

Artificial Intelligence

Realizing the Ultimate Promises of Computing

David L. Waltz

Artificial intelligence (AI) is the key technology in many of today's novel applications, ranging from banking systems that detect attempted credit card fraud, to telephone systems that understand speech, to software systems that notice when you're having problems and offer appropriate advice. These technologies would not exist today without the sustained federal support of fundamental AI research over the past three decades.

Bringing Common Sense, Expert Knowledge, and Superhuman Reasoning to Computers

Although there are some fairly pure applications of AI—such as industrial robots, or the INTELLIPATH (available at zelda.thomson.com/chaphall/medelec.html) pathology diagnosis system recently approved by the American Medical Association and deployed in hundreds of hospitals worldwide—for the most part, AI does not produce stand-alone systems, but instead adds knowledge and reasoning to existing applications, databases, and environments, to make them friendlier, smarter, and more sensitive to user behavior and changes in their environments. The AI portion of an application (e.g., a logical inference or learning module) is generally a large system, dependent on a substantial infrastructure. Industrial R&D, with its relatively short time-horizons, could not have justified work of the type and scale that has been required to build the foundation for the civilian and military successes that AI enjoys today. And beyond the myriad of currently deployed applications, ongoing efforts that draw upon these decades of federally-sponsored fundamental research point towards

even more impressive future capabilities:

Autonomous vehicles: A DARPA-funded onboard computer system (www.cs.cmu.edu/~pomerlea/nhaa.html) from Carnegie Mellon University drove a van all but 52 of the 2849 miles from Washington, DC to San Diego, averaging 63 miles per hour day and night, rain or shine.

Computer chess: Deep Blue (www.research.ibm.com/research/systems.html#chess), a chess computer built by IBM researchers, defeated world champion Gary Kasparov in a landmark performance.

Mathematical theorem proving: A computer system at Argonne National Laboratories (www.mcs.anl.gov/home/mccune/ar/robbins/) proved a long-standing mathematical conjecture about algebra using a method that would be considered creative if done by humans.

Scientific classification: A NASA system learned to classify very faint signals as either stars or galaxies with superhuman accuracy, by studying examples classified by experts.

Advanced user interfaces: PEGASUS (sls-www.lcs.mit.edu/PEGASUS.html) is a spoken language interface connected to the American Airlines EAASY SABRE reservation system, which allows subscribers to obtain flight information and make flight reservations via a large, on-line, dynamic database, accessed through their personal computer over the telephone.

In a 1977 article, the late AI pioneer Allen Newell foresaw a time when the entire man-made world would be permeated by systems that cushioned us from dangers and increased our abilities: smart vehicles, roads, bridges, homes, offices, appliances, even clothes. Systems built around AI components will increasingly monitor financial transactions, predict physical phenomena and economic trends,

This article was written for inclusion in the booklet “Computing Research: A National Investment for Leadership in the 21st Century,” available from the Computing Research Association, cra.org/research.impact.

control regional transportation systems, and plan military and industrial operations. Basic research on common sense reasoning, representing knowledge, perception, learning, and planning is advancing rapidly, and will lead to smarter versions of current applications and to entirely new applications. As computers become ever cheaper, smaller, and more powerful, AI capabilities will spread into nearly all industrial, governmental, and consumer applications.

Moreover, AI has a long history of producing valuable spin-off technologies. AI researchers tend to look very far ahead, crafting powerful tools to help achieve the daunting tasks of building intelligent systems. Laboratories whose focus was AI first conceived and demonstrated such well-known technologies as the mouse, time-sharing, high-level symbolic programming languages (Lisp, Prolog, SCHEME), computer graphics, the graphical user interface (GUI), computer games, the laser printer, object-oriented programming, the personal computer, e-mail, hypertext, symbolic mathematics systems (MACSYMA, MATHEMATICA, MAPLE, DERIVE), and, most recently, the software agents which are now popular on the World Wide Web. There is every reason to believe that AI will continue to produce such spin-off technologies.

Intellectually, AI depends on a broad intercourse with computing disciplines and with fields outside computer science, including logic, psychology, linguistics, philosophy, neuroscience, mechanical engineering, statistics, economics, and control theory, among others. This breadth has been necessitated by the grandness of the dual challenges facing AI: creating mechanical intelligence and understanding the information basis of its human counterpart. AI problems are extremely difficult, far more difficult than was imagined when the field was founded. However, as much as AI has

borrowed from many fields, it has returned the favor: through its interdisciplinary relationships, AI functions as a channel of ideas between computing and other fields, ideas that have profoundly changed those fields. For example, basic notions of computation such as memory and computational complexity play a critical role in cognitive psychology, and AI theories of knowledge representation and search have reshaped portions of philosophy, linguistics, mechanical engineering, and control theory.

Historical Perspective

Early work in AI focused on using cognitive and biological models to simulate and explain human information processing skills, on “logical” systems that perform common-sense and expert reasoning, and on robots that perceive and interact with their environment. This early work was spurred by visionary funding from the Defense Advanced Research Projects Agency (DARPA) and Office of Naval Research (ONR), which began on a large scale in the early 1960s and continues to this day. Basic AI research support from DARPA and ONR—as well as support from NSF, NIH, AFOSR, NASA, and the U.S. Army beginning in the 1970s—led to theoretical advances and to practical technologies for solving military, scientific, medical, and industrial information processing problems.

By the early 1980s an “expert systems” industry had emerged, and Japan and Europe dramatically increased their funding of AI research. In some cases, early expert systems success led to inflated claims and unrealistic expectations: while the technology produced many highly effective systems, it proved very difficult to identify and encode the necessary expertise. The field did not grow as rapidly as investors had been led to expect, and this translated into some temporary disillusionment. AI researchers responded by developing new technologies, including streamlined methods for eliciting expert knowledge, automatic methods for learning and refining knowledge, and common sense knowledge to cover the gaps in expert information. These technologies have given rise to a new generation of expert systems that are easier to develop, maintain, and adapt to changing needs.

Today developers can build systems that meet the advanced information processing needs of government and industry by choosing from a broad palette of mature technologies. Sophisticated methods for reasoning about uncertainty and for coping with incom-

plete knowledge have led to more robust diagnostic and planning systems. Hybrid technologies that combine symbolic representations of knowledge with more quantitative representations inspired by biological information processing systems have resulted in more flexible, human-like behavior. AI ideas also have been adopted by other computer scientists—for example, “data mining,” which combines ideas from databases, AI learning, and statistics to yield systems that find interesting patterns in large databases, given only very broad guidelines.

Case Studies

The following four case studies highlight application areas where AI technology is having a strong impact on industry and everyday life.

Authorizing Financial Transactions

Credit card providers, telephone companies, mortgage lenders, banks, and the U.S. Government employ AI systems to detect fraud and expedite financial transactions, with daily transaction volumes in the billions. These systems first use learning algorithms to construct profiles of customer usage patterns, and then use the resulting profiles to detect unusual patterns and take the appropriate action (e.g., disable the credit card). Such automated oversight of financial transactions is an important component in achieving a viable basis for electronic commerce.

Configuring Hardware and Software

AI systems configure custom computer, communications, and manufacturing systems, guaranteeing the purchaser maximum efficiency and minimum setup time, while providing the seller with superhuman expertise in tracking the rapid technological evolution of system components and specifications. These systems detect order incompletenesses and inconsistencies, employing large bodies of knowledge that describe the complex interactions of system components. Systems currently deployed process billions of dollars of orders annually; the estimated value of the market leader in this area is over a billion dollars.

Diagnosing and Treating Problems

Systems that diagnose and treat problems—whether illnesses in people or problems in hardware and software—are now in widespread use. Diagnostic systems based on AI technology are being built into photocopiers, computer operating systems, and office automation tools to reduce service calls. Stand-

alone units are being used to monitor and control operations in factories and office buildings. AI-based systems assist physicians in many kinds of medical diagnosis, in prescribing treatments, and in monitoring patient responses. Microsoft's OFFICE ASSISTANT (207.68.137.59/office/office97/documents/o97whnew/WNew3.htm), an integral part of every Office 97 application, provides users with customized help by means of decision-theoretic reasoning.

Scheduling for Manufacturing

The use of automatic scheduling for manufacturing operations is exploding as manufacturers realize that remaining competitive demands an ever more efficient use of resources. This AI technology—supporting rapid rescheduling up and down the “supply chain” to respond to changing orders, changing markets, and unexpected events—has shown itself superior to less adaptable systems based on older technology. This same technology has proven highly effective in other commercial tasks, including job shop scheduling, and assigning airport gates and railway crews. It also has proven highly effective in military settings—DARPA reported that an AI-based logistics planning tool, DART, pressed into service for operations Desert Shield and Desert Storm, completely repaid its three decades of investment in AI research.

The Future

AI began as an attempt to answer some of the most fundamental questions about human existence by understanding the nature of intelligence, but it has grown into a scientific and technological field affecting many aspects of commerce and society.

Even as AI technology becomes integrated into the fabric of everyday life, AI researchers remain focused on the grand challenges of automating intelligence. Work is progressing on developing systems that converse in natural language, that perceive and respond to their surroundings, and that encode and provide useful access to all of human knowledge and expertise. The pursuit of the ultimate goals of AI—the design of intelligent artifacts; understanding of human intelligence; abstract understanding of intelligence (possibly superhuman)—continues to have practical consequences in the form of new industries, enhanced functionality for existing systems, increased productivity in general, and improvements in the quality of life. But the ultimate promises of AI are still decades away, and the necessary advances in knowledge and

AI researchers tend to look very far ahead, crafting powerful tools to help achieve the daunting tasks of building intelligent systems.

Universities Space Research Association

The Universities Space Research Association (USRA) is conducting a search for the director of the Research Institute for Advanced Computer Sciences (RIACS). The institute is operated by USRA for the National Aeronautics and Space Administration (NASA).

RIACS is located in the heart of Silicon Valley at NASA's Ames Research Center and nearby the outstanding universities in the San Francisco Bay area. The Ames Research Center has recently been designated NASA's Center of Excellence in Information Technology. In this capacity, Ames has the responsibility to build an information technology research program that is pre-eminent within the Agency.

RIACS is chartered to carry out research and development in computer science. This work is devoted in the main to tasks that are strategically enabling with respect to NASA's missions in space exploration and aeronautics. The current foci for this work are: high-performance computing, cognitive and perceptual enhancements (computational aids designed to enhance these human abilities), and autonomous systems.

An additional effort of RIACS is to broaden the base of researchers in these areas of importance to the nation's aeronautics and space enterprises. In this connection, RIACS works to foster collaborative links between scientists at Ames and RIACS' staff and visitors. In particular, through its visiting-scientist program, RIACS facilitates the participation of university-based researchers, in the U.S. and abroad, in this R&D.

RIACS' researchers have access to the excellent

computational facilities of the Ames Research Center.

Over the next five years, RIACS have the opportunity and the challenge simultaneously to increase substantially its staff, expand the scope of its research and development effort in these important areas of IT, and extend the range of its cooperative undertakings with researchers on the campuses and in industry.

We are seeking an individual to lead the institute in meeting this challenge.

Salary and benefits are competitive and will depend in part upon the qualifications of the individual to be selected for the position. The ideal applicant would hold a Ph.D. in computer science and have a substantial record of productive research as well as experience in the management and administration of research.

Candidates are encouraged to submit their applications by 15 December, but we will continue to accept applications until the position is filled. Please send a complete resume and a listing of three references, by mail to:

Dr. W. D. Cummings, Executive Director
Universities Space Research Association
Suite 212
10227 Wincopin Circle
Columbia, MD 21044
or, by e-mail to: cummings@hq.usra.edu

The Universities Space Research Association is an Equal Opportunity Employer.

technology will require a sustained fundamental research effort.



David Waltz is Vice President of the Computer Science Research Division of NEC Research Institute in Princeton, NJ, and Adjunct Professor of Computer Science at Brandeis University in Waltham, MA. From 1984 to 1993, he was Senior Scientist and Director of Advanced Information Systems at

Thinking Machines Corporation in Cambridge, MA, where he headed a group that built a large-scale text retrieval system, a web search engine, and data mining systems. At the same time he was Professor of

Computer Science at Brandeis. From 1974 to 1983 he was Professor of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign.

Waltz received the S.B., S.M., and Ph.D. degrees from MIT. His research, especially his work on constraint propagation and memory-based reasoning, has helped spawn active R&D fields.

Waltz is President of the American Association of Artificial Intelligence and was elected a Fellow of AAAI in 1990. He was President of the Association for Computing Machinery (ACM) Special Interest Group on Artificial Intelligence from 1977-79, and served as Executive Editor of *Cognitive Science* from 1983-86 and as AI Editor for *Communications of the ACM* from 1981-84. He is a senior member of IEEE.