



Invited Talks and Panels

2002 AAAI Presidential Address:

AI and the Impending Revolution in Brain Science

Tom M. Mitchell, Carnegie Mellon University

AAAI-02 Keynote Address:

Probabilistic AI and Information Retrieval

Michael I. Jordan, University of California, Berkeley

Much progress has been made in recent years in the area of information retrieval, in particular as embodied in Internet search engine technology. Much progress has also been made in probabilistic, graph-theoretic AI. What are the possibilities for bringing these two lines of research together —

for viewing large-scale information retrieval as a core enabling technology for AI systems, and for asking IR systems to exhibit true inferential capabilities? Jordan will discuss research aimed at bridging the AI/IR gap.

AAAI-02/IAAI-02 Joint Invited Talk:

Human Level "Strong" AI: The Prospects and Implications

Raymond Kurzweil, KurzweilAI.net (Kurzweil Accelerating Intelligence Network)

Three-dimensional molecular computing will provide the hardware for human-level "strong" AI well before 2030. The more important software insights will be gained in part from the reverse-engineering of the human brain, a process well under way. Once nonbiological intelligence matches the range and subtlety of human intelligence, it will neces-

sarily soar past it because of the continuing acceleration of information-based technologies, as well as the ability of machines to instantly share their knowledge. The implication will be an intimate merger between the technology-creating species and the evolutionary process it spawned.

AAAI-02 Invited Talk:

Perspectives on Artificial Intelligence Planning

Hector Geffner, ICREA - Universitat Pompeu Fabra (Barcelona)

Planning has always been a key area in artificial intelligence. In its general form, planning is concerned with the automatic synthesis of action strategies (plans) from a description of actions, sensors, and goals. Planning thus contrasts with two other approaches to intelligent behavior: the programming approach, where action strategies are defined by hand, and the learning approach, where action strategies are inferred from experience.

Different assumptions about the nature of actions, sensors, and costs lead to various forms of planning: (1) Planning with complete information and deterministic actions, (2) planning with non-deterministic actions and sensing, and (3) planning with temporal and concurrent actions, etc. Most

work so far has been devoted to "classical" planning (1. above), where significant changes have taken place in the last few years. On the methodological side, the area has become more empirical with experimental evaluation being routine; on the technical side, approaches based on heuristic or constrained-based search have taken over blind-search approaches.

In this talk, Geffner will provide a coherent picture of planning in AI while trying to convey some of the current excitement in the field. He'll make emphasis on the mathematical models that underlie various forms of planning, and the ideas that have been found most useful computationally.

AAAI-02 Invited Talk:

Dimension Reduction that Preserves Information and Neural Coding

Naftali Tishby, The Hebrew University

Many cognitive functions, such as prediction, feature extraction, noise filtering, and learning, can be viewed as special cases of one principle: compression while preserving information. This information theoretic principle was turned into a computational paradigm: the information bottleneck method. This variational method yielded several novel learning and data analysis algorithms, with many applications to information retrieval as well as to analysis of neural coding in several neurobiological systems, that were carried in Tishby's lab. In this talk Tishby will focus on a new approach to data dimensionality reduction that stems from this

principle. Here he searches for low dimensional (nonlinear) reduction of co-occurrence (or contingency) tables that preserve the (mutual) information in the table. He gives a new alternate-projection algorithm for achieving such a reduction and show its convergence to an optimal set of information preserving features. This approach is particularly useful when the data is not naturally quantized but rather represented by low dimension continuous features. Such a reduction may have interesting biological implications. (Based on joint work with Amir Globerson and Noam Slonim.)

IAAI-02 Invited Talk:

Robot-Assisted Urban Search and Rescue at the WTC: Where's the AI?

Robin R. Murphy, University of South Florida

On September 11, 2001, the Center for Robot-Assisted Search and Rescue responded within six hours to the WTC disaster; this is the first known use of robots for USAR. The University of South Florida was one of the four robot teams, and only academic institution. The USF team participated on-site in the search efforts from September 12 through 22, collecting and archiving data on the use of robots.

This talk will provide an overview of the use of robots for USAR as well as discuss what AI techniques were available, what was actually used, and why. It will also summarize the key lessons learned from the robotics efforts at the WTC.

The lessons learned cover the areas of platforms and mobility, sensors and sensing strategies, control, and human-robot interactions. Possibly the most pervasive lesson learned is that robots for USAR must be considered from an "information technology" perspective, where platforms, sensors, control schemes, networks, and interfaces must all be coevolved to ensure the information extracted by the robots is truly usable by the rescue community.

Extensive video footage of the site and "robot's eye" views will be shown.

IAAI-02 Invited Panel:

Pioneering AI Businesses I: A 20-Year Review

Panel Leader: Neil Jacobstein, Teknowledge Corporation

Several AI-based businesses started in the early 1980s. They underwent a classic boom and bust cycle. Hype exceeded expectations, and some investors and technologists lost patience. However, history shows that in cases of disruptive technological innovation, forecasts are usually too optimistic in the short run, and too conservative in the long run. Is that the case with AI businesses? This panel of AI entrepre-

neurs will review the technology base and history of pioneering AI businesses, extract lessons learned, and identify future opportunities. Companies discussed will include IntelliCorp, Teknowledge, Inference, Syntelligence, Carnegie Group, Cycorp, and others. An interactive question and answer session with panel members will follow brief presentations from each panelist.

IAAI-02 Invited Panel:

Pioneering AI Businesses II: Recent Startups

Panel Leader: Craig Knoblock, University of Southern California and Fetch Technologies

This panel will focus on the process of starting an AI company. The barriers to starting a new company include applying the technology to address a specific market need and creating and running a successful business. The speakers on this panel are AI researchers that have recently started their own companies. Some of the issues to be discussed by the

panelists include how to go from a technology to a business, how to get funding for a company, and what are some of the pitfalls to watch out for. An interactive question and answer session with panel members will follow brief presentations from each panelist.

IAAI-02 Invited Panel (Collocated with AI Festival):

PI Entrepreneurs Forum

This forum will provide an open and informal setting for AI pioneers, technologists, entrepreneurs, venture capitalists,

legal, and intellectual property experts to network and discuss issues in starting and running AI-based companies.