

Getting Back to “The Very Idea”

Ron Brachman



Ron Brachman.

■ For many years, the very idea of artificial intelligence has been provocative and exciting. However, with a continually increasing focus on specialized subareas and somewhat narrow technical problems (both of which are inevitable and in many ways healthy), we may be torpedoing our core research agenda: the creation of a true synthetic intelligence. I reflect briefly on the essential interdependencies of the components of intelligence, the important roles of architecture and integration, and the need to get back to thinking about the very idea of AI. AAAI's role in the field has evolved over the years, but after a quarter-century as an organization, and a half-century as a field, it seems like AAAI is in an ideal situation to bring AI as a whole back to its roots.

Each time my eye wanders across the jumble of books on my shelves, the same one always seems to grab me. In 1985, the philosopher John Haugeland wrote a thought-provoking treatise on AI that he titled *Artificial Intelligence: The Very Idea*. I am pretty sure that it's the title that keeps pulling me back. It's hard to glance even briefly at a phrase like that without stopping to reflect on what our field is all about. On top of that, it's always fun to reread the stirring terms in which Haugeland describes AI: "...by now not only psychology but also a host of related disciplines are in the throes of a great intellectual revolution. And the epitome of the entire drama is Artificial Intelligence, the exciting new effort to make computers think."

As I think back on my recent term as AAAI president, I remember musing fairly often, both alone and with the executive council, about where the field was with respect to "the very idea" and about whether the excitement



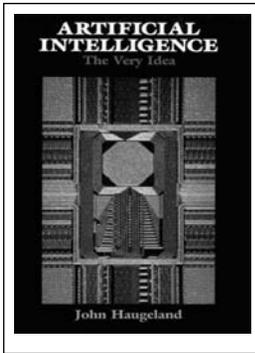
AAAI President Ron Brachman Confers with AAAI-05 Conference Cochair Manuela Veloso at the 2005 National Conference in Pittsburgh, Pennsylvania.

and drama were still with us. In many ways, AI has evolved to being a label on a family of relatively disconnected efforts. Extremely successful and vibrant areas, such as machine learning, robotics, and reasoning under uncertainty, have developed their own conferences and journals and roil with intellectual excitement and new ideas. In some ways—as reflected by several of my colleagues in this issue of *AI Magazine*—this should be counted as a great success for our field and for AAAI as an organization (especially in that AAAI has given birth to and nurtured many of the specialized conferences). The breaking up into subfields is probably an inevitable consequence of the maturity that fifty years of life brings to scientific pursuits. But it has also become common for incoming AAAI presidents and executive councilors to wonder aloud about how we might bring the increasingly disparate subfields of AI back together.

In our excitement about each of our own individual research areas, we may have lost sight of the very idea of AI. As we focus on more spe-

cialized and deeper technical results in learning algorithms and corpus-driven machine translation (for example), we are probably not thinking about the fact that even the world's greatest statistical classifier or the world's best machine translation (MT) algorithm will not end up being a fully intelligent system at the end of its research road. That per se is not an issue for machine learning or natural language processing as individual focal subfields. But the predominance of this style of thinking is indeed an issue for AI as a larger endeavor—and, consequently, for AAAI.

As I reflected in my recent AAAI presidential address (to be transcribed in a future issue of this magazine), my own belief is that the centrifugal intellectual force in our field is not only making AAAI less of a meaningful venue for AI scientists and practitioners—which, by itself, would not be the end of the world: institutions can certainly outlive their usefulness and should be replaced by new ones when they do—it may also be having a more pernicious effect on our very research. Intelligence seems to



Courtesy, The MIT Press

John Haugeland's Book
was Published by The
MIT Press in 1985.

emerge from a complex combination of many specialized abilities, such as reasoning, learning, perception, and language use. But it does not seem to be a simple compositional result of the mere juxtaposition of reasoning, learning, perception, and language “boxes.” Rather, there is some set of deep, essential interdependencies that tie these elements together: learning must result in knowledge that reasoners can use in many different ways, perception seems to be intrinsically guided by cognition, and so on. If we don't have people studying the very idea of artificial intelligence, focusing on architecture and interdependency, all the great work in the world on the various pieces will never add up to the grand dream of our field—a true synthetic intelligence. As we spin off areas, we may actually be destroying our real research agenda.

One of the unique aspects of my time as AAAI president was the concurrent privilege I was afforded of serving as the director of the Information Processing Technology Office at the Defense Advanced Research Projects Agency (DARPA). DARPA—and the Information Processing Technology Office (IPTO) more specifically—has had a major impact on the history of AI since the early 1960s, providing hundreds of millions of dollars to the research community in support of key efforts in the field. While a drop in DARPA funding many years ago was at least partly responsible for somewhat of an “AI winter,” over my term there we were strongly supported by the agency director and Congress in pursuing an aggressive—and well-funded—AI-oriented agenda. Our focus was on what we came to call “cognitive systems,” and this allowed us to do something tangible about bringing different elements of AI together in integrated systems. Among other things, we strongly emphasized overarching architectural considerations for tightly integrating reasoning, learning, perception, and language use, and we were able to help create some new teams of people that normally would not attend the same specialized conferences or share their results with one another. We also put considerable emphasis on the evaluation of complete, integrated AI systems, thereby challenging those teams to collaborate to satisfy metrics that depended heavily on the components working together. The early results are encouraging, and the funding level continues to be significant. Both of these bode well for the future of the very idea of an artificial intelligence, and the cognitive systems agenda may be a useful model for AAAI to use as it moves to reinvigorate its core mission.

There is only so much a single scientific or-

ganization can do to affect the social dynamic of a field of study, but AAAI has arguably had a profound effect on the still young field of artificial intelligence. One of the exciting things about serving as the steward of the organization is one's ability to make a real contribution to its members not only by responding to their needs but also by guiding AAAI's resources in helpful new directions (or in this case, helpful old ones). As we look forward to a time of renewed excitement about the very idea of AI, I can see AAAI playing an essential role as the main venue for the reunification of the spun-off subfields. Having our conference and publications highlight new results in intelligent system architecture and integration of specialized results from different areas would recapture a unique purpose for AAAI and would bring it back to the mission envisioned by its founders twenty-five years ago.

Back in 1985 Haugeland suggested that AI was “neither preposterous nor inevitable.” The former is reassuring for those of us who have invested our careers in the field. The latter says that AAAI can really make a difference. After a period of scientific dispersion in which it may have lost a sense of the very idea it was created to pursue, AAAI now seems on the path to resuming its starring role in the great intellectual drama of minds and machines. I personally feel honored to have been a small part of it. And I, for one, am eager to see how it all turns out.

Reference

Haugeland, John 1985. *Artificial Intelligence: The Very Idea*. Cambridge, MA: Bradford Books.

Ron Brachman earned his B.S.E.E. degree from Princeton University (1971), and his S.M. (1972) and Ph.D. (1977) degrees from Harvard University. He has made numerous important contributions to the area of knowledge representation, including a recent textbook with Hector Levesque, *Knowledge Representation and Reasoning*. Brachman started his career at Bolt Beranek & Newman, spent several years at Fairchild/Schlumberger's Lab for AI Research, and, having developed a world-class AI group at AT&T Bell Laboratories, moved into senior research management jobs at Bell Labs and AT&T Labs. He served as the director of DARPA's Information Processing Technology Office from 2002 to 2005, and there developed IPTO's cognitive systems initiative, which brought hundreds of millions of dollars to the national research community. He recently joined Yahoo! Research as its vice president of Worldwide Research Operations. Brachman is a Fellow of ACM and AAAI, was program chair for AAAI-84, and was president of AAAI from 2003–2005.