Invited Talks

Grounded Language Learning

Raymond J. Mooney (University of Texas at Austin)

Most approaches to semantics in computational linguistics represent meaning in terms of words or abstract symbols. Grounded-language research bases the meaning of natural language on perception and/or action in the (real or virtual) world. Machine learning has become the most effective approach to constructing natural-language systems; however, current methods require a great deal of laboriously annotated training data. Ideally, a computer would be able to acquire language like a child, by being exposed to language in the context of a relevant but ambiguous environment, thereby grounding its learning in perception and action. We will review recent research in grounded language learning and discuss future directions.

Raymond J. Mooney is a professor in the Department of Computer Science at the University of Texas at Austin. He received his Ph.D. in 1988 from the University of Illinois at Urbana/Champaign. He is an author of over 150 published research papers, primarily in the areas of machine learning and natural language processing. He was the president of the International Machine Learning Society from 2008-2011, program cochair for AAAI 2006, and is a AAAI and ACM Fellow. His recent research has focused on learning for natural-language processing, statistical relational learning, active transfer learning, and connecting language, perception and action.

Robert S. Engelmore Memorial Award Lecture

Deborah L. McGuinness (Rensselaer Polytechnic Institute)

Deborah L. McGuinness is the Tetherless World Senior Constellation Chair, professor of computer and cognitive science, and founding director of Rensselaer Polytechnic Institute's Web Science Research Center. McGuinness is a leading authority on the semantic web and has been working in knowledge representation and reasoning environments for over 25 years. Her primary research focuses on making smart systems understandable and usable by a broad range of people. She leads active research efforts in explanation, trust, ontology environments, and provenance. McGuinness is also known for semantic application environments, particularly for eScience frameworks such as the semantic escience framework and demonstration portals including many in natural science and health informatics settings. McGuinness also founded McGuinness Associates — a small woman owned business — that consults on semantic applications in a wide range or areas with recent focus on health and environmental informatics, context-aware mobile computing, and next generation journalism.

Fighting the Tuberculosis Pandemic Using Machine Learning

Kristin P. Bennett and the TB-Insight Team (Rensselaer Polytechnic Institute)

Tuberculosis (TB) infects one third of the world's population and is the second leading cause of death from a single infectious agent worldwide. The emergence of drug resistant TB remains a constant threat. We examine how machine learning methods can help control tuberculosis. DNA fingerprints of *Mycobacterium tuberculosis* complex bacteria (Mtb) are routinely gathered from TB patient isolates for every tuberculosis patient in the United States to support TB tracking and control efforts. We develop learning models to predict the genetic lineages of Mtb based on DNA fingerprints. Mining of tuberculosis patient surveillance data with respect to these genetic lineages helps discover outbreaks, improve TB control, and reveal Mtb phenotype differences. We discuss learning- and visualization-based tools to support public health efforts towards TB control in development for the New York City Health Department.

Title TBA

Vijay Kumar (University of Pennsylvania)

Vijay Kumar is the UPS Foundation professor in the School of Engineering and Applied Science at the University of Pennsylvania, and on sabbatical leave at White House Office of Science and Technology Policy where he serves as the assistant director for robotics and cyber physical systems. He received his bachelors of technology from the Indian Institute of Technology, Kanpur and his Ph.D. from The Ohio State University in 1987. He has been on the faculty in the Department of Mechanical Engineering and Applied Mechanics with a secondary appointment in the Department of Computer and Information Science at the University of Pennsylvania since 1987.

Kumar served as the deputy dean for research in the School of Engineering and Applied Science from 2000-2004. He directed the GRASP Laboratory, a multidisciplinary robotics and perception laboratory, from 1998-2004. He was the chairman of the Department of Mechanical Engineering and Applied Mechanics from 2005–2008. He then served as the deputy dean for education in the School of Engineering and Applied Science from 2008-2012.

Kumar is a Fellow of the American Society of Mechanical Engineers (ASME) and the Institution of Electrical and Electronic Engineers (IEEE). He has served on the editorial boards of the *IEEE Transactions on Robotics and Automation, IEEE Transactions on Automation Science and Engineering, ASME Journal of Mechanical Design,* the *ASME Journal of Mechanisms and Robotics* and the *Springer Tract in Advanced Robotics (STAR).* He is the recipient of the 1991 National Science Foundation Presidential Young Investigator award, the 1996 Lindback Award for Distinguished Teaching (University of Pennsylvania), the 1997 Freudenstein Award for significant accomplishments in mechanisms and robotics, the 2012 ASME Mechanisms and Robotics award, the 2012 IEEE Robotics and Automation Society distinguished service award and a 2012 World Technology Network award. He has won best paper awards at DARS 2002, ICRA 2004, ICRA 2011, and RSS 2011 and has advised doctoral students who have won Best Student Paper awards at ICRA 2008, RSS 2009, and DARS 2010. He is also a distinguished lecturer in the IEEE Robotics and Automation Society and an elected member of the Robotics and Automation Society Administrative Committee (2007-2012). His research interests are in robotics, specifically multirobot systems, and microaerial vehicles.

Socially Assistive Robotics: Human-Robot Interaction Methods for Creating Robots That Care

Maja J. Mataric (University of Southern California)

Socially assistive robotics (SAR) is a new field of intelligent robotics that focuses on developing machines capable of assisting users through social rather than physical interaction. The robot's physical embodiment is at the heart of SAR's effectiveness, as it leverages the inherently human tendency to engage with lifelike (but not necessarily humanlike or otherwise biomimetic) social behavior. People readily ascribe intention, personality, and emotion to robots; SAR leverages this engagement stemming from non-contact social interaction involving speech, gesture, movement demonstration and imitation, and encouragement, to develop robots capable of monitoring, motivating, and sustaining user activities and improving human learning, training, performance and health outcomes. Human-robot interaction (HRI) for SAR is a growing multi-faceted research area at the intersection of engineering, health sciences, neuroscience, social, and cognitive sciences. This talk will describe our research into embodiment, modeling and steering social dynamics, and long-term user adaptation for SAR. The research will be grounded in projects involving analysis of multimodal activity data, modeling personality and engagement, formalizing social use of space and non-verbal communication, and personalizing the interaction with the user over a period of months, among others. The presented methods and algorithms will be validated on implemented SAR systems evaluated by human subject cohorts from a variety of user populations, including stroke patients, children with autism spectrum disorder, and elderly with Alzheimers and other forms of dementia.

Maja Mataric is a professor and Chan Soon-Shiong chair of computer science, neuroscience, and pediatrics at the University of Southern California, founding director of the USC Center for Robotics and Embedded Systems, codirector of the USC Robotics Research Lab and vice dean for Research in the USC Viterbi School of Engineering. She received her MS and PhD in computer science and AI from the Massachusetts Institute of Technolgy and her BS in computer science from the University of Kansas. She is a Fellow of AAAS and IEEE, and recipient of the Presidential Awards for Excellence in Science, Mathematics and Engineering Mentoring, the Anita Borg Institute Women of Vision Award for Innovation, Okawa Foundation Award, NSF Career Award, the MIT TR35 Innovation Award, and the IEEE Robotics and Automation Society Early Career Award. Mataric is an associate editor of three major journals, has published extensively, and served on the NSF CISE Advisory Committee, among other advisory boards. She is actively involved in K-12 educational outreach to engage student interest in science, technology, engineering, and math (STEM) topics. Her research into socially assistive robotics aims to endow robots with the ability to help people through individual noncontact assistance in convalescence, rehabilitation, training, and education, with applications for children with autism spectrum disorders, survivors of stroke and traumatic brain injury, and individuals with Alzheimer's Disease and other forms of dementia.

Poker AI: Algorithms for Creating Game-Theoretic Strategies for Large Incomplete-Information Games

Tuomas Sandholm (Carnegie Mellon University)

Incomplete-information games — such as most auctions, negotiations, and future (cyber)security settings — cannot be solved using minimax search even in principle. Completely different algorithms were needed. A dramatic scalability leap has occurred in our ability to solve such games over the last seven years, fueled

largely by the Annual Computer Poker Competition. I will discuss the key domain-independent techniques that enabled this leap, including automated abstraction techniques and approaches for mitigating the issues that they raise, new equilibrium-finding algorithms, safe opponent exploitation methods, techniques that use qualitative knowledge as an extra input, and endgame solving techniques. I will finish by benchmarking poker programs against the best human poker professionals and by discussing what humans can learn from the programs.

Tuomas Sandholm is a professor at Carnegie Mellon University in the Computer Science Department, Machine Learning Department, and the CMU/UPitt joint Ph.D. program in computational biology. He has published over 450 papers on market design, game theory, search and integer programming, auctions and exchanges, automated negotiation and contracting, coalition formation, voting, safe exchange, normative models of bounded rationality, resource-bounded reasoning, and machine learning. In parallel with his academic career, he was founder, chairman, and CTO/chief scientist of CombineNet, Inc. from 1997 until its acquisition in 2010; during this period the company commercialized 800 of the world's largest-scale generalized combinatorial auctions, totaling over $60 billion. Sandholm's algorithms also run the US-wide UNOS kidney exchange. He is the founder and CEO of Optimized Markets, Inc., a startup that is bringing a new paradigm to advertising sales and scheduling. Among his many honors are the NSF Career Award, inaugural ACM Autonomous Agents Research Award, Sloan Fellowship, Carnegie Science Center Award for Excellence, and Computers and Thought Award. He is a Fellow of the ACM and AAAI.

Joint AAAI/IAAI Invited Talk

Telling Stories at Internet Scale

Larry Birnbaum (Northwestern University)

Taking full advantage of the massive scale of "big data" will require technologies for analyzing and communicating these data to people, in terms they can understand and act on, at an equally massive scale. The automatic generation of narratives from data offers the promise of meeting this critical need. Our technology, which leverages human editorial judgment at scale, is today generating millions of stories from data, including highly personalized stories, in domains varying from business operations, to sports, education, medicine, and finance. The resulting narratives are often indistinguishable from those written by human analysts and writers.

Larry Birnbaum is an associate professor of electrical engineering and computer science, and of journalism, at Northwestern University, as well as a founder and chief scientific advisor of Narrative Science Inc. At Northwestern, Birnbaum is codirector of the Intelligent Information Lab, and a founder of the Knight Lab, an interdisciplinary center for innovation in technology, media, and journalism. He formerly served as chair of the Computer Science Department. At Narrative Science, Birnbaum focuses on technology and IP. Birnbaum received his BS and PhD degrees from Yale, and was on the faculty there before joining Northwestern.

IAAI-13 Invited Talk

Title TBA

Lawrence Hunter (University of Colorado)

Lawrence Hunter is the director of the University of Colorado's Computational Bioscience Program and a professor of pharmacology (School of Medicine) and computer science (Boulder). He received a Ph.D. in computer science from Yale University in 1989, and then joined the National Institutes of Health as a staff scientist, first at the National Library of Medicine and then at the National Cancer Institute, before coming

to Colorado in 2000. Hunter is widely recognized as one of the founders of bioinformatics; he served as the first president of the International Society for Computational Biology (ISCB), and created several of the most important conferences in the field, including ISMB, PSB and VizBi. Hunter's research interests span a wide range of areas, from cognitive science to rational drug design. He has published more than 100 scientific papers, holds two patents and has been elected a fellow of both the ISCB and the American College of Medical Informatics. His primary focus recently has been the integration of natural language processing, knowledge representation, machine learning and advanced visualization techniques to address challenges in interpreting data generated by high throughput molecular biology.

EAAI-13 Invited Talk

Title TBA

Daniel Klein (University of California, Berkeley)

Dan Klein (PhD Stanford) is an associate professor of computer science at the University of California, Berkeley. He works on natural language processing, machine learning, and other areas of AI. There may come a day when his research systems rise up and give his talks for him. Until then, he teaches a lot, from an undergraduate AI class with over 300 people per semester to an online version with over 30,000. Klein's research honors include a Microsoft Faculty Fellowship, a Sloan Fellowship, an NSF CAREER award, the ACM Grace Murray Hopper award for his work on grammar induction, and best paper awards at the ACL, NAACL, and EMNLP conferences. He is particularly proud to have been recognized with multiple teaching honors, including the University of California, Berkeley's Distinguished Teaching Award.