AI Magazine Volume 26 Number 4 (2006) (© AAAI)

AAAI: It's Time for Large-Scale Systems

Daniel G. Bobrow

Daniel G. Bobrow

■ The most important challenge facing AI today is enabling components to interact in larger scale systems, where modules built with multiple alternative methodologies can be incorporated into robust applications.

Today AAAI is again at a time of transition. The infrastructure—computing power, memory, bandwidth, and connectivity—has evolved dramatically. Important theoretical advances have been made in areas such as machine learning, natural language, knowledge representation, task descriptions, sensing, and action in the world. Once again there is substantial demand for AI applications from customers such as DARPA, with a requirement to solve real problems.

We need to find ways making AI components interact in larger scale systems. Researchers need to explore how multiple methodologies can be incorporated into robust applications. We should not treat alternative approaches as competitive, but seek to find ways using such alternatives synergistically.

For example, probabilistic reasoning and symbolic reasoning should be brought together to provide rapid solutions to the likely cases, and explainable alternatives derived from symbolic models for the rarer cases. We need to make results of machine learning understandable and modifiable through symbolic methods. We need a spectrum of methods to support improvement of system performance from experience. Rule-based systems improve through the learning of their creators and further programming. Machine learning systems improve by being given more human annotated examples and feedback on their behavior. We need to incorporate learning between these two— where the user can instruct a system in how its behavior should change, without the need to know how the system functions. Systems need the ability to explain their behavior, reassure users with justifications for their answers, provide guidance to use the system better, and even educate users about how to think about problems.

In other areas, bringing together different AI subfields can potentially solve problems that have resisted more narrowly focused approaches. Natural language systems that integrate computational linguistics with knowledge representation and reasoning engines, for example, may finally enable systems that can understand and interact with humans in robust and generally useful ways. They may help us realize the dream of systems that can learn by reading.

The pressing need now for pulling divergent approaches together stems in part from transitions that occurred during my term as AAAI president (1989–91). We were just recovering from an exuberance engendered by money and attention flowing from business people who were selling the idea that AI was going to transform the world NOW—especially through expert systems. Companies showing LISP machines, knowledge-based programming environments, and AI solutions occupied large booths at jammed AI conferences. AI had some notable successes—the intermediate Fredkin prize for chess was won when Deep Thought beat an international master; IBM's AI-based LMS (Logistics Management System) was helping run its Vermont semiconductor facility. The first Innovative Applications in Artificial Intelligence Conference (1989) described a significant number of deployed applications with measurable profit.

But in late 1989, there was a predicted "AI winter," fueled by increasing public realization that the "expert" in "expert systems" was an aspiration rather than a guarantee. In the field, there was a clear appreciation of a need for deeper theory, better engineering, broader coverage, and understanding of limits of system capabilities. Paul Cohen's seminal article, analyzing the AAAI-90 papers, reflected on the need for a more effective methodology for AI. Papers focused on systems showed results on realistic problems but provided no testable hypotheses or formal analyses. Those centered on models and algorithms used primarily abstract examples but provided some analyses of the complexity or coverage. Cohen provided guidelines for a more mature science of artificial intelligence.

AAAI as an organization also started to take a longer view. AAAI presidents went from a one-year term to serving for two years as president, with continuity provided by two years as president-elect and two years as past-president. An AAAI fellows program was created to honor significant contributions to the field, since larger, more diffuse societies such as the ACM or IEEE were not in a position to evaluate these contributions. Carol Hamilton became executive director and built a stable organization to support the AAAI activities; the AAAI office became the preferred means of support for workshops in many specialized subfields of AAAI.

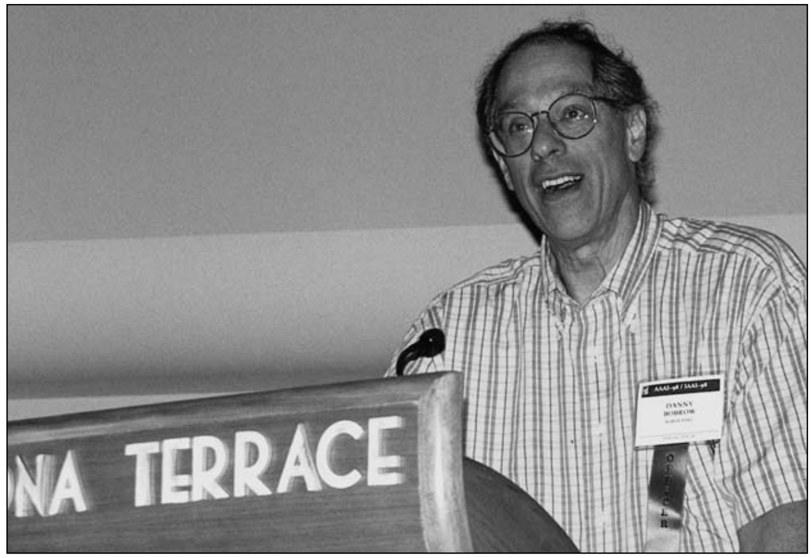

Bobrow Speaking at the Fifteenth National Conference in Madison, Wisconsin, July, 1998.

The workshops led to an increasing number of focused conferences, from knowledge representation to machine learning; from constraint-based programming to intelligent user interfaces; from planning to computational linguistics. Many of these still receive logistical support from the AAAI office and sometimes are colocated with AAAI. Some are bigger than AAAI now. The growth of these specialized conferences represents both a triumph and a failure for AAAI. While the support and focus for these diverse areas has improved, it has meant a loss of a unifying vision that facilitates their integration. The challenge is to bring together all these perspectives to enable large-scale effective systems that interact intelligently with people, the physical environment, and the ubiquitous digital systems that surround us.

Daniel G. Bobrow is a research fellow in the Intelligent Systems Laboratory of the Palo Alto Research Center. He received his Ph.D. in artificial intelligence from the Massachusetts Institute of Technology and has published more than 100 papers. Bobrow served as president of AAAI from 1989–1991, and is a fellow of the AAAI and of the ACM. His research includes topics in knowledge-based systems, language understanding, programming language design, computer-supported collaborative work, and systems that bring together social and technical design in a single participatory framework. His current research focus is on systems to read and understand natural language.