

Whither AI

Identity Challenges of 1993–95

Barbara J. Grosz



Barbara J. Grosz.

■ The 1993–95 period presented various “identity challenges” to the field of AI and to AAAI as a leading scientific society for the field. The euphoric days of the mid-1980s AI boom were over, various expectations of those times had not been met, and there was continuing concern about an AI “winter.” The major challenge of these years was to chart a path for AI, designed and endorsed by the broadest spectrum of AI researchers, that built on past progress, explained AI’s capacity for addressing fundamentally important intellectual problems and realistically predicted its potential to contribute to technological challenges of the coming decade. This reflection piece considers these challenges and the ways in which AAAI helped the field to move forward.

During the years I served as AAAI president (1993–95), the society was 13–15 years old, I was entering my third decade as an AI researcher, and AI as a field was entering its forties. Adolescence, the twenties, and the forties each bring particular “developmental” challenges to people, and, though surely coincidentally, elements of those life stages seem also to characterize the period of my presidency. As do people in their early teens, researchers in AI were grappling with the field’s identity and how to distinguish it, in a positive way, from related fields addressing similar problems; like the twenty-something’s search for a meaningful career, our field was in search of how best to characterize its potential to contribute to the various national needs that were driving funding agencies’ programs; and, like those in middle age, we searched for ways to explain AI’s past contributions not only within the AI context but also to computer science both scientifically and technologically as

well as more generally to important intellectual challenges.

The euphoric days of the mid-1980s AI boom, which was fueled in part by reaction to the Japanese Fifth Generation project, were over. AI had gone from an obscure field—I was no longer asked if artificial intelligence was something the CIA created—to a highly visible one, but expectations of the late 1980s had not been met. There was continuing concern about reactions to unfulfilled promises leading to an AI “winter.” Our field had made great progress, but there was a large gap between what AI applications and systems were able to do and the expectations generated in the 1980s. Industry support had declined, and funding agency program directors were asking for help justifying sustaining then-current levels of support for AI research. AAAI conference attendance and membership were also down. There was much discussion and debate about what AAAI could do as well as about how AI related to the rest of computer science and the roles that AI research could play in building the “information superhighway,” that is, in new national ventures, most especially the “national information infrastructure” and “high performance computing” programs being launched by the federal government.

The major challenge of these years for the field’s leaders was to chart a path for AI, one that built on past progress, explained AI’s capacity for addressing fundamentally important intellectual problems, and realistically (insofar as possible) predicted its potential to contribute to technological challenges of the coming decade. The major challenge of my presidency, which it seemed crucial to tackle immediately, was to muster AI’s leaders to meet this challenge. It was important to bring career-younger people into this process, both to build on their expertise and enthusiasm and to ensure that they too felt ownership of whatever new paths were charted. This much of the “state of AI” came to mind immediately when I was asked to contribute my recollections to this issue of *AI Magazine*.

In trying to ground these memories, to see if they were indeed accurate, I reviewed minutes of various AAAI Executive Council meetings in the early 1990s. The minutes recorded that two workshops took place during the first year of my presidency, which, fortunately for these current reflections, led to reports that were printed in *AI Magazine*. The first, “A Report to ARPA on Twenty-First Century Intelligent Systems” (21st century; *AI Magazine* 15:3, 10–20, 1994), resulted from a meeting AAAI organized “to assist ARPA in defining an agenda for foundational AI research.” The second, “The Role of



*The Report to ARPA on Twenty-First Century Intelligent Systems
Appeared in this AI Magazine Special Issue on Agents.*

Intelligent Systems in the National Information Infrastructure” (NII; *AI Magazine*, 16:3, 45–64, 1995), reports on a workshop organized by AAAI and cosponsored by the National Science Foundation with two complementary purposes: to increase awareness of the “opportunities presented by the National Information Infrastructure activities” in the AI community and to “identify key contributions of research in AI to the NII.”

With some trepidation I downloaded and reread these reports. There are great dangers in defining a research agenda for a field and predicting five to ten years into the future. Bringing any small group of people together to formulate plans for a whole field runs the risk of producing recommendations that promote individual agendas. Hindsight often reveals great ignorance, narrow perspectives, or lack of imagination. These reports, though, seem to have survived the test of time. Each report starts with characterizations of new kinds of systems or families of systems to which AI re-

search and technology could contribute; they then move to specific discussions of research problems. Each also documents the solid AI research base on which subsequent research could build. In speaking of the future intelligent systems it describes, the twenty-first century report notes that “[u]seful, if limited members of these families should be possible within five years, although the full visions are at least one to two decades away.”

Our field has made significant progress in most of the areas described in these two reports, and, perhaps surprisingly (at least to this skeptic of crystal balls), we have developed important, interesting albeit limited versions of many of the systems envisioned. For instance, there are more sophisticated (“intelligent”) simulation systems; robot teams have not only played soccer, but also assisted in disaster recovery operations; and algorithms developed for machine learning and natural language processing are embedded in systems in everyday use by a broad spectrum of the world’s population. There remains much yet to be done, but the basis on which we can build is even more solid and richer. I was, however, struck by how little has been done in developing one family of intelligent systems, “intelligent project coaches” that participate with people in the design and operation of complex systems. As computer systems themselves become more complex, and their design and maintenance an ever more critical problem, building the capabilities for computer systems to participate intelligently in these endeavors seems even more important and the intellectual challenges of building such systems just as exciting. Thus, as I look to the future, I hope reflections of AAAI presidents ten or twenty-five years hence will report on developments in this area.

In reflecting on progress since my presidency, I was also struck by the maturation of various AI subfields that were emerging in the late 1980s and early 1990s. The machine learning, uncertainty in AI, and multiagent areas now have their own conferences, held annually; their own journals; and, in at least two cases, their own associations or boards. The challenges of bringing together researchers in disparate areas of AI, so that they can benefit from each others’ advances, is ever greater and more important.

After my 1994 AAAI presidential address (see *AI Magazine* 17(2): 67–85), a highly valued and respected senior colleague thanked me for making him optimistic again about the field. It is a comment I remember frequently as I ponder AI’s future. In the address, I highlighted the “solid base established by research in a number

of areas over the last decade,” and I tried, implicitly if not explicitly, to stress the importance of AI being open to diverse approaches to modeling intelligence, the value of links between our scientific and engineering endeavors, the need for people who focused on building systems to respect theories and for those developing theories to appreciate the challenges of building systems, and for us to collaborate with one another both in research and in supporting our field. I argued that a major lesson of the 1980s was that AI could not stand alone but that AI capabilities needed to be designed as parts of systems built collaboratively with those in other areas of computer science, but I also pointed out that there was much “the rest of computer science could benefit from knowing about what we know how to do.” In the last decade, AI has proved its importance not only to computer science and to the individual-oriented cognitive sciences, but also to those social sciences (for instance, economics) that model group behaviors.

In commenting on the workshops that led to the twenty-first century and NII reports, I noted that “Each workshop brought together people from across the spectrum of AI research and applications. The participants were asked to become familiar enough with work outside their own individual research interests to be able to explain it to funders and to justify funding of work on key problems. Participants worked for the common good, even if their individual cost was higher ... I was struck by the enthusiasm everyone exhibited, and by the perseverance with which they stuck to the task.... Given the current climate in Washington and the U.S. more generally for research funding, I expect AAAI will be called on to do more of this, and will in turn need to call on you, our members, to help.” Ten years later, this too remains true.

Barbara J. Grosz is Higgins Professor of Natural Sciences in the Division of Engineering and Applied Sciences and dean of science of the Radcliffe Institute for Advanced Study at Harvard University. Her seminal contributions include the development of some of the earliest and most influential computer dialogue systems, establishing the research field of computational modeling of discourse; pioneering research on models of collaboration; and the design of techniques for collaborative human-computer interface systems. She has been elected to the American Philosophical Society and the American Academy of Arts and Sciences, is a Fellow of the American Association for Artificial Intelligence, the ACM, and the American Association for the Advancement of Science, and has received the University of California, Berkeley Computer Science and Engineering Distinguished Alumna Award as well as distinguished service awards from both AAAI and IJCAI. She served as AAAI President from 1993–95.