

Strategic Directions in Artificial Intelligence¹

Jon Doyle and Thomas Dean

■ This report, written for the general computing and scientific audience and for students and others interested in artificial intelligence, summarizes the major directions in artificial intelligence research, sets them in context relative to other areas of computing research, and gives a glimpse of the vision, depth, research partnerships, successes, and excitement of the field.

What Is Artificial Intelligence?

The field of artificial intelligence (AI) consists of long-standing intellectual and technological efforts addressing several interrelated scientific and practical aims:

- constructing intelligent machines, whether or not these operate in the same way as people do;
- formalizing knowledge and mechanizing reasoning, both common sense and refined expertise, in all areas of human endeavor;
- using computational models to understand the psychology and behavior of people, animals, and artificial agents; and
- making working with computers as easy and as helpful as working with skilled, cooperative, and possibly expert people.

Even considering only the first two of these aims, AI has perhaps the broadest concerns of all areas of computing research, covering investigations ranging from the natural sciences to the social, from engineering to economics, physics to psychology. Its very nature forces AI to grapple with the complexity of the natural world as well as that of the artificial systems

that form the subject matter of computing studies generally. Its main theoretical questions stand as peers of the deepest and most important questions in any scientific field, and its practical impact on living promises to equal that of any known technology.

The aims of AI reflect ancient dreams of using minds and hands to create beings like ourselves. In centuries past, pursuit of these dreams gave rise to both mechanical automata and formal theories of reasoning, eventually yielding the spectacularly successful modern artificial computers that, in calculating and computing, replicate and surpass abilities that people of earlier times regarded as intellectual activities on a par with writing letters and playing good chess. Using these computers over the past four decades, modern AI has built on the best thinking in a number of areas—especially computer systems, logic, the mathematical theory of computation, psychology, economics, control theory, and mathematical problem solving—to construct concrete realizations of devices that

- solve intellectual problems both theoretical and practical, common and esoteric;
- control robot motions through planning, sight, touch, and self-awareness;
- interpret human language, both written and spoken; and
- learn new skills and knowledge through instruction, from experience, and by analyzing other data.

Some of these realizations have proven highly successful, others rudimentary and incomplete, but each captures recognizable and significant elements of familiar human capa-

This article is reprinted with permission from *ACM Computing Surveys* 28(4), December 1996. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. See footnotes 2 and 3 for portions based on text copyright © 1995 American Association for Artificial Intelligence, and reprinted with permission. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept, ACM Inc., fax +1 (212) 869-0481, or permissions@acm.org. © 1996 ACM 0360-0300/96/1200-0001 \$03.50.

bilities, and provides a skeleton upon which future research may enlarge.

One can divide present-day AI research into the following primary (and overlapping) areas ²:

1. *knowledge representation and articulation* seeks to discover expressive and efficient forms and methods for representing information about all aspects of the world and to use these methods to create and compile explicit, formal, multipurpose catalogs of substantive knowledge;
2. *learning and adaptation*, which extends statistical, analytical, and scientific discovery techniques and hypothesized neurophysiological mechanisms to procedures that extract a wide range of general trends, facts, and techniques from instruction, experience, and collected data;
3. *deliberation, planning, and acting*, which concerns methods for making decisions, constructing plans or designs to achieve specified goals, and monitoring, interpreting, diagnosing, and correcting the performance of the plans and implementations of the designs;
4. *speech and language processing*, which seeks to create systems capable of communicating in and translating among natural written and spoken languages;
5. *image understanding and synthesis*, which develops algorithms for analyzing photographs, diagrams, and real-time video image streams as well as techniques for the customized presentation of quantitative and structured information;
6. *manipulation and locomotion*, which seeks to replicate and surpass the abilities of natural hands, arms, feet, and bodies;
7. *autonomous agents and robots*, which integrates the other areas to create robust, active entities capable of independent, intelligent, real-time interactions with an environment over an extended period;
8. *multiagent systems*, which identifies the knowledge, representations, and procedures needed by agents to work together or around each other;
9. *cognitive modeling*, which focuses on contributing techniques and constructing integrated architectures that replicate structural or behavioral features of human cognition; and
10. *mathematical foundations*, which takes the concepts and techniques of the other areas as subjects for formalization, distillation, analysis, and reconceptualization.

The aims of AI reflect ancient dreams of using minds and hands to create beings like ourselves.

These areas form but one categorization of work in the field; other such lists have been suggested in the past, and categorizations will again change as the field solves problems and identifies new ones. Consequently, no such list constitutes a definition of AI. For that matter, neither do other common misapprehensions, such as defining AI in terms of the activities of a particular group of people, or the use of a particular set of methods or programming languages, or as a label for computing research that does not come under some other heading. The first two of these mistake accident for essence, and the last incorrectly divides the computing field. AI is not a list of areas or a methodology, even less a group or a catch-all label, but consists of work on the enduring and intriguing aims of understanding intelligent beings and constructing intelligent systems.

Problems of Formulation

AI has won its successes only with great effort. In earlier times, researchers used informal means to specify the problems under investigation, and their work revealed the great difficulty of formulating these problems in precise terms. Solving these problems of formulation (Minsky 1962) required considerable experimentation with and exploration of alternative conceptualizations in order to find appropriate ways of making them amenable to technical investigation and solution. Although logic, game theory, and other disciplines contributed formal approaches to specifying these problems, their methods often missed the mark in essential ways, especially by begging AI's question through presuming too much reasoning power and coherence on the part of the agent. In coming to new formulations, AI has often advanced these other fields, providing the first precise means for addressing problems shared with them.

Researchers in AI have traditionally met problems of formulation joyfully, courageously, and proudly, accepting the severe risks ensuing from such exploratory work in pursuit of the proportionately large gains that can result from finding successful formulations. The willingness to cultivate problems lacking ready formalizations has also engendered some disrespect for AI, as observers focus on the failures rather than on the successes. However, this adventurousness has proven highly fruitful, creating whole new subfields for formal investigation. Though some important problems still lack adequate formalizations, for many others AI has successfully

provided formal foundations supporting rich areas of technical investigation.

AI has undergone a sea-change in the general character of its research methodology since about 1980, partly through progress on its problems of formulation, and partly through increasing integration with related areas of computing research and other fields. Speculative, exploratory work remains necessary in investigations of many difficult issues. In particular, the natural or useful scope of the formalized knowledge employed in an investigation does not always admit simple, formally satisfying characterizations, so the field retains an element of conceptual exploration. The more typical research effort today, however, relies on formal, theoretically precise, and experimentally sophisticated methods for investigation and technical communication. Rigorous science, engineering and mathematics now overshadow other work in much of the literature. Recent AI also replaces the focus of the early analytical studies on using isolated “toy” domains with a focus on using realistically broad and large-scale problem domains, and concentrates much more on integrating its ideas, systems, and techniques into standard computing theory and practice. These changes not only complement the increase in precision and formality, but demand additional rigor in order to enforce the conventions and coherence necessary in scaling up and integrating systems.

Accompanying this change in the character of AI results and research, accepted methods of educating students in AI have changed to recognize many prerequisites sometimes overlooked in years past. To understand the literature and make good in their own work, modern AI students must learn basics of a number of fields: logic, statistics, decision theory, stochastic processes, analysis of algorithms, complexity theory, concurrency, and computational geometry, to name but a few.

Contributions

Some highlights of the major contributions of AI to computing, and to science more generally, include artificial neural networks, automated deduction, autonomous and semi-autonomous mobile robots, computational qualitative reasoning (about physical systems), constraint programming, data-mining systems, decision-tree learning methods, description logics (structured declarative representations going beyond those structures common in traditional logic), design and configuration systems, evolutionary compu-

tation, expert or knowledge-based systems (based on corpora of explicit mainly declarative knowledge), fuzzy logic and control systems, graphical representations of uncertain information (Bayesian belief networks and others), heuristic search, logic and rule-based programming systems, mechanized symbolic mathematical calculation, natural language understanding and generation systems, non-monotonic logics (a new category of logic formalizing assumption making), planning and scheduling systems, program synthesis and verification methods, real-time speaker-independent speech understanding, reason or truth maintenance systems (systematic recording and reuse of reasoning steps), robotic assembly systems, text processing and retrieval systems, and visual classification and registration systems.

One can appreciate the intellectual productivity of AI through the subjects it launched or has helped launch as independent areas of research, including artificial neural networks, automated deduction, constraint programming, heuristic search, integrated software development environments, logic programming, object-oriented programming, mechanized symbolic mathematical calculation, and program synthesis and verification methods. One should also note the major contributions AI has made to symbolic computing and functional programming. Both have been stimulated in fundamental ways through the sustained development and use of LISP and its relatives in AI research. AI has made important contributions to computational linguistics, to the area of epistemic logics (especially through nonmonotonic logics, theories of belief revision, and the computational applications now also heavily used in the theory of distributed systems), and to economics and operations research (where AI methods of heuristic search, especially stochastic heuristic search, have caused something of a revolution). AI has also served computing research as a prime exporter to other scientific fields of the notion of studying processes in their own right. AI models of process and information processing in language, reasoning, and representation have caused major shifts in linguistics, psychology, philosophy, and organization theory (e.g., with rule-based systems and artificial neural networks providing a “rehabilitation” of the impoverished and formerly stagnating behavioristic approach to psychology), and AI models now figure prominently in each of these fields. In addition to changing scientific fields, some AI methodologies (especially expert knowledge-based systems,

AI has undergone a sea-change in the general character of its research methodology since about 1980, partly through progress on its problems of formulation, and partly through increasing integration with related areas of computing research and other fields.

One need not hazard such risks, however, to identify the core challenges facing the next generation of AI systems, namely exhibiting robust operation in hostile environments, broad and deep knowledge of large domains, the ability to interact naturally with people, and a degree of self-understanding and internal integrity.

artificial neural networks, and fuzzy systems) have changed the perspective of many engineers, who now go beyond the traditional concerns of algorithms and data to capture the knowledge or expertise underlying desired functionalities.

The manifold practical applications of AI continue to expand every year. The following few examples give the flavor of current successes, but one may find many more in the proceedings of the annual AAAI conference on Innovative Applications of Artificial Intelligence, in recent issues of Communications of the ACM (for example, the November 1995 and January, February, April, May, and August 1996 issues), and in other sources in the literature. Probabilistic diagnostic systems, based on graphical uncertainty representations, form a large class of successful applications, including the Intellipath pathology diagnosis system approved by the American Medical Association (Heckerman 1991), the VISTA monitoring and analysis system used by NASA for space shuttle mission control (Horvitz et al. 1992), and even the printer diagnosis and “wizard” subsystems of Microsoft software (Heckerman, Breese, and Rommelse 1995). Artificial neural networks also appear in many successful systems, from automated PAP smear diagnosis to online handwriting recognition (Lyon and Yaeger 1996) and vehicle navigation (Jochem and Pomerleau 1996). Fuzzy logic systems have been applied to many problems including camera and appliance control. Design and configuration systems form a large class in everyday use, with the largest, such as AT&T’s PROSE and QUESTAR systems, processing orders worth billions of dollars (Wright et al. 1993). Expert knowledge-based systems abound, with applications from credit authorization and detection of money laundering (Senator et al. 1995) to highly skilled simulations of helicopter pilots (Tambe et al. 1995) and great numbers of knowledgeable help-desk (customer service) systems. The automatically-synthesized KTS (Kestrel Transportation Scheduler) software has proven startlingly efficient in large-scale scheduling applications (Smith, Parra, and Westfold 1996), and knowledge-based planning and scheduling systems now yield dramatic improvements in manufacturing efficiency and productivity (Naj 1996). Speech-understanding technology has begun to have commercial impact, from control aids for the manually impaired to replacing telephone operators. Machine-vision systems now find routine use in industrial inspection and assembly processes and play increasingly

important roles in the analysis of medical imagery, from the analysis of radiographs to helping surgeons operate on the correct areas of the brain (Grimson et al. 1996). Clinical trials now in progress seek to evaluate a wide range of computer-aided surgical procedures, including the use of surgical robotic devices in hip replacement surgery. Applications in automatic vehicle control have only reached the demonstration stage (CMU’s RALPH vehicle drove across the continental United States with minimal human intervention (Jochem and Pomerleau 1996)), but promise more widespread applications in the near future. Machine-learning methods have successfully automated the analysis of astronomical data and found new classifications for astronomical objects (Goebel et al. 1989). To these one must add the “impractical” but impressive success of game-playing systems, which through systematic exploitation of AI search techniques and special-purpose hardware now hold the position of the world’s best checkers player (Schaeffer et al. 1992), have tied the world-champion backgammon player (Tesauro 1995), and seriously challenged the world chess champion (Kasparov 1996).

Directions

Predicting the results of the next generation of fundamental research requires either bravery or foolishness. One need not hazard such risks, however, to identify the core challenges facing the next generation of AI systems, namely exhibiting robust operation in hostile environments, broad and deep knowledge of large domains, the ability to interact naturally with people, and a degree of self-understanding and internal integrity.

- The world, although not necessarily malicious, provides a messy and hostile environment with a structure that cannot be completely anticipated. Whether a robot moving through the physical world, a “softbot” wandering through cyberspace, or an intelligent collaborative or analytical agent, the system confronts problems, opportunities, situations, and external (user or collaborator) demands that change over time. The system will surely fail in its purpose unless it can quickly learn from its experience and adapt to changes in its environment.
- The system will require large amounts of knowledge reflecting a broad cross-section of the world. It must capture or embody this knowledge in a form suitable for use in many different tasks, for tailoring portions

to achieve high performance on specific tasks, and for acquiring knowledge through its own problem-solving experience and through collaboration with people and other agents. Exploiting large bodies of poorly structured information, such as the World Wide Web and older information systems, will demand efficiency and dexterity in handling extremely large amounts of complex information.

- To interact with its human collaborators, the system must engage in extended dialogues that progressively clarify and enrich the depth of understanding. Doing this requires using natural languages and appropriate visual displays and tactile modalities to communicate, rather than simply exchanging single correct sentences as demanded by today's limited systems. Interacting with its artificial collaborators calls for similar economy and clarity and need not, but may, involve the same languages as used for communication with humans.
- The system also must understand itself as well as its collaborators and the world around it, the better to integrate its diverse components, facilitate their smooth interaction, maintain and improve its knowledge and skills, and dynamically regulate the use of its scarce resources.

Building systems with these characteristics poses the same challenges that have driven AI research throughout its history, and each of the areas of technical investigation introduced earlier—knowledge representation and articulation, learning and adaptation, deliberation, planning, and acting, speech and language processing, image understanding and synthesis, manipulation and locomotion, autonomous agents and robots, multiagent systems, cognitive modeling, and mathematical foundations—supports a vigorous research effort contributing to meeting these challenges. This brief survey cannot present a complete picture of all the important directions of research in each of these areas (see Weld, Marks, and Bobrow 1995) for a more generous, though still abbreviated, summary, the challenge problems listed by (Selman et al. 1996), and the 1994 Turing Award lectures of Feigenbaum (1996) and Reddy (1996) for other perspectives). Instead, the following sections sketch some broad directions—spanning many of the narrower areas of investigation—that characterize much of the work in the field today and, in all probability, for the next decade. These directions consist of pursuing systemic and intellectual integration, of which building robots (both physical and computational) and

modeling rationality (mainly in the sense of economics and decision theory) form two broad special cases; supporting collaboration; enhancing communication; obtaining the broad reaches of knowledge needed for intelligent action; and deepening the mathematical foundations of the field. The robot direction also constitutes a technical area, but a broad one that touches on most of the other areas as well. The remaining directions each make essential use of several of the technical areas of investigation.

Pursuing Integration

AI today vigorously pursues integration along several dimensions: integrating systems that support different capabilities, combining theories and methodologies that concern different facets of intelligence, coordinating subfields within AI, and reconciling, accommodating, and exploiting ideas from other disciplines.

Making progress on hard problems requires analysis, and AI has made substantial progress by isolating and understanding many of the important subtasks and subsystems of intelligent behavior in terms of knowledge representation, learning, planning, vision, and like subjects. Much current research seeks to put the pieces back together by constructing integrated systems that incorporate major capabilities drawn from several or all of these areas. For example, natural language processing systems now incorporate learning techniques, recent planning systems incorporate methods for reasoning under uncertainty, and “active” vision systems combine planning control of robot motions with analysis of the resulting sensor data. Integration offers a special opportunity both to test the component theories and also to constrain further the requirements on them. Integration takes special prominence in work on building robots and supporting collaboration, detailed in the following, and in work on complete cognitive architectures, such as SOAR (Rosenbloom, Laird, and Newell 1993).

Apart from the engineering challenge of building complex, hybrid systems capable of accomplishing a wide range and mixture of tasks, AI's scientific challenge consists of providing integrated computational theories that accommodate the wide range of intellectual capabilities attributed to humans and assumed necessary for nonhuman intelligences. Many efforts at theoretical integration occur among the subfields of AI. Common logical underpinnings help integrate theories of knowledge representation, planning, problem

... research on control of autonomous robots has shifted toward less detailed representations that make simpler demands on sensory and actuation systems.

solving, reasoning, and some aspects of natural language processing, while economic concepts of rationality and the mathematics of Markov decision processes help unify recent theories of probabilistic planning, fault diagnosis and repair, reinforcement learning, robot control, and aspects of speech recognition and image processing. Of necessity, many of these efforts at theoretical integration cross disciplinary boundaries and lead to integration with other fields. AI has drawn on and contributed to logic, philosophy, psychology, and linguistics for some time. Integration with economics, decision theory, control theory, and operations research has served as a focus for more recent efforts, detailed in the section on rationality.

The most novel case, but perhaps of the greatest immediate practical importance, consists of integration with related areas of computing research and practice. Integration with these areas has progressed steadily, but slower than one might hope; the areas of tightest integration include theory, databases, and programming languages (especially for logic and object-oriented programming). No one in AI today views AI systems as standing alone; instead, most view AI techniques as supplying components of complex computer systems, components that provide key elements of the capabilities, flexibility, and cooperativeness of an overall system. To realize their benefits fully, AI techniques and the theories underlying them must be integrated much more completely into the warp and woof of computing theory and practice. Representative long-term goals for integration with related areas of computing research include:

- dramatically change the nature of programming so that, for most tasks and applications, the programmer works with a collaborative, knowledge-based agent that provides direct and extensive support for specifying, designing, implementing, maintaining, and re-engineering reliable, robust, secure, efficient, and intelligible hardware and software systems.
- remove artificial distinctions between knowledge bases and databases, which current systems treat separately (the knowledge bases in the running program, the comparatively large databases in specialized servers). Construct efficient, uniformly transparent mechanisms for representing large amounts of knowledge and data, for translating among these representations, and for applying knowledge-based inference, learning, and discovery mechanisms to information appearing in a variety of

forms in extremely large-scale knowledge and data repositories.

- better integrate reasoning systems with traditional programs, making it easy to mix computation and reasoning from knowledge to achieve desired results. Integrate the corresponding programming and specification tools as well, removing artificial distinctions between description logics and object-oriented systems and between logic programming, rule-based programming, and traditional programming systems.
- lessen the tension between speed and quality of action by continuing adaptation and extension of knowledge-based reasoning and learning techniques to real-time operation and control of complex real-world systems that involve hard deadlines.
- make computers easier to use: more cooperative and customizable, with interfaces that employ natural languages and other modalities to communicate in familiar and convenient ways.

Building Robots (Physical and Computational)

Building integrated agents that perceive and act in extant complex and dynamic environments requires integrating a wide range of subfields of AI and computing research. These environments include both physical environments and the “virtual” worlds of information systems. By focusing on everyday worlds of interest to people, such as office buildings or the Internet, researchers avoid the methodological hazards of designing and simulating toy worlds unwittingly tailored to the designs they were supposed to validate. They also avoid the opposite problem of focusing on problems so hard even humans cannot solve them.

The term “robot” traditionally refers to automated agents acting in physical environments, with terms like “softbot” and “software agent” introduced to refer to agents acting purely within information systems, but this distinction promises to fade in importance as physical agents enter into electronic communication with each other and with online information sources, and as informational agents exploit perceptual and motor mechanisms (such as interpretation of graphical images and synthesis of gestures and other animations). Accordingly, this report calls both types of agents robots, returning to the original sense of the word as an artificial worker in Karel Čapek’s 1921 play *R.U.R.* (Rossum’s Universal Robots).

Many of the major areas of AI and comput-

ing research play essential roles in work on robots, from planning, sensing and learning to high-performance numerical computing and interacting with multiple databases across networks. Robots working in informational environments require little investment in additional expensive or unreliable robotic hardware, since existing computer systems and networks provide their sensors and effectors. Robots with physical abilities, in contrast, require mechanization of various physical sensory abilities, including vision, hearing, touch, taste, smell, thermoreceptivity, and mechanization of various physical motor abilities, including manipulation and locomotion. These areas comprise some of the major efforts of AI and provide some of its most impressive successes.

Recent work points toward new directions and applications in physical perception and motor abilities. Maturing work on vision as inverse graphics now finds applications in medicine and industry, while research on vision for autonomous robots now takes as its focus less well understood approaches employing more qualitative and “purposive” analyses that select which portions or aspects of images to look at based on what the robot is trying to do. Work on motor abilities now yields unexpected applications in rational drug design for traditional techniques like configuration-space planning, while research on control of autonomous robots has shifted toward less detailed representations that make simpler demands on sensory and actuation systems. Other work actively seeks to transfer the new representation techniques to applications such as industrial cleaning and ordnance disposal.

Scaling the operation of autonomous robots to more complicated tasks, and to natural environments in which the robots operate safely in the presence of humans, requires further integration of perception, action, and reasoning. High-level reasoning about what to do requires developing new perceptual systems that generate the kinds of data needed by the reasoning system, but the reasoning system in turn must make realistic demands on perception. The marriage of these abilities aims to produce robots that combine the high-level programmability of traditional AI systems with the fault tolerance of current autonomous robots.

The area of computer vision exhibits increasing integration with other disciplines. The subfield of active vision, for example, seeks to radically simplify the process of information extraction by closely coupling it to the

control of action for a particular task, thus exploiting the practical constraints imposed by the domain of operation. Other approaches exploit theoretical and technological integration. For example, inverse optics—roughly, the use of images to build models like those used in computer-aided design systems—now draws on collaborations with computer graphics, medical image processing, computational geometry, and multimedia.

Representative long-term goals in this direction include building robots that:

- combine planning, learning, vision, touch, speech and other senses in performing everyday tasks, e.g., housecleaning, cooking, shopping, answering the telephone, making appointments, and negotiating or bargaining with other agents (human or otherwise) for commodities and information.
- adaptively monitor, select, tailor, and rewrite the contents of electronic information sources (TV, faxes, newswires, the World Wide Web) to inform one of news and events in accord with one’s changing personal interests, plans, and purposes.
- record, monitor, and analyze one’s medical history and condition over one’s entire lifetime, helping to explain and maintain treatment plans, to detect physician mistakes, and to guide interactions with healthcare providers.
- perform tasks people cannot or do not want to do, such as mining, firefighting, handling hazardous material, and planetary exploration.
- operate within large-scale distributed systems to monitor and maintain the overall system operation, learning how to detect and defend against malicious external (criminal or terrorist) or internal (disgruntled or corrupt employee) attacks.

Modeling Rationality

Formal and informal notions of rationality from psychology (reasoning and argument) and logic (semantic consistency, deductive closure) have served AI well from its earliest days. They supply concepts useful in mechanizing several forms of reasoning, and provide the basis for major cognitive-modeling explorations of hypotheses about the psychology of human ratiocination and its integration with other mental faculties. These large-scale, detailed cognitive theories have already begun to change the face of psychological theory, while nonmonotonic, probabilistic, and new modal logics continue to expand conceptions of logical rationality. The

... recent work in AI has drawn on economic theory in many ways while remaining cognizant of its limitations.

main new direction here, however, seeks integration of rationality in the logical and psychological senses with the economic sense of rationality (maximum utility, optimal allocation of resources). Rationality in the economic sense has made only sporadic appearances in AI until recently, even though it subsumes the logical sense from a formal point of view and provides explanations of important aspects of rationality in the psychological sense. Rationality in the economic sense offers many attractions as an organizing principle for both intelligent system construction and intellectual integration. It contributes to the system's coherence (in terms of explanation, justification, and verification), to its competence (offering performance advantages), and to its construction methodology (design and development advantages). Researchers in many areas of AI have recognized these advantages and begun work on exploiting rationality in the economic sense. In consequence, economic rationality promises to permeate much of AI; indeed, this work also promises to contribute to economics as well, as AI and economics work together on their shared problems.

Early work in AI largely rejected formal economic models in favor of psychological ones because the standard economic theory focuses on an idealization in which rational agents suffer no limitations of memory or time in coming to decisions, and which, for these reasons and others, may not be realizable in the world. Economic approaches generally presupposed possession of utility and probability functions over all contingencies, which did not help in AI's need to construct these functions at the outset. Moreover, economics formalized preference and probability information in terms of very abstract representations that, through a lack of much structure, supported only very inefficient algorithms for making rational choices. In contrast, the psychological problem-solving methodology quickly adopted in AI starts with an easily realizable notion of rationality that is much weaker than the standard economic notion (one sanctioned, moreover, by Herbert Simon, a heretical economist founder of AI). Rather than seeking to maximize the numerical utility or expected utility across all conceivable actions, problem-solving rationality simply seeks to find actions meeting less stringent aspirations, such as satisfying designated conditions ("goals") on the resulting states. Building on this approach, researchers now work towards ideal rationality through several means: by increasing the sophistication of rea-

soning about goals, by adopting explicit notions of utility, and by performing tractable optimizations that take into account the limited knowledge and abilities of the decision maker.

As this approach to rationality suggests, recent work in AI has drawn on economic theory in many ways while remaining cognizant of its limitations. The first and major exploitation came about through partially solving the problem of representing probabilistic information that stymied early attempts to use decision-theoretic ideas directly. The popular graphical formalisms, especially Bayesian networks and influence diagrams, now support great numbers of successful applications, from sophisticated medical reasoners to mundane printer-diagnostic subsystems of personal computer operating systems. Indeed, the decision-theoretic notions of preference, utility, and expected utility now play important roles in many areas of AI research, as they help to shape learning and adaptation, to guide the plans and actions of autonomous agents and robots, and to reconcile and integrate AI planning methods with those of operations research. As interest in collaboration and multiagent systems has increased, many AI researchers have adopted the tools of game theory and the theory of social choice to analyze and design agent interaction protocols, to understand computational decision-making methods, and to analyze functional decompositions of mental organization. In the most explicit borrowing from economics, some work employs computational market price systems to allocate resources in a decentralized manner, and uses theoretical analyses of different economic systems to tailor multiagent organizations to achieve high efficiency in performing specific tasks.

Just as AI has contributed to logic, the intellectual trade with economics flows both ways, though unequally at present. Bayesian networks and other AI methods have improved practice in statistics. The anticipated but as yet unrealized prize contribution, however, lies in using the precise, detailed models of mental organization developed in AI in formulating a realistic and useful theory of the rationality of limited agents (such as people) and organizations composed of such agents, something that has evaded economics throughout its history. The AI theories relating goals and preferences provide one step in this direction, as they augment the traditional economic theories of preference with new qualitative languages for modeling the incomplete and conflicting desires of agents.

Recent work on control of deliberation, balancing the costs of further deliberation against the expected benefits, also points in this direction. More immediately, AI and computing research might help economists get a handle on costs and value of information, computation and communication, factors too often neglected in economics.

Representative long-term goals in this direction include:

- continued development of efficient representations and algorithms for rational decision and action that integrate, extend, and improve on current structured representations for probabilities, preferences, decisions, and (game-theoretic) games.
- construction of effective computational techniques that allow trading amounts of computational commodities—such as time, memory, or information—for gains in the value of computed results.
- understanding the relationship between the rationality of the design of a system (problem formulation, knowledge representation, computational resources) and the rationality of the resulting system.
- extending the application of theories of rationality to learning and adaptation, especially in situations where the learning process must both use and learn preference and utility information.

Supporting Collaboration

Quite apart from the research collaborations within and without AI just described, the subject matter of collaboration and coordination of multiple agents (human or artificial) forms one of the main directions for AI research in coming years (Grosz 1996). To prove useful as assistants, AI systems must interpret the words and deeds of people to model the desires, intentions, capabilities and limitations of those people, and then use these models to choose the most appropriate or helpful actions of their own. Making these interpretations often means relying on statistical properties of past behavior, and choosing how to cooperate often means assessing or negotiating the preferences and tradeoffs held by the various participants.

Studies of collaboration have a long history in sociology, economics, politics, linguistics, and philosophy. AI has studied collaboration issues in four primary contexts: understanding dialogue, constructing intelligent assistants, supporting collaborative and group work, and designing “artificial societies.” In the longest-studied of these contexts, under-

standing dialogue, the normal rules of conversational implicature presuppose cooperative intent on the part of the listener. Asking a computer “Can I see the accounts receivable summary?” should yield either presentation of the summary or an explanation of the reason for its unavailability, not a less-than-helpful “yes” or “no.” Aggravation with the stupidity of computers will never cease without such cooperative interpretation of requests and statements.

In the more recent context of designing intelligent assistants, the assistant systems must seek to understand and support the aims of the user. These systems go beyond mere decision support by attempting to anticipate and satisfy the needs of the user whenever possible and appropriate. The ARPA/Rome Laboratory Planning Initiative (Fowler, Cross, and Owens 1995) and NASA’s VISTA ground-control support system (Horvitz et al. 1992) provide good examples of such assistants.

In a broader context, AI research contributes to providing supportive environments for collaboration and group-cooperative work. As in understanding discourse and designing intelligent assistants, these supportive environments must model processes and plans, but they must also supply methods which reason from these models to coordinate projects, manage workflow constraints, filter and broker information, answer questions, notify participants as appropriate, translate “utterances” between different interface modalities, and generate summaries to quickly bring offline participants up to date.

The newest context, designing artificial societies, introduces a design perspective into economics by seeking to tailor the preferences of agents, the protocols of interaction, and the environmental constraints so as to automatically yield collaboration, non-interference, and other desirable properties of group behavior.

Research on collaborative systems draws together many of the research areas of AI, especially planning, multi-agent learning, speech and language, and image understanding and presentation, and involves fundamental issues of modeling commitment, communication requirements, constraints and tradeoffs, negotiation methods, and methods for resolving conflicts among the intentions of collaborating agents. Collaborative systems also provide an interesting environment for attacking a core problem of knowledge representation, that of amassing enough knowledge about a broad domain, including many application tasks, to improve performance significantly.

*Efficient
and
natural
communi-
cation
holds the
key to
many
of the
promises
of
computers
....*

Situating people and artificial agents in a common environment with a shared domain model, even a rudimentary one, creates the opportunity for large numbers of collaborators to convey their knowledge to and share their discoveries with one another and with the artificial agents, and for each participant to learn from the collaborative experience.

Representative long-term goals in this direction include:

- construct DWIM (do what I mean) capabilities for household, educational, commercial, and industrial appliances, yielding machines that infer the desires and intentions of the users and cooperate with them in achieving their aims.
- construct intelligence “amplifiers” or “protheses” that act with a person to overcome limitations of knowledge, memory, or speed in achieving the person’s aims.
- construct societies of human and artificial agents that collaborate in efficient ways to achieve complex ends.

Enhancing Communication

Efficient and natural communication holds the key to many of the promises of computers, given that relying on command languages, menus, textual display, and other traditional media stymies many potential applications³. The activities these applications support normally rely on many different communication modalities, such as spoken utterances, written texts, and the gestures that accompany them, and effective participation in these activities requires the ability to understand and generate communications in these modalities. In addition, the ability to read would greatly simplify the task of imparting knowledge to artificial agents, considering the vast amount of human knowledge encoded in written form. AI has long addressed these issues, and has contributed to great progress on realizing linguistic and visual communication mechanisms involving multiple modalities, including natural language, gestures, and graphics. The most general form of these abilities, however, lies far beyond current scientific understanding and computing technology.

Ambiguity, intent, and thinking while speaking form some of the main obstacles to achieving the desired communication. Human languages all use a small set of resources (such as words, structures, intonations, and gestures) to convey an exceedingly wide, rich, and varied set of meanings. Speakers often use the same word, structure, or gesture in many different ways, even in the same sentence or episode. Although people rarely notice such

ambiguities, their identification and resolution challenge current speech- and language-processing systems. Intent, or the difference between what people say (or write) and what they actually mean, arises because people rely on their audience to infer many things left unsaid or unwritten from context and common knowledge. Furthermore, people often begin to speak or write before thinking through their ideas completely, using the formulation of utterances as a step in understanding their own partially formed ideas. Both practices result in partial and imperfect evidence for what people really mean to communicate.

Recent developments include the use of statistical models, typically generated automatically, to predict with good accuracy simple grammatical features of utterances such as the part of speech of a word, as well as semantic properties such as the word sense most likely in a given context. These models thus reduce problems caused by ambiguities in the grammatical and semantic properties of words. In other work, idealized models of purposive communicative action support improved discourse modeling.

Much of the success of current natural language processing technology stems from a long and tedious process of incremental improvement in existing approaches. Extracting the best possible performance from known techniques requires more work of this kind, but exploration of new and combined approaches supplies additional opportunities. For example, although statistical and machine-learning techniques in natural language processing offer broad (but shallow) coverage and robustness with respect to noise and errors, grammatical and logical techniques offer deeper analyses of meaning, purpose, and discourse structure. These two types of techniques could complement one another, with the symbolic techniques serving to specify a space of interpretation possibilities and the statistical techniques serving to evaluate efficiently the evidence for alternative interpretations. The results of such integration should prove of value to all natural language processing applications, from information extraction and machine translation to collaborative interfaces. Another opportunity involves determining the most effective combination of natural language processing technology with other technologies to forge effective multimodal user interfaces.

Representative long-term goals in this direction include:

- provide spoken and gestural control for common appliances (lighting, heating and

air conditioning, computers, televisions, automobiles, etc.) in many settings: the smart house, office, factory, etc.

- automate the formalization of knowledge from books and other texts.
- provide simultaneous translation between languages, automatic translations of written texts, and natural spoken renditions of written or electronic texts.
- permit the use of natural spoken or written languages in interacting with large-scale databases and sources of knowledge: automate telephone operators, librarians, travel agents, and other services.

Obtaining Knowledge

The most widespread benefit so far of putting AI into practice consists of the bodies of human knowledge formalized with an eye to mechanizing reasoning. Though the idea of writing down expert knowledge in explicit form goes back at least to the code of Hammurabi, if not to the earlier Egyptian and Babylonian inventors of geometry and arithmetic, the knowledge formalized and codified through AI methods has a very different character and purpose. AI compilations go beyond mere books by representing not just the "factual" knowledge about the subject but also the reasoning processes appropriate to specific uses of the knowledge. Authors of books focus on conveying propositional knowledge, normally leaving it up to the reader to learn how to apply and interpret the knowledge. Authors of traditional computer programs focus on representing processes, necessarily leaving it to the documentation (if any) to convey the facts used or presupposed in the design or operation of the programs. The efficient mechanization, maintenance, and explication of expertise requires expressing both types of knowledge in declarative representations. Reasoning systems may then manipulate these representations in a variety of ways to support explanation, guidance, maintenance, and learning. The novel opportunities created by capturing reasoning processes as well as factual knowledge have stimulated great effort in this area, and construction of knowledge-based systems today goes on in hundreds if not thousands of sites. Most of this work stays invisible, as businesses and organizations view these bodies of articulated expertise as trade secrets and competitive advantages they do not wish to see their competitors replicate.

The problem of formalizing knowledge remains one of the principal challenges to AI research. Current successful knowledge-based

systems rely on carefully limiting the scope and domain of the formalized knowledge, in order to make it tractable to collect, codify, and correct this knowledge. The experience of AI shows two key lessons about this task: formalizing knowledge is difficult, and adequate formalizations are feasible. The current formalizations, although adequate to the specific tasks addressed so far, fail to support the integration aims of AI research in several ways, and overcoming these limitations forms a major task for AI research that forces consideration of many fundamental issues in knowledge representation.

First, current formalizations do not cover the broad scope of knowledge needed for intelligent activity outside of carefully circumscribed circumstances, in particular, the knowledge needed by integrated systems acting in everyday household, social, workplace, or medical situations; nor do current formalizations fit together smoothly, since the conceptualizations adequate to one domain rarely do justice to the concepts from peripheral domains. Addressing these problems calls for constructing formal "ontologies" or conceptual organizations adequate to the broad scope of human knowledge that include propositional, uncertain, and algorithmic and procedural knowledge; finding ways for efficiently structuring, indexing, and retrieving large-scale bodies of knowledge; reasoning across multiple domains, and across the same knowledge represented for different purposes; and efficiently representing the contexts or foci of attention that form the specific portions of the large bodies of interest in episodes of reasoning. To prove useful in practice, the structures and methods developed here will require (and benefit from) smooth integration with extant databases and database organizations, as well as a closer integration between declarative knowledge about formalized procedures and the use of typical procedural programming languages.

Second, most extant bodies of formalized knowledge presuppose, but avoid formalizing, the commonsense knowledge so characteristic of people. Although expert performance often does not depend on common sense (as any number of jokes about experts illustrate), commonsense knowledge and reasoning appear crucial, both for tying together domains of expert knowledge and for recognizing the boundaries of specialized expertise in order to avoid acting inappropriately. Thus constructing broadly knowledgeable and capable systems requires formalizing and mechanizing commonsense reasoning. The

The most widespread benefit so far of putting AI into practice consists of the bodies of human knowledge formalized with an eye to mechanizing reasoning.

As a field, AI embarks on the next fifty years excited about the prospects for progress, eager to work with other disciplines, and confident of its contributions, relevance, and centrality to computing research.

amount of knowledge needed for intelligent action across the broad range of human activity promises to dwarf even the large body developed in the long-running CYC project (Lenat 1995).

Third, current methods for constructing bodies of formalized knowledge require much (often heroic) human labor on the part of the best (and least available) people knowledgeable in each area, as does their maintenance or adjustment as circumstances change. Though some applications may command the resources these methods demand, realizing the benefits of knowledge-based systems in the broad spectrum of applications requires developing methods in which the necessary mass of knowledge accumulates through many small contributions made by a range of people, both the ordinary many and the expert few, and through the exploitation of machine labor.

The goal of enabling people to make incremental contributions to knowledge bases motivates research on simplifying and streamlining the process of updating and maintaining the system's knowledge and abilities. Performing the primary tasks—identifying gaps in knowledge, expressing the knowledge needed to fill those gaps, and checking new knowledge against old—requires knowledge about the system's own knowledge and operation. Accordingly, methods for these tasks rely on declarative formalizations of both the processes for carrying out each of these steps and of the structure and function of each part of the knowledge base, rather than on the mainly procedural representations found in most programming languages. Such formalizations, and methods for using them, form the basis of the extensively investigated KADS methodology and library (Schreiber, Wielinga, and Breuker 1993). Automating these methods as part of the system's own reasoning permits the system to exhibit limited forms of self-understanding, and makes the processes of reasoning and acquisition quite synergistic.

Of course, people do not always possess the knowledge they need, and even with automated helps may still find it extremely hard to articulate the knowledge they do have. Work on machine learning and discovery techniques bridges the gap in many cases. This work builds on statistical methods and “connectionist” models inspired by neurophysiology, but extends them to cover a much richer class of models and to combine symbolic and numerical methods in useful ways. Current methods can capture some expert behavior, but often do so in a way that does not provide useful explanations of the behavior.

Using these bits of embodied expertise in many cases requires further analysis to transform the knowledge (e.g., “turn to the right if $E > 0$ ” for some complex numerical expression E) into a more explicit and sensible form (“turn to the right if the road turns right”). For example, one important new area uses Bayesian networks to summarize prior knowledge in an understandable way, Bayesian inference to combine prior knowledge with new data, and techniques of compositional representation to learn (construct) new networks when the prior network fails to accommodate the new data adequately. Another new area, knowledge discovery in databases (or “data mining”), finds regularities and patterns in extremely large data sets by integrating techniques from machine learning and statistics with modern database technology.

Representative long-term goals in this direction include:

- construct evolving formal encyclopedias of knowledge and methods and techniques for tailoring them to particular ends, starting with the most commonly used and sharable categories but ultimately covering all human knowledge and methods.
- determine the most effective ways to represent information for different purposes, together with means for combining representations of different types of knowledge and for translating among these.
- develop automated tutors that use formal encyclopedias to help educate humans in all topics and at all levels (within the scope of the encyclopedias), that use questions and observations of the student to model the student's knowledge, abilities, and learning style, and that use the model to tailor construction of successive lessons or exercises to the particular needs of the student.
- design organizations and factories that improve themselves, automatically analyzing experience to learn hidden efficiencies and inefficiencies.
- automate the more routine and data-intensive areas of commercial, industrial, statistical, and scientific research.

Deepening Foundations

Mathematical work in AI has long swum in the same waters as the theory of computation, logic, and mathematical economics. Early mathematical work focused on the theory of search and the power of statistical and neural-net models of recognition, but later work has added deep and rich theories of nonmonotonic reasoning; of the expressiveness, inferential

complexity, and learnability of structured description languages; and of stochastic search techniques. Some of this work employs notions taken from or developed in concert with the theory of computation, such as time-space classifications of computational complexity and epistemic theories of distributed systems. AI theories must consider richer classifications of systems, however, since the properties distinguishing minds (belief, desire, intent, rationality, consciousness, sensory and motor faculties, etc.) constitute a larger and more puzzling set than those distinguishing computations. Although reasonable formalizations exist for some of these distinguishing properties, others remain problems for formulation. AI shares some of these problems with the mathematical sides of logic, economics, physics, and the theory of computation, but alone among the disciplines aims to characterize the full range of possible psychological organizations for minds, from the trivial to the superhuman. Since conceptual analysis flourishes best in the context of solving specific problems, the concrete complex systems developed in AI research bestow an advantage on AI over its relatives, which typically lack nontrivial yet tractable examples to study. These concrete, complex examples continue to attract the attention of workers in other disciplines, and this comparative advantage promises a stream of AI contributions to these other fields.

Representative long-term goals in this direction include:

- identify realistic yet theoretically comprehensible and tractable theories of rationality appropriate to agents of limited knowledge and abilities.
- properly formalize all aspects of psychological theories: not just reasoning and decision-making, but habit, memory, emotion, motivation and other aspects as well.
- find appropriate common mathematical forms that reconcile theories of informational (computational, cognitive, economic) and material (physical) agents.

Conclusion

These studies are an impetus to youth, and a delight to age; they are an adornment to good fortune, refuge and relief in trouble; they enrich private and do not hamper public life; they are with us by night, they are with us on long journeys, they are with us in the depths of the country.

Cicero, *Pro Archia*, VII.xvi

By addressing both the underlying nature of intelligence and the development of theories, algorithms, and engineering techniques necessary to reproduce reliable, if rudimentary, machine intelligence, AI research makes numerous, large, and growing contributions to computing research and to the evolving social and industrial information infrastructure. Some contributions come through study of the deep scientific issues that concern our understanding of computation, intelligence, and the human mind. Others come through practical applications that help make computer systems easier and more natural to use and more capable of acting as independent intelligent workers and collaborators. Continued progress requires pursuing both types of contributions. The practical applications alone offer some of the strongest motivations for pursuing the scientific studies, as achieving the practical benefits seems hopeless without obtaining a deeper scientific understanding of many issues. At the same time, success in many of the scientific investigations calls for developing broad bodies of knowledge and methods—and practical applications provide the most natural context for developing these bodies of intelligence.

AI researchers retain enthusiasm about their field, both about the problems it addresses and about the ongoing progress on these problems, even as it has matured into a field of substantial content and depth. AI has needs that intersect with all areas of computing research, and a corresponding interest in partnerships with these areas in advancing knowledge and technique on these shared problems. It offers techniques and theories providing leverage on hard problems and also offers large important problems that might well serve as target applications for much of computing research. Only a few of these have been described in this short summary, and many opportunities remain for joint exploration with other areas of computing. As a field, AI embarks on the next fifty years excited about the prospects for progress, eager to work with other disciplines, and confident of its contributions, relevance, and centrality to computing research.

Write the vision; make it plain upon tablets, so he may run who reads it.

Habakkuk 2:2, *RSV*

Acknowledgments

The editors thank the group members (see Footnote 1) for their essential contributions and tireless efforts. The editors also thank Ger-

hard Brewka, Peter Denning, Eric Grimson, Chris Hankin, Patrick Hayes, James Hendler, Larry Rudolph, Elisha Sacks, Joseph Schatz, Pietro Torasso, and Peter Wegner for providing additional comments and suggestions.

This report draws on two longer ones prepared by the American Association for Artificial Intelligence, namely “A Report to ARPA on Twenty-First Century Intelligent Systems” (Grosz and Davis 1994) and “The Role of Intelligent Systems in the National Information Infrastructure” (Weld, Marks, and Bobrow 1995). The editors thank AAAI for permission to use portions of the latter report (see footnotes 2 and 3).

Notes

1. This report represents the efforts of the Strategic Directions in Computing Research AI Working Group—Ronald Brachman, Thomas Dean (co-chair), Johan de Kleer, Thomas Dietterich, Jon Doyle (co-chair), Cordell Green, Barbara Grosz, Ian Horswill, Leslie Pack Kaelbling, Daphne Koller, Joseph Marks, Fernando Pereira, Bart Selman, Yoav Shoham, Howard Shrobe, William Swartout, Michael Wellman, and Shlomo Zilberstein—formed as part of the ACM Workshop on Strategic Directions in Computing Research, held at the Massachusetts Institute of Technology Laboratory for Computer Science on June 14-15, 1996. Group members worked out the underlying ideas of this report at the meeting, aided by prior discussions and their individual position statements. Ron Brachman presented a talk on strategic directions in AI to the workshop as a whole. The general structure of this report derives from his talk, augmented through many contributions by the above. The editors, of course, bear the responsibility for any flaws.

2. The list of areas of AI and their descriptions draws some text from the corresponding list appearing in (Weld, Marks, and Bobrow 1995). The borrowed text is Copyright © 1995 American Association for Artificial Intelligence, and is reprinted with permission from AAAI.

3. Much of the text in the section on enhancing communication was taken directly from sections of (Weld, Marks, and Bobrow 1995), with some revisions. The borrowed text is Copyright © 1995 American Association for Artificial Intelligence, and is reprinted with permission from AAAI.

References

- Feigenbaum, E. A. 1996. How the “What” Becomes the “How.” *Communications of the ACM* 39(5): 97–104.
- Fowler, N., III; Cross, S. E.; and Owens, C. 1995. The ARPA-Rome Knowledge-Based Planning and Scheduling Initiative. *IEEE Expert* 10(1): 4–9.
- Goebel, J.; Volk, K.; Walker, H.; Gerbault, F.; Cheeseman, P.; Self, M.; Stutz, J.; and Taylor, W. 1989. A Bayesian Classification of the IRAS LRS Atlas. *Astronomical Astrophysics* 222:L5–L8.
- Grimson, W. E. L.; Lozano-Perez, T.; Wells, W. M., III; Ettinger, G. J.; White, S. J.; and Kikinis, R. 1996. An Automatic Registration Method for Frameless Stereotaxy, Image Guided Surgery, and Enhanced Reality Visualization. *IEEE Transactions on Medical Imaging* 15(2): 129–140.
- Grosz, B. 1996. Collaborative Systems. *AI Magazine* 17(2): 67–85.
- Grosz, B., and Davis, R., eds. 1994. A Report to ARPA on Twenty-First Century Intelligent Systems. *AI Magazine* 15(3): 10–20.
- Heckerman, D. E. 1991. *Probabilistic Similarity Networks*. Cambridge, Mass.: MIT Press.
- Heckerman, D. E.; Breese, J.; and Rommelse, K. 1995. Decision-Theoretic Troubleshooting. *Communications of the ACM* 38(3): 49–57.
- Horvitz, E.; Srinivas, S.; Rouokangas, C.; and Barry, M. 1992. A Decision-Theoretic Approach to the Display of Information for Time-Critical Decisions: The Vista Project. In Proceedings of the SOAR-92 Conference on Space Operations Automation and Research. Washington, D.C.: National Aeronautics and Space Administration.
- Jochem, T., and Pomerleau, D. 1996. Life in the Fast Lane: The Evolution of an Adaptive Vehicle Control System. *AI Magazine* 17(2): 11–50.
- Kasparov, G. 1996. The Day That I Sensed a New Kind of Intelligence. *Time* 147(13): 55.
- Lenat, D. B. 1995. CYC: A Large-Scale Investment in Knowledge Infrastructure. *Communications of the ACM* 38(11): 33–38.
- Lyon, R. F., and Yaeger, L. S. 1996. On-Line Hand-Printing Recognition with Neural Networks. In Proceedings of the Fifth International Conference on Microelectronics for Neural Networks and Fuzzy Systems, 201–212. Washington, D.C.: IEEE Computer Society Press.
- Minsky, M. 1962. Problems of Formulation for Artificial Intelligence. In Proceedings of a Symposium on Mathematical Problems in Biology, 35–46. Providence, R.I.: American Mathematical Society.
- Naj, A. K. 1996. Manufacturing Gets a New Craze from Software: Speed. *Wall Street Journal*, 13 August 1996, B4. See also <http://www.i2.com/>.
- Reddy, R. 1996. To Dream the Possible Dream. *Communications of the ACM* 39(5): 105–112.
- Rosenbloom, P. S.; Laird, J. E.; and Newell, A. 1993. *The Soar Papers: Readings on Integrated Intelligence*. Cambridge, Mass.: MIT Press.
- Schaeffer, J.; Culberson, J.; Treloar, N.; Knight, B.; Lu, P.; and Szafron, D. 1992. A World Championship Caliber Checkers Program. *Artificial Intelligence* 53:273–289.
- Schreiber, A. T.; Wielinga, B. J.; and Breuker, J. A., eds. 1993. *KADS: A Principled Approach to Knowledge-Based System Development*. San Diego, Calif.: Academic.
- Selman, B.; Brooks, R. A.; Dean, T.; Horvitz, E.; Mitchell, T.; and Nilsson, N. J. 1996. Challenge Problems for Artificial Intelligence. In Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI-96), 1340–1345. Menlo Park,

Calif.: American Association for Artificial Intelligence.

Senator, T. E.; Goldberg, H. G.; Wooton, J.; Cottini, M. A.; Khan, A. F. U.; Klinger, C. D.; Llamas, W. M.; Marrone, M. P.; and Wong, R. W. H. 1995. The Financial Crimes Enforcement Network AI System (FAIS): Identifying Potential Money Laundering from Reports of Large Cash Transactions. *AI Magazine* 16(4): 21–39.

Smith, D. R.; Parra, E. A.; and Westfold, S. J. 1996. *Synthesis of Planning and Scheduling Software*, *Advanced Planning Technology*, ed. Austin Tate, 226–234. Menlo Park, Calif.: AAAI Press.

Tambe, M.; Johnson, W. L.; Jones, R. M.; Koss, E.; Laird, J. E.; Rosenbloom, P. S.; and Schwamb, K. 1995. Intelligent Agents for Interactive Simulation Environments. *AI Magazine* 16(1): 15–39.

Tesauro, G. 1995. Temporal Difference Learning and TD-Gammon. *Communications of the ACM* 38(3): 58–68.

Weld, D. S.; Marks, J.; and Bobrow, D. G., eds. 1995. The Role of Intelligent Systems in the National Information Infrastructure. *AI Magazine* 16(3): 45–64.

Wright, J. R.; Weixelbaum, E. S.; Vesonder, G. T.; Brown, K. E.; Palmer, S. R.; Berman, J. I.; Moore, H. H. 1993. A Knowledge-Based Configurator That Supports Sales, Engineering, and Manufacturing at AT&T Network Systems. *AI Magazine* 14(3): 69–80.



Jon Doyle works as principal research scientist in the Clinical Decision Making Group in the Laboratory for Computer Science at the Massachusetts Institute of Technology (MIT), following similar appointments at Stanford University and Carnegie Mellon University. He received a bachelor's

degree with honors in mathematics in 1974 from the University of Houston and a master's (1977) and a doctorate (1980) from MIT. Elected a fellow in 1991 and a councilor in 1996 of the American Association for Artificial Intelligence, Doyle has served as chair of ACM SIGART (1989–1991); co-program chair of KR'94; conference chair of KR'96; associate chair of Strategic Directions in Computing Research (1996); and associate editor of *JAIR*, *Computational Intelligence*, *JOLLI*, and *ACM Computing Surveys*. His research interests include rational psychology, theories of limited rationality, qualitative decision theory, nonmonotonic reasoning, reason maintenance, reasoned deliberation, representation, logic, philosophy, physics, and history.



Thomas Dean received a B.A. in mathematics from Virginia Polytechnic Institute and State University in 1982 and his Ph.D. in computer science from Yale University in 1986. He is currently professor of computer science and cognitive and linguistic science at Brown University in Providence, Rhode Island. His general research interests include temporal and spatial reasoning, planning, robotics, learning, and probabilistic inference. His current research is concerned with theories of temporal and spatial inference for reasoning about actions and processes. Of particular interest are problems in which the notion of risk is complicated by there being limited time for both deliberation and action.

Dean is a fellow of the American Association for Artificial Intelligence; a member of the board of directors of the Computing Research Association; and a member of the board of trustees for the International Joint Conferences on Artificial Intelligence, Inc. He was the recipient of a National Science Foundation Presidential Young Investigator Award in 1989; he served as the program co-chair for the 1991 National Conference on Artificial Intelligence; and he will be the program chair for the 1999 International Joint Conference on Artificial Intelligence, to be held in Stockholm, Sweden. He is co-author, with Mike Wellman, of the Morgan Kaufmann text entitled *Planning and Control*, which ties together techniques from AI, operations research, control theory, and the decision sciences. His latest book is *Artificial Intelligence: Theory and Practice* (Benjamin Cummings) with James Allen and John Aloimonos.

Dean is a fellow of the American Association for Artificial Intelligence; a member of the board of directors of the Computing Research Association; and a member of the board of trustees for the International Joint Conferences on Artificial Intelligence, Inc. He was the recipient of a National Science Foundation Presidential Young Investigator Award in 1989; he served as the program co-chair for the 1991 National Conference on Artificial Intelligence; and he will be the program chair for the 1999 International Joint Conference on Artificial Intelligence, to be held in Stockholm, Sweden. He is co-author, with Mike Wellman, of the Morgan Kaufmann text entitled *Planning and Control*, which ties together techniques from AI, operations research, control theory, and the decision sciences. His latest book is *Artificial Intelligence: Theory and Practice* (Benjamin Cummings) with James Allen and John Aloimonos.