Invited Talks

AAAI Presidential Address

Eric Horvitz (Microsoft Corporation)

Eric Horvitz is a principal researcher and research area manager at Microsoft Research. He has had a lifelong interest in perception, reasoning, and action under uncertainty. He has pursued insights about intelligence via studies of inference and decision making under limited and varying computational resources, including investigations of bounded optimality, value of computation, utility-theoretic metareasoning, and flexible procedures and representations. His current interests span theoretical and practical challenges in machine reasoning and learning, principles of human-computer collaboration, and search and information retrieval. Before his service as AAAI president, he was elected a Fellow and is a past councilor of the organization. He has been chair of the Association for Uncertainty and Artificial Intelligence (AUAI), and has served on the DARPA Information Science and Technology Study Group (ISAT), and the Naval Research Advisory Committee (NRAC). He has been active on numerous editorial boards and program committees and with the organization of multiple conferences and workshops. He received his Ph.D. and M.D. degrees from Stanford University.

Robert S. Engelmore Memorial Award Lecture

Toward Cognitive Prostheses

Kenneth M. Ford (Florida Institute for Human & Machine Cognition [IHMC])

The emerging concept of human-centered (HCC) computing represents a significant shift in thinking about intelligent machines and, indeed, about information technology in general. Human-centered computing embodies a "systems view," in which human thought and action and technological systems are seen as inextricably linked and equally important aspects of analysis, design, and evaluation. From an AI perspective, the HCC framework is focused less on stand-alone exemplars of mechanical cognitive talent, and is concerned more with computational systems designed to amplify human cognitive and perceptual abilities. This approach results in systems that can be regarded as cognitive or perceptual prostheses, much as eyeglasses are a sort of ocular prosthesis. These systems fit the human and machine components together in ways that exploit their respective strengths and mitigate their respective weaknesses. Building cognitive prostheses is fundamentally different from AI's traditional Turing Test ambition — it doesn't set out to imitate human abilities, but to extend them. This shift in perspective places human/machine interaction issues at the center of the subject. The "system" in question isn't "the computer," but instead includes cognitive and social systems, computational tools, and the physical facilities and environment. Thus, human-centered computing provides a new research outlook for AI applications, with new research agendas and goals.

   Kenneth Ford is the founder and director of the Florida Institute for Human & Machine

Cognition (IHMC), an independent not-for-profit research institute. He is also is the author or coauthor of hundreds of scientific papers and six books. Ford's research interests include artificial intelligence, cognitive science, human-centered computing, and entrepreneurship in government and academia. He received a Ph.D. in computer science from Tulane University. He is emeritus editor-in-chief of AAAI/MIT Press, involved in the editing of several journals. Ford is a Fellow of the AAAI.

In January 1997, Kenneth Ford was asked by NASA to develop and direct its new Center of Excellence in Information Technology and to serve as the associate center director at Ames Research Center. In July 1999, Ford was awarded the NASA Outstanding Leadership Medal. That same year, Ford returned to private life and to the IHMC.

In October of 2002, President Bush nominated Kenneth Ford to serve on the National Science Board. In September 2005, Ford received the Doctor Honoris Causas from the University of Bordeaux. Also in 2005, Ford was appointed a member of the Air Force Science Advisory Board. In June 2007, he was appointed to the NASA Advisory Council.

AAAI-08 Invited Talk

From Images to Scenes: Using Lots of Data to Infer Geometric, Photometric and Semantic Scene Properties from a Single Image

Alexei A. Efros (Carnegie Mellon University)

Reasoning about a scene from a photograph is an inherently ambiguous task. This is because a single image in itself does not carry enough information to disambiguate the world that it's depicting. Of course, humans have no problems understanding photographs because of all the prior visual experience they can bring to bear on the task. How can we help computers do the same? Efros proposes to "brute force" the problem by using massive amounts of visual data, both labeled and unlabeled, as a way of capturing the statistics of the natural world.

In this talk, Efros will present some recent results on inferring geometric, photometric, and semantic scene properties from a single image. He will first briefly describe the system for estimating the rough geometric surface layout of a scene as well as the camera viewpoint. He will show how this information, in turn, can be useful for modeling objects in the scene. Next, he will describe a very simple way of using the surface layout information as a way of estimating a rough illumination map for the scene. Finally, he will describe a new system that uses millions of unlabeled photographs from Flickr to capture some implicit semantic scene structure of an image.

Alexei (Alyosha) Efros is an assistant professor at the Robotics Institute and the Computer Science Department at Carnegie Mellon University. His research is in the area of computer vision and computer graphics, especially at the intersection of the two. He is particularly interested in using data-driven techniques to tackle problems, which are very hard to model parametrically but where large quantities of data are readily available. Efros received his Ph.D. in 2003 from the University of California, Berkeley and spent the following year as a fine fellow at Oxford, England. He is a recipient of the NSF CAREER award (2006), the Sloan Fellowship (2008), and the Guggenheim Fellowship (2008).

100 Million Years of Evolutionary History of the Human Genome

David Haussler (University of California, Santa Cruz)

With our ability to sequence entire genomes, we have for the first time the opportunity to compare the genomes of present day species, and deduce the trajectories by which they diversified from a common ancestral genome. Starting with a small shrewlike ancestor in the Cretaceous period approximately 100 million years ago, the different species of placental mammals radiated outward, creating a stunning diversity of forms from whales to armadillos to humans. From the genomes of present-day species, it is possible to computationally reconstruct what most of the DNA bases in the genome of the common ancestor of placental mammals must have looked like. We can then deduce most of the changes that lead to humans. In so doing, we discover how Darwinian evolution has shaped us at the molecular level.

Because most random mutations to functionally important regions of DNA reduce fitness, these changes usually disappear over time in a process known as negative selection. From its unusually high conservation between species, it is immediately evident that at least 5 percent of the human genome has been under negative selection during most of mammalian evolution, and is hence likely to be functionally important. Protein-coding genes and structural RNA genes stand out among the negatively selected regions because of their distinctive pattern of restricted DNA base substitutions, insertions and deletions. However, most of the DNA under negative selection in mammalian genomes, and indeed vertebrate genomes in general, does not appear to be part of protein-coding genes, and shares no sequence similarity with any DNA in the genomes of invertebrates. Experimental evidence suggests that many of these unclassified functional elements serve to regulate genes involved in embryonic development.

Overlaid on the background of negative selection, we occasionally see a short segment of widely conserved DNA that has rapidly changed in a particular lineage, suggesting possible positive selection for a modified function in that lineage. The most dramatic example of this in the last 5 million years of human evolution occurs in a previously unstudied RNA gene expressed in the developing cerebral cortex, known as Human Accelerated Region 1 (HAR1). This gene is turned on only in a select set of neurons, during the time in fetal development when these neurons orchestrate the formation of the substantially larger cortex of the human brain. It will be many years before the biology of such examples is fully understood, but right now we relish the opportunity to get a first peek at the molecular tinkering that transformed our animal ancestors into humans.

David Haussler is an investigator with the Howard Hughes Medical Institute and professor of biomolecular engineering at the University of California, Santa Cruz, where he directs the Center for Biomolecular Science and Engineering. He is also affiliated with the Departments of Computer Science and Molecular, Cell, and Developmental Biology. He serves as scientific codirector for the California Institute for Quantitative Biomedical Research and is a consulting professor at both Stanford Medical School and the University of California, San Francisco Biopharmaceutical Sciences Department. Haussler's research lies at the interface of mathematics, computer science, and molecular biology. He develops new statistical and algorithmic methods to explore the molecular evolution of the human genome, integrating cross-species comparative and high-throughput genomics data to study gene structure, function, and regulation. He has focused on computational analysis and classification of DNA, RNA, and protein sequences. As a collaborator on the international Human Genome Project, his team posted the first publicly available computational assembly of the human genome sequence on the Internet. His

group now maintains a web browser for the genome sequence that is used extensively in biomedical research. Most recently, Haussler has focused on broadly exploring the functional elements of the human genome, primarily through interspecies comparisons; he tests the resulting findings in his wet laboratory. His findings have shed light on the possible functionality of what was once considered to be "junk" DNA. He has also begun to computationally reconstruct the genome of the ancestor common to placental mammals. Haussler received his BA in mathematics from Connecticut College in 1975, an MS in applied mathematics from California Polytechnic State University at San Luis Obispo in 1979, and his Ph.D. in computer science from the University of Colorado at Boulder in 1982. He was recently elected to both the National Academy of Sciences and the American Academy of Arts and Sciences. He is also a fellow of both AAAS and AAAI. He has won a number of prestigious awards, most recently the 2006 Dickson Prize for Science from Carnegie Mellon University.

Sense and Sensibility: Sentiment Analysis, Opinion Mining, and the Computational Treatment of Subjective Language

Lillian Lee (Cornell University)

"What do other people think?" has always been an important consideration to most of us when making decisions. Long before the world wide web, we asked our friends who they were planning to vote for and consulted *Consumer Reports* to decide which dishwasher to buy. But the Internet has (among other things) made it possible to learn about the opinions and experiences of those in the vast pool of people that are neither our personal acquaintances nor well-known professional critics — that is, people we have never heard of. Enter sentiment analysis, a flourishing research area devoted to the computational treatment of subjective and opinion-oriented language. Sample phenomena to contend with range from sarcasm in blog postings to the interpretation of political speeches. Lee's talk will cover some of the motivations, challenges, and approaches in this broad and exciting field.

Lillian Lee is an associate professor of computer science at Cornell University. Her research interests include natural language processing, information retrieval, and machine learning. She is the recipient of the inaugural best paper award at HLT-NAACL 2004 (joint with Regina Barzilay), a citation in "Top Picks: Technology Research Advances of 2004" by Technology Research News (also joint with Regina Barzilay), and an Alfred P. Sloan Research Fellowship, and her group's work has been featured in the *New York Times*.

Making Sense of Complex Networks

Mark Newman (University of Michigan)

There are networks in almost every part of our lives. Some of them are familiar and obvious: the Internet, the power grid, the road network. Others are less obvious but just as important. The patterns of friendships or acquaintances between people form a social network. Boards of directors join together in networks of corporations. Communities of scientists and other aca-

demics join together in networks of collaboration. Recent years have seen an explosion of interest in networks among mathematicians, sociologists, computer scientists, physicists, biologists, and others. Newman's talk will describe some of the successes and challenges of the study of networks and discuss a promising new line of research in the application of methods from machine learning to the analysis and understanding of networked systems.

Mark Newman received his Ph.D. in theoretical physics from the University of Oxford in 1991 and conducted postdoctoral research at Cornell University before moving to the Santa Fe Institute, where he was a resident scientist until 2002, and then to his present position at the University of Michigan. He is currently a professor of physics and complex systems at the University of Michigan as well as being a member of the external faculty of the Santa Fe Institute.

AAAI-08 Invited Talk

What Is To Be Done?

Stuart Russell (University of California, Berkeley)

Much has been achieved in the field of AI, yet much remains to be done if we are to reach the goals imagined by the early pioneers. Russell will examine some of what we have recently understood, as a means of identifying what might be understood next.

Stuart Russell received his B.A. with first-class honors in physics at Oxford University in 1982 and his Ph.D. in computer science at Stanford University in 1986. He then joined the faculty of the University of California at Berkeley, where he is is a professor and chair of computer science and holds the Smith-Zadeh chair in engineering. He is a Fellow and former executive council member of AAAI, a Fellow of ACM, and winner of the NSF Presidential Young Investigator Award, the Computers and Thought Award, and the ACM Karl Karlstrom Outstanding Educator Award. He is the author of over 150 papers and three books.

IAAI-08 Invited Talk

Realizing Claytronics: A Challenge for AI

Seth Copen Goldstein (Carnegie Mellon University)

In this talk, Seth Goldstein will describe the hardware and software challenges involved in realizing Claytronics, a form of programmable matter. The goal of the claytronics project is to create ensembles of cooperating submillimeter robots, which work together to form dynamic 3D physical objects. For example, claytronics might be used in telepresense to mimic, with high-fidelity and in 3-dimensional solid form, the look, feel, and motion of the person at the other end of the "telephone" call. To achieve this long-range vision we are investigating hardware mechanisms for constructing sub-millimeter robots, which can be manufactured en masse using photolithography. In parallel with our hardware effort, we are developing novel distributed programming languages and algorithms to control the ensembles.

Seth Copen Goldstein's research focuses on computing systems and nanotechnology. Broadly speaking, Seth's research is aimed at understanding systems nanotechnology. Among his research efforts are three projects: the Phoenix project, the Claytronics project, and brain in a bottle. The common theme among these projects is to understand how to design, manufacture,

program, and use robust reconfigurable systems built with massive numbers of similar, and often unreliable, programmable units. Goldstein joined the faculty at Carnegie Mellon University in 1997. He received his Masters and Ph.D. in computer science at the University of California at Berkeley. Before attending UC Berkeley, Seth was CEO and founder of Complete Computer Corporation. His undergraduate work was at Princeton University.

IAAI-08 Invited Talk

Boss, the Urban Challenge, and the Promise of Autonomous Driving

Chris Urmson (Carnegie Mellon University)

The DARPA Urban Grand Challenge was a 60-mile race through an urban roadway where vehicles had to follow the same rules of the road that human drivers are expected to respect. This challenge differed from the two previous Grand Challenges in that the robots not only had to navigate the course, but they were required to drive safely in the presence of other human-driven cars as well as the other robot competitors. Urmson will describe the details of the DARPA Urban Challenge and describe Carnegie Mellon's entry Boss, the autonomous vehicle that won the challenge. He will describe the overall system architecture and highlight the many component technologies that made up Boss.

Chris Urmson researches and develops autonomous navigation algorithms, architectures and systems. He is the director of technology for the Urban Challenge at Carnegie Mellon University and was a principal architect of the Red Team and Red Team Too Grand Challenge entries. He has been part of developing numerous field robots and has tested them in places as exotic as the arctic circle, the Atacama Desert and Pittsburgh, Pennsylvania. He earned his Ph.D. in 2005 from Carnegie Mellon University and his B.Sc. in computer engineering from the University of Manitoba in 1998.