

---

**Message from the  
Vice President,  
Services and Software,  
Research Division**

The quest to understand the nature of intelligence is one of the great intellectual challenges of humankind. Of all the great questions—the nature of life, the structure of matter, the origins of the universe—the study of intelligence may be the most difficult, and perhaps the most interesting.

For thousands of years, this question was the focus of philosophers. With the advent of computers, we gained some important new tools for understanding and modeling various aspects of intelligent behavior, and the field of artificial intelligence (AI) was born. For the first time, we could hope not only to understand the nature of human thought, but to use that understanding to produce artificial systems exhibiting intelligent behavior, at least to some limited degree.

The grand challenge is still out there, and no one can confidently predict when—or if—AI research will yield a complete answer. AI has made considerable progress in emulating some higher-level functions of the human intellect—puzzle-solving, chess, and so on—but we still cannot begin to match the common sense of a five-year-old child or the sensory-motor capabilities of a housefly.

Fortunately, we have not had to await ultimate success to gain great benefit from AI research. Over the years, AI has given rise to a number of extremely valuable technologies: rule-based expert systems, speech and handwriting recognition, search engines for natural-language documents, aids for machine translation, and friendlier user interfaces, just to name a few. It is likely that the demands of AI applications have placed demands on systems that have catalyzed significant advances.

A successful research program in AI must combine work on practical near-term projects having significant commercial benefit with a proper appreciation for the field's ambitious long-term agenda. Most of the researchers, most of the time, should work on projects that are likely to lead to some valuable, tangible results. This provides a steady stream of accomplishments that justify the whole enterprise, both for the sponsors and for the researchers themselves. But it is the long-term goals of AI that provide the greatest excitement and the radical new ideas. It is the dream of ultimate success that attracts some of the most brilliant and creative researchers to the field. So we must also provide some resources for pursuit of that longer-term dream. If we lose sight of this balance, the field will soon stagnate.

This special issue of the *IBM Systems Journal* reflects this mixture of goals in AI research. On the one hand, we have reports of exciting new AI-inspired applications in areas as diverse as document management, fault analysis, and the pricing of service contracts. On the other hand, we have reports on two recent IBM workshops that addressed the field's longer-term goals. It is an exciting time in AI research, both at IBM and in other research labs around the world. I believe that this special issue illustrates that point.



Alfred Spector