

# Invited Talks

*AAAI-06 Keynote Address*

## **AI and Semantic Web**

Tim Berners-Lee, *World Wide Web Consortium*

The relationship between AI and the semantic web has been something that has provoked a lot of heated corridor discussion over the years. The talk will try to outline what the semantic web is and is not, at a conference where there may be some anniversary reflection on what AI is and is not. It is not always obvious how to transfer existing AI techniques into a fractal weblike space, or what the effect will be. But it is certainly exciting.

A graduate of Oxford University, England, Berners-Lee now holds the 3Com Founders chair at the Laboratory for Computer Science and Artificial Intelligence Lab (CSAIL) at the Massachusetts Institute of Technology (MIT), where he leads the Decentralized Information group (DIG). He directs the World Wide Web Consortium, an open forum of companies and organizations with the mission to lead the Web to its full potential.

With a background of system design in real-time communications and text processing software development, in 1989 he invented the World Wide Web, an internet-based hypermedia initiative for global information sharing, while working at CERN, the European Particle Physics Laboratory. He wrote the first web client (browser-editor) and server in 1990.

*IAAI / AI Magazine Robert S. Engelmore Memorial Lecture*

## **What Do We Know About Knowledge?**

Bruce G. Buchanan, *University of Pittsburgh*

Intelligent systems need knowledge. However, the simple equation “knowledge is power” leaves three major questions unanswered. First what do we mean by “knowledge,” second, what do we mean by “power,” and third, what do we mean by “is?”

In this talk Buchanan will examine these questions. In particular he will focus on some of the milestones in understanding the nature of knowledge and some of what we have learned from fifty years of AI. The discipline and detail required to write programs that use knowledge have given us some valuable lessons for implementing the knowledge principle.

But there are still interesting challenges ahead.

Bruce Buchanan received a B.A. in mathematics from Ohio Wesleyan University (1961), and M.S. and Ph.D. degrees in philosophy from Michigan State University (1966). He was on the faculty at Stanford University and the University of Pittsburgh, with appointments in computer science, philosophy, medicine, and intelligent systems. He is a Fellow of the American Association for Artificial Intelligence, a Fellow of the American College of Medical Informatics, and a member of the National Academy of Sciences Institute of Medicine. He has served on the editorial boards of several journals and has served as secretary-treasurer and president of the American Association of Artificial Intelligence.

*AAAI-06 Invited Talk*

## **Unifying Logical and Statistical AI**

Pedro Domingos, *University of Washington*

Intelligent agents must be able to handle the complexity and uncertainty of the real world. Logical AI has focused mainly on the former, and statistical AI on the latter. Markov logic combines the two by attaching weights to first-order formulas and viewing them as templates for features of Markov networks. Inference algorithms for Markov logic draw on ideas from satisfiability, Markov chain Monte Carlo and knowledge-based model construction. Learning algorithms are based on the voted perceptron and inductive logic programming. Markov logic has been successfully applied to problems in entity resolution, link prediction, information extraction and others, and is the basis of the open-source Alchemy system.

Pedro Domingos is an associate professor of computer science and engineering at the University of Washington. His research interests are in AI, machine learning, and data mining. He received a Ph.D. in information and computer science from the University of California at Irvine, and is the author or coauthor of over 100 technical publications. He was program cochair of KDD-2003, and has served on numerous editorial boards and program committees. He has received several awards, including a Sloan fellowship, an NSF CAREER award, a Fulbright scholarship, an IBM faculty award, and best paper awards at KDD-98, KDD-99, and PKDD-2005.

AAAI-06 Invited Talk

## **Cognitive Tutors and Opportunities for Convergence of Human and Machine Learning Theory**

Ken Koedinger, *Carnegie Mellon University*

Cognitive tutors are computer-based intelligent tutors that are based on cognitive psychology theory and employ artificial intelligence methods. Cognitive tutors have led to significant improvements in student learning and are now in regular use in mathematics classrooms in over 2,000 schools across the country. They have also been an important basic research vehicle to test and advance both psychology and AI theory. Besides illustrating this important application of AI, I will also discuss how we have diverged from a time when theoretical explorations of machine and human learning were more closely aligned and how it is now time reconvene and to fruitfully compare and contrast the independent theoretical advances both fields have made.

Ken Koedinger specializes in human-computer interaction and psychology at Carnegie Mellon University. His research focuses on human problem solving and learning and how better cognitive models of human learning can be used in technology-enhanced courses to demonstrably increase student achievement. He is a cofounder of Carnegie Learning ([CarnegieLearning.com](http://CarnegieLearning.com)) and is the CMU director of the Pittsburgh Science of Learning Center ([LearnLab.org](http://LearnLab.org)). PSLC is a \$25 million NSF center that provides researchers with "LearnLab," an international resource for creating, running, and mining data from rigorous classroom experiments to apply and advance theories of human and machine learning.

AAAI-06 Invited Talk

## **Developing an Intelligent Personal Assistant: The CALO Project**

Karen Myers, *SRI International*

Knowledge workers today must juggle a range of tasks and responsibilities while maintaining awareness of deadlines, resources and events that could impact objectives. I will describe a collaboration by a team of 25 organizations to produce an intelligent personal assistant for improving the productivity of such workers. This assistant, called CALO (cognitive assistant that learns and organizes), both performs tasks on its user's behalf and assists with the filter-

ing, organization, and preparation of information. CALO is seeded with default problem-solving knowledge but learns to expand and improve its capabilities over time by observing and interacting with its user.

Karen Myers is a principal scientist in the AI Center at SRI International, where she directs a research group focused on planning and agent technologies. Her current interests lie with the development of mixed-initiative AI systems that enable humans and machines to solve problems cooperatively. Myers is on the editorial boards of *JAIR* and *Artificial Intelligence*, and is a member of the AAAI executive council. She earned her Ph.D. in computer science from Stanford University.

AAAI-06 Invited Talk

## **Global Inference and Learning: Towards Natural Language Understanding**

Dan Roth, *University of Illinois at Urbana-Champaign*

The maturity of machine learning techniques allows us today to learn many low level natural language predicates and generate an appropriate vocabulary over which reasoning methods can be used to make significant progress in natural language understanding. I will describe research on a framework that combines learning and inference. Our inference with classifiers approach allows the output of local classifiers for different problem components to be assembled into a whole that reflects global preferences and constraints. Examples will be drawn from "wh" attribution in natural language processing (determining who did what to whom when and where) and textual entailment (determining whether one utterance is a likely consequence of another).

Dan Roth is a professor in the Department of Computer Science at the University of Illinois at Urbana-Champaign. Roth has published broadly in machine learning, natural language processing, and knowledge representation and reasoning and has developed advanced machine learning based tools for natural language applications that are being used widely by the research community. Among his paper awards are the best paper award in IJCAI-99 and the 2001 AAAI Innovative Applications of AI Award. Roth was the program chair of CoNLL'02 and of ACL'03 and is an associate editor for *JAIR* and the *Machine Learning Journal*. Roth earned his Ph.D in computer science from Harvard University in 1995.

## **Electrifying Knowledge Work: 362 Innovative Applications of Artificial Intelligence 1989-2006**

Neil Jacobstein, *Teknowledge Corporation*

Early stage artificial intelligence has already produced a wide range of valuable applications in industry and government. Many of these applications have performed complex tasks such as planning, monitoring, design, risk assessment, diagnosis, training, process control, classification, and analysis. AAAI's Innovative Applications of AI Conference has published over 360 successful applications of AI in fields as diverse as biotechnology, space flight, manufacturing, security, paleontology, construction, energy, music, military, intelligence, banking, telecommunications, news media, management, law, emergency services, agriculture, treaty verification, and many others. This talk will review the patterns that connect these applications over 18 years of the IAAI conference: what worked, what didn't, and what were the key trends. None of these systems exhibited general intelligence, but each documented our ability to codify and distribute human problem solving knowledge, and put it to work.

Neil Jacobstein is CEO of Teknowledge Corporation, a 25 year-old AI company. He was the chair of AAAI's 2005 Innovative Applications of AI conference. He has been a technical consultant on projects for NSF, DARPA, NASA, NIH, EPA, DOE, NIST, NRO, GM, Ford, P&G, Boeing, Applied Materials, NPR, NAS, the Army's STRICOM Technical Advisory Board, and others. In 1999, Jacobstein became a Crown Fellow at the Aspen Institute. He received his B.S. summa cum laude from the University of Wisconsin, and an M.S. from the University of Texas. He attended Stanford University's Advanced Management College and Managing Innovation programs. Jacobstein was a graduate research intern and consultant at Xerox PARC.

## **Winning the DARPA Grand Challenge**

Sebastian Thrun, *Stanford University*

The DARPA Grand Challenge was the most significant event in the field of robotics in more than a decade. A mobile ground robot had to traverse 132 miles of punishing desert terrain in less than ten hours. In 2004, the best robot only made 7.3 miles. A year later, Stanford won this historical challenge and cashed the \$2 million prize. This talk, delivered by the leader of the Stanford Racing Team, will provide insights into the software architecture of Stanford's winning robot "Stanley." The robot heavily relied on advanced sensor technology, and advanced artificial intelligence to make sense out of the massive amounts of sensor data acquired by the vehicle. The talk will introduce you into the fascinating world of autonomous robotics, share with you many of the race insights, and discuss with you some of the implications for the future of our society.

Sebastian Thrun is an associate professor of computer science and director of the Stanford Artificial Intelligence Laboratory (SAIL). Prior to winning the Grand Challenge, Thrun published seven books, nearly 300 refereed papers, won numerous best paper and career achievement awards, was named one of the "Brilliant Ten" by *Popular Science*, and served as principal investigator on 6 major DARPA research initiatives. Thrun's research focuses on robotics, machine learning, and artificial intelligence. Stanley, the robot, was recently named "Best Robot Ever" by *Wired Magazine*.