

Advances in Cognitive Systems

Pat Langley

The AAAI Fall Symposium on Advances in Cognitive Systems was held in Arlington, Virginia, from November 4 to 6, 2011. The meeting's goal was to bring together researchers interested in developing intelligent systems that demonstrate the full range of human cognitive abilities and to report progress on this daunting task.

The original aims of artificial intelligence, when it was launched in the late 1950s, were to explain intelligence in information-processing terms and to reproduce the entire range of human cognitive abilities in computational artifacts. Although the field has seen impressive advances in the last few decades, many researchers have, in the process, forgotten or abandoned these important goals. The purpose of the Fall Symposium on Advances in Cognitive Systems (<http://www.cogsys.org/acs/2011/>) was to congregate scientists who remain committed to AI's original vision. The meeting received 50 paper submissions and it was attended by over 75 participants, indicating that there remains substantial interest in this view on the discipline.

Research in cognitive systems, as represented by contributors to the meeting, differs from what has become mainstream AI in five basic ways. First, it retains a concern with high-level aspects of cognition, such as the ability to engage in multi-step inference, understand the meaning of natural language, generate novel plans that achieve goals, and even reason about one's own reasoning. Second, cognitive systems emphasizes the role of structured representations in intelligence, recognizing that both computers and humans are general symbol manipulators, not just numeric processors. Third, it approaches intelligence from a systems perspective, examining how different abilities fit together rather than focusing on component algorithms. Fourth, the field draws ideas and inspiration from findings about human cognition, both for candidate mechanisms and for challenging problems to drive research. Finally, it adopts flexible views on evaluation, welcoming papers that demonstrate new functionality, present novel approaches to established problems, analyze challenging cognitive tasks, and report on architectures for integrated intelligence. Together, these characteristics identify cognitive systems as a separate paradigm that remains committed to the original vision for artificial intelligence.

The symposium included 22 talks and 23 posters that reflected a number of distinct but connected themes. Some presentations dealt with flexible mechanisms for conceptual inference over commonsense knowledge that move beyond traditional deductive approaches. A second theme dealt with deep language processing, at both the sentence and dialogue levels, that combines linguistic and domain knowledge to produce genuine representations of meaning. A related topic was reasoning about the mental states of other agents, an ability essential to dialogue and joint activity, with a number of frameworks relying on partitioned memories (divided into 'worlds' or 'contexts') to keep mental models distinct. A fourth emphasis involved mechanisms for acquiring structural knowledge from experience and instruction, most using some form of explanation to generate candidate structures. Together, these themes clarified the distinctive character of research on cognitive systems and its concern with producing human-like computational artifacts.

Open discussions during the meeting focused on three important issues. One concerned criteria for evaluating research progress, and for publishing papers, on cognitive systems. There was general agreement that existing conferences are not supportive of work in the paradigm, and that the community should find ways to broaden existing venues or establish new ones so that submissions will receive appropriate treatment. Another discussion topic dealt with challenges that might drive

research in the area. Proposals included mechanisms for flexible and scalable inference, interactive approaches to acquiring knowledge over extended periods, techniques that interface general-purpose reasoning with specialized methods for spatial and visual inference, and motivational processes that influence reasoning and goal generation. Finally, participants discussed resources that would support cognitive systems research, with the main candidate being a Wiki that incorporates pointers to software, course material, challenge problems, and data sets.

There was general consensus among participants that the symposium served a genuine need within the AI community and that additional meetings were desirable. There was also a sense that, although the cognitive systems paradigm adopts the original aims of AI, its modern incarnation is relatively new and, to ensure its success as a scientific discipline, we must clarify its distinctive characteristics, foster a community of like-minded researchers, identify and make progress on challenge problems, establish venues for publishing results, and train promising new scientists. The symposium was only one step on the road toward the audacious goal of computational systems that exhibit the same forms of intelligence as humans, and this task will keep the community occupied for years to come.

Pat Langley served as the chair of the meeting, with Paul Bello (Office of Naval Research), Nicholas Cassimatis (Rensselaer Polytechnic Institute), Kenneth Forbus (Northwestern University), John Laird (University of Michigan, Ann Arbor), and Sergei Nirenburg (University of Maryland, Baltimore County) serving as co-organizers. Papers from the symposium appear in Technical Report FS-11-01 from AAAI Press.

Pat Langley is Professor of Computer Science and Engineering at Arizona State University in Tempe, Arizona, and Director of the Institute for the Study of Learning and Expertise in Palo Alto, California.