted inference mechanisms based on very different knowledge representations and operating on different text collections. The workshop helped foster solutions to problems in several intelligent question and answer systems that can now justify their extracted answers. The so-

lutions discussed will further improve such systems and enable them to tackle more complex questions.

The workshop papers were published as an AAAI technical report and posted in AAAI's digital library.

Integrating Planning into Scheduling

*Mark Boddy, Amedeo Cesta,
and Stephen Smith*

The topic of the workshop was how to integrate planning capabilities with scheduling algorithms and frameworks. In recent years, the AI planning community has focused increasingly on extending classical planning formalisms to incorporate notions of resources and time. Recently published work and the results achieved in the most recent International Planning Competition (IPC) at ICAPS 2004, have demonstrated considerable progress on incorporating metric quantities and durative actions into the classical planning framework, increasing the relevance of classical planning techniques to scheduling problems. However, because the focus in classical planning is on individual actions, rather than organizing or synchronizing with operations in the larger environment, and on discrete state changes, rather than multiple, interacting asynchronous processes, augmenting planning systems with models that include durative actions and resource capacity constraints is unlikely by itself to be an effective solution for problems in which resource allocation is central. On the other hand, the techniques schedulers typically use to solve embedded planning problems tend to be problem-specific and are difficult to extend and transfer to new contexts.

The 2005 workshop was a follow-up to the workshop on the same topic held at ICAPS 2004. At the previous workshop, the differences in perspective on this integration issue were evident from the outset.

Planning-centric papers received good reviews from people with strong roots in the planning research community and were deemed irrelevant by scheduling researchers. Papers with a very strong scheduling focus received

the same treatment, with the roles for each community reversed. However, the several papers that did address the integration we are seeking were universally seen as relevant, and discussion at the 2004 workshop itself led to a more integrated view. At the workshop's conclusion, it was commonly acknowledged that there is a wide variety of possible integrations of planning and scheduling, in some cases involving adding limited forms of resource reasoning to planning.

At the 2005 workshop, several papers advanced the themes identified and discussed at the 2004 workshop, while others highlighted an area that received little attention last year: integrated planning and scheduling for distributed systems. The summary panel and discussion concluded that collectively the two workshops have yielded a sufficient understanding of both kinds of integration strategies, and the domain characteristics for which those strategies are likely to be useful, to provide a general overview and a preliminary synthesis. The current intent is to write up these insights for publication. We also reached a general consensus that a more appropriate term to describe work in this area would be "integrating planning and scheduling."

The workshop papers were published as an AAAI technical report and posted in AAAI's digital library.

Learning in Computer Vision

Bir Bhanu

Computer vision was one of the first areas that the AI community worked on, and now the two communities (learning and computer vision) are marching along their own paths with little interaction. The objective of the first one-day workshop on learning in computer vision was to bring the two communities together to address interdisciplinary research issues. Such an interaction would help increase the competence of AI vision systems to be used in complex real-world applications.

The goal of computer vision research is to provide computers with human-

like perception capabilities so that they can sense the environment, understand the sensed data, take appropriate actions, and learn from this experience to enhance future performance. The computer vision field has evolved from the application of classical pattern recognition and image-processing techniques to advanced applications of image understanding, model-based vision, knowledge-based vision, and systems that exhibit learning capability. The ability to reason and the ability to learn are the two major capabilities associated with these systems. In recent years theoretical and practical advances have been made in the field of computer vision by new techniques and processes of learning, representation, and adaptation. Learning represents the next challenging frontier for computer vision.

During the workshop there was a discussion of the development of flexible, robust AI-based computer vision systems for real-world dynamic scene understanding. There were talks on learning in computer vision, multi-robot interaction, sensor planning, incremental 3D modeling, semantic feature extraction from video, recognizing activities in video, sensorimotor learning, visually guided control, and statistical learning.

The workshop papers were published as an AAAI technical report and posted in AAAI's digital library.

Link Analysis

Dunja Mladenic, Natasha Milic-Frayling, and Marko Grobelink

Link analysis has been developed over the past 20 years in various fields, including discrete mathematics (graph theory), social sciences (social network analysis), and computer science (graph as a data structure). Recently this area has attracted wider attention for its applicability in areas such as law enforcement investigations (terrorism and money laundering), fraud detection (insurance and banking), web analysis (search engines and marketing), and telecommunications (routers, traffic, and connectivity). Particularly interesting are problems and issues that fall within the intersection of link analysis and fields such as Web and text min-

ing, relational data mining, and, more generally, data mining. Typical examples are in the areas of trend analysis, community identification, web user profiling, media clipping, and marketing, where link analysis complements other research fields and derives additional value from information processing. Another interesting scenario is the extraction of information from unstructured data, representation of the extracted data in graphical form, and further analysis of the resulting graph structure to derive and discover new knowledge.

This workshop followed a series of text mining and link analysis workshops that we have organized over the last 15 years at main international conferences, including the International Conference on Machine Learning (ICML), the Knowledge Discovery and Data Mining Conference (KDD), the IEEE International Conference on Data Mining (ICDM), and the International Joint Conference on Artificial Intelligence (IJCAI) (see <http://kt.ijs.si/dunja/TextWebJSI/>).

Presentations at this workshop covered a range of topics, attesting to the richness and versatility of this research area. Methods for manipulation of graphs were covered in several papers, including research on pattern matching efficiency of semantic graphs using higher-order constructs and extraction of relevant semantic subgraphs to facilitate learning using Bayesian networks. Data analysis workflows generally include a variety of tools to solve a problem. Graph analysis is just one of the components typically used. As expected, different application areas impose different types of workflow for analysts, and thus the tools need to be flexible. This was addressed by specifying a representation language to enable exchange of patterns, hypotheses, and evidence among analysis tools.

The workshop included several papers that show cross-application benefits. For example, one applied the results of social networks research to the problem of ranking autonomous systems on the Internet. Another addressed the management of servers based on the power law relationships observed in the network topologies themselves.

The discovery of link structure and exploitation of graph properties is becoming a common trend in information retrieval. A paper on topic-specific scoring of documents combined the standard methods with link properties of the topics structure to enhance document retrieval. Similarly, a paper on summarization of broadcast news video exploited the link analysis of named entities. Because these and related techniques can be boosted by the availability and quality of a domain ontology, the issue of automatic extraction and structuring of domain-specific terms is of great importance. The workshop included a paper on a graph-based ranking algorithm that identifies domain keywords and exploits the dependencies among terms to structure them as concepts and attributes. Applications often challenge the standard research methods; for example, the application of relational graphs analysis to the problem of detecting tax fraud illustrates a prototype that has been deployed in practice.

The workshop papers were published as an AAAI technical report and posted in AAAI's digital library.

Mobile Robot Workshop

Sheila Tejada and Paul E. Rybski

The mobile robot workshop was an extension of the AAAI 2004 robot program centered around the theme "Robots Interacting with Humans." Participants from the robot competition and exhibition presented highlights of their work at the workshop. The robot program included three competitions—the robot challenge, a scavenger hunt, and an open interaction task—as well as a general robot exhibition.

The robot challenge task was to develop a robot that can attend the conference. This task included a number of subtasks: finding the registration desk from the entrance to the conference center, registering for the conference, performing volunteer duties as required. LABORIOUS, from the Université de Sherbrooke led by François Michaud, earned first place for their work on Spartacus. Spartacus exhibited an impressive variety of algorithms in a motivated behavior architecture

to perform the navigation and processing of visual and auditory inputs required for high-level interaction.

Participants in the scavenger hunt were given a list of objects that would be in specified locations at specified times. The task required robots to map and navigate a dynamic area with moving obstructions, such as people, to acquire the desired objects. This task was designed to use some of the same capabilities as urban search-and-rescue robots. The panel of judges awarded first place to the Harvey Mudd College robot, HMC Hammer. The team of undergraduates used an Evolution ER1 robot with relatively inexpensive equipment in a robot-as-computer-peripheral design that successfully found objects in a very challenging environment.

The open interaction event encouraged researchers to demonstrate techniques for interactive, entertainment, and social robotics. Participants in the competition were encouraged to design robots that grabbed and sustained attendees' attention. Hanson Robotics won the competition for their human emulation robot, an android that depicted the late science fiction writer Phillip K. Dick. The Hanson Robotics team showed that a robot with a very human face could interact well with people and not be viewed as disturbing.

The robot exhibition was designed to showcase current research that does not strictly fit any of the competition tasks. The following teams earned technical achievement awards: LABORIUS for work on map building and human-robot interaction; Harvey Mudd College for overall excellence in a fully autonomous system; University of Massachusetts Lowell for control interface usability and robust path finding and object recognition; Carnegie Mellon University Claytronics for a visionary hardware concept; a Naval Research Laboratory collaboration with University of Missouri for engaging interaction using a cognitive model and innovative interface; the Swarthmore College Academic Autonomy team for adaptive vision for lighting conditions; and the Carnegie Mellon University Tekkotsu project for visualization for education robots. The

following teams earned honorable mention: the Carnegie Mellon University Pink Team, the Drexel University Autonomous Systems Lab, Kansas State University, the Stony Brook Robot Design Team, the Carnegie Mellon University CMDash'05, and the University of Pittsburgh.

The workshop papers were published as an AAAI technical report and posted in AAAI's digital library.

Modular Construction of Humanlike Intelligence

Kristinn R. Thórisson, Hannes Vilhjálmsson, and Stacy Marsella

From the birth of AI, one of the central challenges of the field has been to understand and model human intelligence. The primary motivation for this workshop was the belief that meeting that challenge requires not only studying many separate skills but also integrating them into a coherent whole. In particular, the development of machines that collaborate and interact socially with people necessitates integration of numerous complex technologies. This integration in turn requires better tools and an increased focus on architecture.

Progress toward these goals is best ensured by a healthy balance between theory, tools, and implementation. Indeed, the papers presented at the workshop fell roughly equally into these three categories. Some tools presented took the form of specifications and software libraries. Examples are OpenAIR, developed by MINDMAKERS.ORG, and NetP by Kai-Yuh Hsiao, Peter Gorniak, and Deb Roy. Both are geared toward building complex software systems and proposing methods for easier connection of programming languages and platforms.

Among other tools were the whiteboards by Kristinn Thorisson, Thor List, Christopher Pennock, and John DiPirro, which provide semantic publishing, subscribing, and routing of messages and streams. Thor List and his collaborators showed how modularity can be brought to the construction of vision architectures, while Hannes Vilhjálmsson and Stacy Marsella proposed common representations for enabling easy construction

of multimodal dialogue planners. Among the theoretical papers presented was a proposal by Alexei V. Samsonovich and Kenneth A. De Jong to use neural networks on top of symbolic representations to learn groupings and allow the system to evolve over time. Another was Zippora Arzi-Gonczarowski's ISAAC, with a special module construct that allows incremental, compositional system growth.

Many of the implemented systems presented were formidable attempts at large-scale integration. For example, the MARCO system by Matt MacMahon, a system for following route instructions, addresses shortcomings of the multicomponent GRACE and GEORGE robots (AAAI Robot Challenge 2002). A robot designed by Maren Bennewitz, Felix Faber, Dominik Joho, Michael Schreiber, and Sven Behnke, can maintain dialogue with two humans at once and the architecture created by Nikolaos Mavridis and Deb Roy enables a robot to interact with a user through speech, updating its beliefs about the surroundings using a range of heterogeneous perceptual inputs. Addressing somewhat different questions, a robot built by Jesse Gray and Cynthia Breazeal runs simulations of people to infer their mental state.

There is significant potential for collaboration in the research presented. The theoretical work needs to be tested in context—some of the robot architectures, such as Mavridis and Roy's, could provide that test bed. Eric Baumer and Bill Tomlinson's work on synthetic social construction could quite possibly benefit from Vilhjálmsson and Marsella's work on social communication or Andrew Gordon's classification of cognitive architectures. Rakesh Gupta and Ken Hennessey's work on commonsense reasoning provides an interesting piece to the puzzle of designing robots to work effectively indoors. All of the work presented, especially the multimodule robot architectures, could benefit from using OpenAIR and NetP, which would simplify integration and ease module reuse.

During the AAAI conference our workshop got a nice boost "from above": Both Marvin Minsky's

**For information on 2006
AAAI Workshops, please
visit:**

**[www.aaai.org/
Workshops/2006](http://www.aaai.org/Workshops/2006)**

keynote address and Ronald J. Brachman's presidential address emphasized integration and large-scale modeling and urged the field to set its sights on human intelligence. Practically all of the researchers in our workshop do so. We find this very exciting and hope they make real progress fast.

The workshop papers were published as an AAAI technical report and posted in AAAI's digital library.

Multiagent Learning

Eduardo Alonso

When designing agent systems, it is impossible to foresee all the potential situations an agent may encounter and specify an optimal agent behavior in advance. Agents therefore have to learn from and adapt to their environment. This task is even more complex when nature is not the only source of uncertainty, and the agent is situated in an environment that contains other agents with potentially different capabilities, goals, and beliefs. Multiagent learning, that is, the ability of agents to learn how to cooperate and compete, becomes crucial in such domains.

The workshop on multiagent learning held July 10 as part of the AAAI 2005 Conference in Pittsburgh achieved all its goals: It increased awareness and interest in adaptive agent research, encouraged collaboration between machine learning and agent system experts, and gave a representative overview of current research in the area of adaptive agents.

The workshop was an inclusive forum for discussion about ongoing or completed work in both theoretical and practical issues. In particular, work was presented about learning to communicate and to get organized in net-

works following game theory approaches and others. Techniques varied from reinforcement learning to evolutionary algorithms. Presentations included an invited talk entitled "Learning for Multiagent Decision Problems" by Geoff Gordon from the Center for Automated Learning and Discovery at Carnegie Mellon.

The workshop papers were published as an AAAI technical report and posted in AAAI's digital library.

Question Answering in Restricted Domains

Diego Molla Aliod and Jose Luis Vicedo

Early question-answering systems of the 1960s and 1970s were based on handcrafted databases and knowledge systems of very specific domains, such as the analysis of lunar rock samples from the Apollo missions or of U.S. baseball league statistics. In those days the choice to use restricted domains was made out of necessity. Now theoretical advances in AI and natural language processing (NLP), together with increases in computing power and the availability of corpora and linguistic tools and resources, have enabled the development of open-domain question-answering systems. However, restricted domains provide an interesting mix of challenges and opportunities that have yet to be fully explored. For example, the limited data available makes it difficult to apply current question-answering techniques based on data redundancy, and generic lexical resources do not cover the specific terminology of some domains. On the other hand, some restricted domains (such as the clinical domain) have high-quality, comprehensive ontologies and resources that can be used for high-precision question-answering tasks. Following the 2004 Association of Computational Linguistics (ACL) workshop on question answering in restricted domains, this one-day AAAI workshop explored some of the issues involved in this increasingly active area of research and development.

The workshop papers were published as an AAAI technical report and posted in AAAI's digital library.

Spoken Language Understanding

*Gokhan Tur, Dilek Hakkani-Tür, and
Srinivas Bangalore*

Natural language processing (NLP) has been one of the defining subtopics of AI since its early days. In recent times, NLP has predominantly been about text understanding and building associated resources for the purposes of information extraction, question answering, and text mining. Many of these tasks have nourished the creation and development of extensive ontologies, practical semantic representations, and novel machine-learning techniques.

In a spirit similar to the workshop at the Human Language Technology Conference / North American chapter of the Association for Computational Linguistics (HLT-NAACL) 2004 on this topic, our attempt was to broaden the scope of language understanding to include spoken language understanding (SLU) in the context of applications such as speech mining and human-machine interactive spoken dialogue systems. We tried to bring together techniques that address the issue of robustness of SLU to speech recognition errors, language variability, and dysfluencies in speech with issues of semantic representation that provide greater portability to a dialogue model. We believe spoken language understanding is an especially attractive topic for cross-fertilization of ideas between the AI, information retrieval (IR), NLP, speech, and semantic Web communities. We thank Dr. Alexander I. Rudnicky, from the School of Computer Science at Carnegie Mellon University for accepting our invitation to share his group's work on spoken language understanding.

We thank the program committee members for their informative reviews of the submitted papers. We thank the authors for electing to present their work at this forum. And we thank the AAAI for inviting us to organize this follow-up workshop.

The workshop papers were published as an AAAI technical report and posted in AAAI's digital library.