IJCAI-09 Awards and Distinguished Papers

Awards

The IJCAI-09 Award for Research Excellence

Scaling AI through Multi-Agent Organizations

Victor Lesser (University of Massachusetts Amherst, USA)

Scaling remains one of the grand challenges for AI. Lesser has been using organizational control to build multiagent systems with hundreds to thousands of intelligent agents. This approach can also be used to structure complex AI systems with extensive and heterogeneous knowledge. Organizational control is a multi-level approach in which organizational goals, roles, and responsibilities are dynamically developed, distributed, and maintained to serve as guidelines for making detailed operational control decisions by the individual agents. Lesser will illustrate the use of organizational control in three distributed application areas: (1) an adaptive sensor and interpretation vehicle-tracking network, (2) a peer-to-peer information search and retrieval system, and (3) a self-improving task allocation system. He will highlight the important balance between externally-directed and self-directed agent activities in uncertain and dynamic environments. Then he will present the continuing research challenges, including how to automate the design of an organization and evolve it as conditions change, create an organizationally situated agent, and evaluate and predict an organization's performance.

The IJCAI-09 Computers & Thought Awards

How Optimized Environmental Sensing Helps Address Information Overload on the Web

Carlos Guestrin (Carnegie Mellon University, USA)

In this talk, Guestrin tackles a fundamental problem that arises when using sensors to monitor the ecological condition of rivers and lakes, the network of pipes that bring water to our taps, or the activities of an elderly individual when sitting on a chair: Where should we place the sensors in order to make effective and robust predictions?

Such sensing problems are typically NP-hard, and in the past, heuristics without theoretical guarantees about the solution quality have often been used. In this talk, Guestrin presents algorithms, which efficiently find provably near-optimal solutions to large, complex sensing problems. His algorithms are based on the key insight that many important sensing problems exhibit submodularity, an intuitive diminishing returns property: Adding a sensor helps more the fewer sensors we have placed so far. In addition to identifying most informative locations for placing sensors, his algorithms can handle settings, where sensor nodes need to be able to reliably communicate over lossy links, where mobile robots are used for collecting data or where solutions need to be robust against adversaries and sensor failures.

Guestrin presents results applying his algorithms to several real-world sensing tasks, including environmental monitoring using robotic sensors, activity recognition using a built sensing chair, and a sensor placement competition. He concludes with drawing an interesting connection between sensor placement for water monitoring and addressing the challenges of information overload on the web. As examples of this connection, he addresses the problem of selecting blogs to read in order to learn about the biggest stories discussed on the web, and personalizing content to turn down the noise in the blogosphere.

STAIR: The Stanford Artificial Intelligence Robot Project

Andrew Y. Ng (Stanford University, USA)

This talk will describe the STAIR home assistant robot project, and the satellite projects that led to key STAIR components such as (1) robotic grasping of previously unknown objects, (2) depth perception from a single still image, (3) practical object recognition using multimodal sensors, and (4) a software architecture for integrative AI.

Since its birth in 1956, the AI dream has been to build systems that exhibit broad-spectrum competence and intelligence. STAIR revisits this dream, and seeks to integrate onto a single robot platform tools drawn from all areas of AI including learning, vision, navigation, manipulation, planning, and speech and NLP. This is in distinct contrast to, and also represents an attempt to reverse, the 30 year old trend of working on fragmented AI sub-fields. STAIR's goal is a useful home assistant robot, and over the long term, we envision a single robot that can perform tasks such as tidying up a room, using a dishwasher, fetching and delivering items, and preparing meals.

In this talk, Ng will describe our progress on having the STAIR robot fetch items from around the office, and on having STAIR take inventory of office items. Specifically, he'll describe learning to grasp previously unseen objects (including unloading items from a dishwasher); probabilistic multi-resolution maps, which enable the robot to open or use doors; and a robotic foveal plus peripheral vision system for object recognition and tracking. Ng will also outline some of the main technical ideas—such as learning 3-D reconstructions from a single still image, and reinforcement learning algorithms for robotic control—that played key roles in enabling these STAIR components.

The Donald E. Walker Distinguished Service Award

Luigia Carlucci Aiello (Università di Roma 'La Sapienzá, Italy)

IJCAI Awards Selection Committee

Jaime Carbonell (Carnegie Mellon University, USA)
Anthony G. Cohn (University of Leeds, UK) (Chair)
Tom Dietterich (Oregon State University, USA)
Malik Ghallab (INRIA-CNRS, France)
Fausto Giunchiglia (University of Trento, Italy)

IJCAI-09 Distinguished Papers

Learning Conditional Preference Networks with Queries
Frédéric Koriche and Bruno Zanuttini

Consequence-Driven Reasoning for Horn *SHIQ* Ontologies
Yevgeny Kazakov