

# AI magazine

Volume 26, Number 4

Winter 2005



 An official publication of the  
American Association for Artificial Intelligence

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