

A Report to ARPA on Twenty-First Century Intelligent Systems

The American Association for Artificial Intelligence

Edited by Barbara Grosz, President and Randall Davis, President-Elect

Report Committee Members:

Ruzena Bajcsy (University of Pennsylvania), Piero Bonissone (GE Corporate Research and Development), Bruce Bullock (ISX), Larry Hunter (National Library of Medicine), Steve Minton (NASA Ames Research Center), Tom Mitchell (Carnegie Mellon University), Ray Perrault (SRI International), T. Lozano-Perez (Massachusetts Institute of Technology), Martha Pollack (University of Pittsburgh), Paul Rosenbloom (University of Southern California), Stuart Shieber (Harvard University), Howie Shrobe (Massachusetts Institute of Technology), Dan Weld (University of Washington)

ARPA Observers:

Steve Cross, Gio Wiederhold

This report stems from an April 1994 meeting, organized by AAI at the suggestion of Steve Cross and Gio Wiederhold.¹ The purpose of the meeting was to assist ARPA in defining an agenda for foundational AI research. Prior to the meeting, the fellows and officers of AAI, as well as the report committee members, were asked to recommend areas in which major research thrusts could yield significant scientific gain—with high potential impact on DOD applications—over the next ten years. At the meeting, these suggestions and their relevance to current national needs and challenges in computing were discussed and debated. An initial draft of this report was circulated to the fellows and officers. The final report has benefited greatly from their comments and from textual revisions contributed by Joseph Halpern, Fernando Pereira, and Dana Nau.

Introduction

Computer systems are becoming commonplace; indeed, they are almost ubiquitous. We find them central to the functioning of most business, governmental, military, environmental, and health-care organizations. They are also a part of many educational and training programs. But these computer systems, while increasingly affecting our lives, are rigid, complex, and incapable of rapid change. To help us and our organizations cope with the unpredictable eventualities of an ever-more volatile world, these systems need capabilities that will enable them to adapt readily to change. They need to be intelligent.

Our national competitiveness depends increasingly on capacities for accessing, processing, and analyzing information. The computer systems used for such purposes must also be intelligent. Health-care providers require easy access to information systems so they can track health-care delivery and identify the most recent and effective medical treatments for their patients' conditions. Crisis management teams must be able to explore alternative courses of action

and support decision making. Educators need systems that adapt to a student's individual needs and abilities. Businesses require flexible manufacturing and software design aids to maintain their leadership position in information technology, and to regain it in manufacturing.

Advanced information technology can help meet these and many other needs in our society. Advances in computer and telecommunications have made available a vast quantity of data, and given us computational power that puts the equivalents of mainframes on our desktops. However, raw information processing power alone, like brute strength, is useful but insufficient. To achieve their full impact, computer systems must have more than processing power—they must have intelligence. They need to be able to assimilate and use large bodies of information and collaborate with and help people find new ways of working together effectively. The technology must become more responsive to human needs and styles of work, and must employ more natural means of communication.

To address the critical limitations of today's systems, we must understand the ways people reason about and interact with the world, and must develop methods for incorporating intelligence in computer systems. By providing computer programs that amplify human cognitive abilities and increase human productivity, reach, and effectiveness, we can help meet national needs in industries like health care, education, service, and manufacturing.

Artificial intelligence (AI) is a field that studies intelligent behavior in humans using the tools—theoretical and experimental—of computer science. The field simultaneously addresses one of the most profound scientific problems—the nature of intelligence—and engages in pragmatically useful undertakings: developing intelligent systems. The concepts, techniques, and technology of AI offer us a number of ways to discover what intelligence is—what one must know to be smart at a particular task—and a variety of computational techniques for embedding that intel-

ligence in a program.

This report describes AI research areas where fundamental scientific advances could enable intelligent systems to meet national needs. It sets this research in context by presenting four families of intelligent systems that make concrete the excitement of such systems and their potential payoff. Useful, if limited members of these families should be possible within five years, although the full visions are at least one to two decades away. These systems have applications in the full range of grand challenge and national applications areas, including health care, education and training, and the environment. In this report we will refer to these four types of systems as "high-impact applications systems."

Intelligent Simulations: Systems that generate realistic simulated worlds would enable extensive, affordable training and education that can be made available anytime and anywhere. A new generation of intelligent simulation capabilities could support the construction of programs that model complex situations, involving both complicated devices and significant numbers of intelligent simulated people. Uses of these capabilities range from crisis management to product evaluation and entertainment.

Intelligent Information Resources: Information-resource specialist systems would support effective use of the vast resources of the national information infrastructure. These systems would work with their users to determine users' information needs, navigate the information world to locate appropriate data sources—and appropriate people—from which to extract relevant information. They would adapt to changes in users needs and abilities as well as changes in information resources. They would be able to communicate in human terms in order to assist those with limited computer training.

Intelligent Project Coaches: Software designed to act as an intelligent, long-term team member could help to design and to operate complex systems. An intelligent project coach system can assist with design of a complex device (such as an airplane) or a

large software system by helping to preserve knowledge about tasks, to record the reasons for decisions, and to retrieve information relevant to new problems. It could help at the operational level to improve diagnosis, failure detection and prevention, and system performance. Project coach systems do not need to be experts themselves; rather, they could significantly boost capability and productivity by collaborating with human experts, assisting them by capturing and delivering organizational memory.

Robot Teams: Intelligent robot system teams can perform tasks that are dangerous, such as environmental clean-up, mine removal, fire-fighting, and rescue operations. They can also perform those tasks that, while essential to the smooth functioning of our society, are mundane, repetitive, or unappealing to human workers. Individual robot team members may have limited capabilities; the teams need not be fully independent. Instead they can work together under human supervision, with robots doing the work and people providing direction and guidance.

Systems like these are motivated by and responsive to pressing national needs. Their technical requirements are considerably broader than AI alone can provide² and will, in the coming years, be the goal of system developers in much hardware and software research. Advances across computing areas will contribute substantial power to the solutions developed. AI capabilities will be key to making the systems intelligent, adaptable, far more accessible to the general public, and, thus, dramatically more effective.

The Underlying Foundation and the Research Need

A common core of capabilities is needed to construct intelligent systems in all the aforementioned categories. These include abilities to reason about the task being performed and basic common sense facts that affect it; to reason about the collaborative process and the knowledge and

capabilities of other systems and people participating in an interaction; to communicate with users in human terms, producing and understanding combinations of spoken and written language, drawings, images, and gestures; to perceive the world; to coordinate perception, planning, and action; and to learn from previous experience and adapt behavior accordingly.

Understanding these capabilities in humans and developing computational techniques to embody them in programs has been a central focus of AI research. A solid foundation has been developed in the large body of previous research. This work produced the technology that underlies the few thousand knowledge-based expert systems used in industry today; it also made major contributions to the DART system, which was used in deployment planning in the Desert Shield effort, as well as many applications in planning, learning, perception, and language processing.³

A major challenge for the next decade is to significantly extend this foundation to make possible new kinds of high-impact application systems. Although the development of systems with the most sophisticated capabilities will require long-term effort, in each category more restricted but still usefully intelligent systems can and will be developed. The planning techniques used in DART exemplify this kind of nearer-term payoff and more immediate contribution to our society's needs.⁴

In 1969, prescient thinkers in the government and military saw the potential for payoff in national computer networks. The dream is now coming true, 25 years later, in a way that is having an enormous impact on America's global economic and strategic position during the information age. It took long-term vision and long-term support to realize the potential of networks.

What can we do *today* for the next 25 years? What can we undertake that will have an explosion of similar importance in five, ten, or even 25 years? What will seem a wise investment in 2019? This document sets out one promising research agenda.

In the next section of this report, we will describe in greater detail the different high-impact application systems and the intelligent-systems capabilities each needs. The Research Thrust Areas section then presents seven cross-cutting research areas that incorporate the major scientific problems that must be solved before such systems can be really useful. It also briefly characterizes the scientific and technological base on which this research can draw. In addition to the scientific challenges, there are issues of large systems engineering to be confronted and infrastructure needs that must be addressed; these are discussed in the Building Large Systems section.

High-Impact Application Systems

Intelligent Simulation Systems

For many tasks, on-the-job training is extremely effective, providing the trainee with the chance to make real, on-the-spot decisions and see the consequences. On-the-job training is impossible, however, when a bad decision can be disastrous—for example, in controlling a steel mill, or making diagnoses and prescribing treatment in an operating room, or running a large company, or making battle management decisions. Simulation systems that could portray realistic simulated worlds, and in particular that had the capability to produce realistic simulations of people, would enable development of training systems for such situations. These same simulation capabilities are also important when the cost of assembling large groups of people for training is prohibitive.

Many educational, commercial, military, entertainment, and scientific applications require the capability of generating realistic simulated worlds.

Training: A training system for crisis management teams could provide first-hand practice in handling problems like those that arose during Hurricane Andrew; a large-scale battlefield simulation could be used to train commanders for new types of terrain, equipment, and tactics.

Education: An interactive history book could allow students to discuss the underlying causes and effects of

the American Civil War with Abraham Lincoln and Robert E. Lee; an environment for learning Japanese could take the language learner on a simulated visit to Tokyo where he or she could interact with shopkeepers, taxi drivers, and business people; simulations of real markets that take into account information states of agents could allow students to explore the consequences of different economic theories.

Industry and Commerce and Military Systems: An evaluation environment for new products, such as automobiles or airplanes, could use simulations of people to test the feasibility of the product's construction, use, and maintenance before it has been built; a new product design could be "used" by simulated people although it existed only on paper; potential customers could try out the product in a simulation. Military systems design requires such capabilities. A new submarine, for example, that might cost one billion dollars to design, will contain miles of pipes and ducts. These conduits are often run through cramped passageways, which must also be used by people. Realistic simulations of a person's abilities to move through this maze can save millions of dollars of design cost and could significantly reduce the time to deployment of a new design.

Entertainment: An interactive mystery book could enable the reader and Sherlock Holmes to match wits with Dr. Moriarty.

Simulation, both computerized and manual, has a long history. Simulated worlds of many types are already becoming widely available. Current applications range from simple video games and building walkthroughs to the SIMNET battlefield simulations involving thousands of real and simulated agents. However, the simulated worlds that can be generated today have limited physical realism and lack realism in their simulations of people. Current simulation technologies are also expensive to program.

The systems we anticipate differ in both scale and function from those that exist today. The anticipated scale of next-generation simulations is illustrated by the problem of providing accurate simulations of a crisis like Hurricane Andrew that would be used

in training crisis managers. Such simulations might require thousands of actors to play the role of victims, fire fighters, police, and emergency rescue squads. It might be economical to use actual people for only a few of these roles; the rest could be simulated.

The advanced functionality we seek is illustrated by use during an actual crisis. An emergency coordinator might employ a simulation of the crisis to analyze different potential responses and predict their likely outcomes. Intelligent simulation technology can assist people in such stressful, time-pressured situations to look further ahead in determining the consequences of proposed actions.

A key challenge in achieving the potential of simulated worlds is constructing realistic humanlike agents. These agents must be able to coordinate perception, planning, and action (a research topic discussed in the Coordination of Perception, Planning, and Acting subsection), learn (see the Learning, Information, Elicitation, and Automatic Adaption subsection), understand and interact with their world (see the subsections on Perception and Reasoning and Representation), deal with other agents (see the subsection on Coordination and Collaboration), and use natural language (see the subsection on Human-Computer Communication in Multiple Modalities). Even generating a believable animation of a humanlike agent is a complex task: real-time generation of realistic facial expressions is difficult both conceptually and computationally, as is producing coordinated movement of the hundreds of joints in a human body used in such simple actions as getting into a car.

Providing all or even a significant portion of this functionality is a challenging mission. However, useful agents can be constructed with only some of these capabilities—even in limited form. For example, early results using the semi-automated OPFOR capability in SIMNET demonstrate the practical utility of even very simple simulated agents.

Information-Resource Specialist Systems

The accumulation of information that

is available electronically presents a major dilemma. The good news is that all of the world's electronic libraries are now at your disposal; the bad news is that you're on your own—there's no one at the information desk. For the NII to be useful, people will have to be able to find information relevant to their problems and tasks, in a reasonable amount of time, with reasonable effort.

A recent report by the Office of Science and Technology Policy noted that “[i]n the near future, every home and business could have an information appliance that combines the capabilities of telephone, television, newspaper, computer, and Internet services such as electronic mail.” But the real power envisioned in this statement: the ability of an information appliance to provide users easy access to the information they need, and access to potential human collaborators—indeed the very notion of an “information appliance”—will be realized only if the appliances are intelligent. We need little more than the traditional example of the flashing 12:00 on thousands of VCRs to appreciate how a complex and unintelligent device can thoroughly undermine our efforts to benefit from it. Anyone frustrated by voice-mail systems that require navigation of long sequences of button presses recognizes the need to provide better ways of communicating with information systems.

An information-resource specialist system (IRSS) could meet a wide range of needs at home, at work, and at school. Such systems would be tailored to individual users rather than a single project and its needs; consequently an IRSS would be able to assist its user with a broad range of information needs.

The clearest uses of IRSSs are as aids to finding information related to a specific problem. For example, a teacher designing a new course might use his or her IRSS to find relevant background materials, slides to use in his or her presentation, or even information about similar courses taught elsewhere. Some of the resources the teacher locates might be used later by students in the course. A rural doctor whose patient presents a rare condi-

tion might use his or her IRSS to help assess different treatments or identify new ones or to locate and consult with an appropriate specialist. A neurosurgeon might use his or her IRSS to search a national database for cranial tomography images that resemble the image for a gunshot victim to locate techniques that were used in similar situations and the doctors who used them.

An IRSS could help track information over long periods of time. For example, a market analyst for a major department store chain might use his or her IRSS to help identify emerging industry trends. The IRSS would automatically monitor national trade databases, and detect and alert its user to changing market patterns.

To realize their enormous potential, IRSSs must be powerful, flexible, and easy to use. Users must be able to communicate in whatever way is most natural to them: typing or speaking, for example, in their native language rather than some artificially designed language. The IRSS must allow the use of diagrams and gestures, combining media and modalities in whatever mix is best for getting the message across (see the subsections Perception, Human-Computer Communication in Multiple Modalities, and Content-Based Retrieval). The commands that users issue will be general and often vague; nevertheless the IRSS must accurately determine how to perform such commands (see subsection Coordination of Perception, Planning, and Acting). The information that a user needs will often not be stored at any one site; thus the IRSS will need to be able to access multiple sites and recognize common information. To actively and continuously seek out useful information, an IRSS will need to learn which topics are of long- and short-term interest to each user (see the subsection Learning, Information Elicitation, and Automatic Adaptation). Finally, several specialist systems may need to coordinate to locate the relevant information (see the subsection Coordination and Collaboration).

Intelligent Project Coaches

The world is becoming increasingly complex and information rich. Not

long ago, automobile engines had carburetors with manual adjustments; now they have microprocessor controlled components and local area networks operating inside the engine compartment. The richness and complexity of these systems are both appealing and daunting. Automobiles can be more energy efficient and less polluting if they are carefully designed and controlled, but complex designs are more difficult to create and harder to debug, modify, and operate. The problems of complexity are nowhere more evident than in the problem of design, modification, and maintenance of the hardware and software that make up complex computer systems.

Recent world-wide political and financial events have intensified the need to renew the competitiveness of manufacturing industries. In response to this need, computerized tools are being developed for a wide variety of design and manufacturing activities, including design development, design analysis, process planning, production planning and scheduling, and production control. However, many of these tools have not been widely accepted by industry because they are difficult to use, inflexible, or lack power.

The Boeing 777 aircraft illustrates major advances in design technology: new tools enabled designers to check spacing and clearance so accurately that a physical mock-up version of the plane was not needed. Even so, these tools still had limitations. They did not incorporate, for example, vast quantities of design information. As a result, engineers had to manually consult printed documents. Other information, such as some of the compromises made in the design process, was never recorded. Now lost, this information will be greatly missed when the design is revised (as all designs are) in the future. Nor did the tools automate, to any extent, methods for assessing the habitability of the designed aircraft or the usefulness of the design.

Intelligent project coach systems [IPCs] can address such needs. They could function as coworkers, assisting and collaborating with design or

operations teams for complex systems. They could also supply institutional memory. The IPC could remember and recall the rationale of previous decisions, and, in times of crisis, explain the methods and reasoning previously used to handle that situation. IPCs would typically incorporate intelligent simulation and information resources systems as components.

Design associates are an IPC type that could assist business, governmental, and scientific design activities. While engineers design artifacts, politicians design public policies, and programmers design software systems, central characteristics of the design process are shared by all. Significant design projects are typically accomplished by teams; designs are almost always redesigned; effective redesign requires an understanding of why previous design choices were made and of how these choices achieved or compromised the desired goals; and all are vulnerable to loss of important information from changes in design-team membership.

For example, an IPC for aircraft design could enhance collaboration by keeping communication flowing among the large, distributed design staff, the program managers, the customer, and the subcontractors. Such an IPC could also assist in adapting existing design during modifications and subsequent generations; support concurrent simulations of an overall design whose components might be in various stages of completion; and capture design rationales (such as for wing design), making them readily available during the entire design lifetime and accessible for maintenance and repair.

One critical area in which IPCs could assist is software development. An IPC could keep track of specifications, design proposals, and implementations for a software project throughout its life cycle. It could record the design decisions of a constantly changing team and also be a repository of solutions and components for new projects. Reasoning techniques could be used to track the (mis)match between specifications and implementations, while analogy

techniques could be used to look for existing specifications, components, or implementations that match some new requirement. Textual analysis and information-retrieval techniques might be used to keep links between informal documentation and formal specifications and representations of development processes.

IPCs could also assist with many of the problems that arise when using a complex system, including diagnosis, planning, and operational tasks. An IPC for a transportation system, for example, could add significant value to the operational control of ships, trains, trucks, or airplanes. During both normal operations and emergencies, the IPC could monitor information derived from the increasing array of electronic sensors in the control room or cockpit, providing guidance and advice based on previous experience to the captain, driver, or pilot. IPCs can also support safe operation of complex processing systems, such as chemical plants and refineries whose interactions can defy even the most skilled human controller during times of crisis.

Enabling IPCs to reason about their task, environment, and team partners requires significant improvements in our technology for representing and reasoning about designs, plans, and goals; in particular, it demands that systems be able to reason about designs at multiple levels of abstraction (see the Coordination of Perception, Planning, and Acting subsection). Assisting design evaluation requires advances in techniques for simulating and projecting possible outcomes (see the Reasoning and Representation subsection). Enabling IPC systems to communicate using familiar means will require advances in natural-language processing, image understanding, and the cognitive aspects of communication with humans (see the subsections Perception and Human-Computer Communication in Multiple Modalities). Enabling these systems to improve their performance and exploit past experience will require improvements in machine learning (see the subsection Learning, Information Elicitation, and Automatic Adaptation). Final-

ly, enabling the systems to function effectively within organizations will require advances in capabilities for reasoning about social organizations, collaborative behavior, rules, and regulations (see the subsection Coordination and Collaboration).

Robot Teams

A focus on robots working in teams, allows for solutions in which knowledge, expertise, and motor capability may be distributed in time and space. While individual robots may have only limited capacities, together in groups they might be able to perform complex tasks. Teams of cooperative robots could assist society in a variety of ways.

- Teams of smart mobile vehicles could be taught to collaborate in surveillance tasks at factories or in the military. They would share information among themselves and with humans, and would distribute parts of the task.
- Robots—particularly mobile ones—have natural applications in dangerous environments, such as those encountered in environmental cleanup, mine removal, and planetary exploration.
- With a gradually aging population comes a market for automated household assistants—even limited ones that are capable of little more than fetching, opening doors, and doing basic cleaning tasks. Because of variability of the home environment, such robots will require considerable intelligence so that they can interact naturally and flexibly with humans.
- Large-scale laboratory experiments, like the Human Genome Project, will be made far more practical by automating some of the sophisticated laboratory procedures that require sensing, manipulation, planning, and transport.
- Manufacturers will make use of assembly robots that can be adapted easily to new tasks and delivery robots that can operate in a dynamic environment with minimal instrumentation. This will provide economies of scale benefits without

the need for a major investment in specialized factories.

The capabilities necessary to realize such systems go well beyond traditional computer-controlled machines such as the industrial robots used in automobile spot-welding. These new robots must be able to move safely and effectively in natural human environments; determine which obstacles can be moved, which can be avoided, and which can be moved over without causing damage; and coordinate with or react to other agents. The hardware for such advanced robots is being developed today, but the ability to employ them in unaided situations or in effective teams remains to be developed.

AI techniques can enable robots to evolve from objects that must be strictly controlled to objects that can be managed. Advances in the traditional robotic disciplines will be necessary (to provide more affordable sensors, for example) but these will not be sufficient. Progress requires more sophisticated perception, integration of input from a range of sensors, and the joint use of symbolic and sensory information (see the Perception subsection); enriched capabilities for robots to communicate with each other and with humans (see the Human-Computer Communication in Multiple Modalities subsection); abilities to plan individual and collective actions and to monitor and control their execution (see the subsections Coordination of Perception, Planning, and Acting, and Coordination and Collaboration); and the ability to acquire new behaviors by learning or by being told (see the subsection Learning, Information Elicitation, and Automatic Adaption).

Research Thrust Areas

The focus of this report is strategic: to define an AI research agenda that will support the development of high-impact application systems like those described in the previous section. The challenges presented by these systems cannot be overcome by improved engineering alone. Solutions will require improved understanding of the processes to be created and modeled,

including better understanding of the processes that underlie intelligent behavior in people. The computational study of intelligence—one of the fundamental scientific challenges of our time—is key to this undertaking.

The seven research thrusts we identify cross traditional AI-research boundaries. They form a bridge between high-impact application systems and underlying research problems. Each represents a significant opportunity in AI. A substantial research investment now will provide a solid base for constructing intelligent systems and will result in considerable payoff in both the long- and shorter-term. The systems context provided by the high-impact applications is likely to spawn new problems and stimulate new kinds of work.

These research thrusts encompass many major issues in understanding the fundamental nature of intelligence, both human and machine. Extensive and ambitious as this research agenda is, like other strategically defined research, it needs the complement of research that is not strategically defined. Both the problems and the payoff of such research are difficult to predict. But a scientific understanding of information processes such as learning, reasoning, and perceiving could change our understanding of ourselves and the world about us, with consequences that are as difficult to foresee as those of other fundamental breakthroughs in science.

Each research thrust area contains a substantial base of existing techniques, tools, and well-defined problems from which to draw. In the sections that follow, we briefly characterize the problem, the research base, and the key scientific challenges for the future.

Learning, Information Elicitation, and Automatic Adaption

Systems that can generalize, learn from experience, and adapt to new circumstances have the potential to reach higher levels of performance than systems that must be modified manually to deal with situations their designers did not anticipate. Virtually

all high impact application systems can be more powerful if they can learn from experience. For example, information assistants that can learn will be able to tailor their information-retrieval process to a user's needs without having to be told exactly what to do; they will instead generalize from previous interactions with the user. Learning skills will enable an intelligent coach system to deal with new types of problems, for example, drawing on its experience in the design of one type of automobile and applying it to the design of another. Networks of robots and computer-based agents in simulated worlds can avoid future coordination problems by learning from their experience in interacting with other robots.

Basic research has steadily advanced the fundamental technology of machine learning for more than two decades. A wide variety of learning methods—including decision-tree induction, neural networks, genetic algorithms, explanation-based learning, and case-based reasoning—have empirically demonstrated their utility on a broad array of real-world problems, from real-time evasive maneuvering to prediction of protein secondary structure. Significant progress has also been made by theoretical computer scientists in mathematically characterizing the scope and computational complexity of such algorithms.

Areas that have recently generated excitement in the machine-learning community include goal-directed learning, in which programs make decisions about what, when, and how to learn; practical methods for learning in the presence of a significant number of irrelevant features; the use of knowledge the system already has to improve the quality of learning; use of machine-learning techniques for scientific discovery and other kinds of data mining; the integration of learning with planning, language processing, and perception-action; and active learning, in which programs design experiments and other information-gathering activities that supplement the analysis of presented data.

Coordination of Perception, Planning, and Acting

Intelligent systems must be able to plan—to determine appropriate actions for their perceived situation, then execute them and monitor the results. Planning, in turn, requires advanced capabilities to represent and reason about time, action, perception and the mental states of other agents. To cope with realistic situations, systems must be able to deal with incomplete, uncertain, and rapidly changing information and must have mechanisms for allocating resources between thinking and acting. Information-resource systems, for example, need to plan the best way to acquire information, trading off the urgency of the request against the cost of accessing different databases, the expected time required, and the likelihood of success. Simulated agents, intelligent coaches for operating complex systems, and taskable robots will all have to cope with complex physical processes and situations in which the actions of other agents conspire to create a complex, unpredictable, dynamic environment.

Basic research in planning has provided a substantial base on which to develop intelligent planning capabilities. Expressive action-representation languages have been devised and given precise semantics. The formal bases of planning are well understood, as is the computational complexity of the problem in general. A variety of algorithms have been developed for constructing plans to satisfy a given set of goals. Learning techniques have been applied to reduce the time planners take to solve problems by enabling them to effectively apply previously derived solutions to new problems. Recently, a new class of planning systems was developed that combines perception, planning, and action and guarantees a response in bounded time. These *reactive planning systems* function in dynamic worlds to which they are connected by their perceptual system; they are more easily linked to traditional control mechanisms for the low-level operation of effectors. Practical systems that have been crafted to take advantage of domain-specific con-

straints can automatically develop plans consisting of thousands of actions, both sequential and parallel, in domains such as logistics, assembly, and manufacturing.

Research in this area currently focuses on several key challenges. One focus is managing the trade-offs among acting, planning, and acquiring further information to reduce uncertainty. For example, will interrogating more databases improve the quality of the answer that an IRSS needs to deliver within two hours without raising the cost prohibitively? In many applications, an extra moment's thought might lead to a better plan, but each delay can also make the problem more difficult. Techniques have been developed that combine decision- and game-theory techniques with classical AI techniques, but significant representation and efficiency questions must still be addressed.

Another research focus grows out of the differences between reactive and classical planners, which excel at different types of tasks. Some intelligent systems will need to perform both types; hence, another area of investigation concerns the appropriate partitioning between reactive behaviors and deliberative behaviors and the development of techniques that integrate planning, perception, information seeking, execution, and plan modification. Other investigations aim to examine ways in which systems can effectively decide among conflicting goals and to analyze ways in which systems can effectively build models of the world, apply them to situations, and modify them from experience.

Coordination and Collaboration

The ubiquity of computers, networks, and distributed information resources means that collaboration is itself ubiquitous. For example, the anticipated national information infrastructure (NII) environment is too large, complex, dynamic, and open to be managed centrally. The people using the NII will be as diverse as the citizenry of this country: they represent many backgrounds, levels of comput-

er experience, and cultures. Consequently, the systems that operate on this infrastructure must be flexible in the ways in which they work and communicate with users. In almost every undertaking, multiple intelligent agents will need to collaborate both with other intelligent systems and with people to provide the services and information needed.

All of the high-impact application systems will require both human-computer and computer-computer collaborations. Machines can be designed to follow fixed protocols, but any collaboration that involves people will have to take into account the ways that people work. The amount and types of communication between agents will also vary widely, depending on both the types of participants and the costs of communication.

Research in AI has already produced a number of models of cooperative and collaborative behavior as well as techniques for allocating tasks and resources among multiple agents and negotiation protocols for coordinating their activities. A range of organizational structures have been examined, from tightly coordinated distributed systems to spontaneous partnerships of independently designed and self-motivated systems. Practical success has been achieved in areas such as distributed sensing and planning and network diagnosis and management, in which agents with a limited set of functions can be made to work together using relatively simple communication protocols. Substantial theoretical understanding has been achieved with the methods that systems can use to represent and reason about other systems' objectives, capabilities, and priorities and about the information that must be communicated to guarantee successful collaborative action. Representation languages are being developed to support collaborative processes such as concurrent engineering. Some steps have been taken in the commercial sector (such as Telescript, OLE, and CORBA), but these steps address only low-level coordination; they cannot adapt to changes in the environment and are consequently unable to automatically incorporate new, indepen-

dently designed resources.

One active area of current research is the development of methods for more effectively reasoning about other agents' abilities and the value to them of different actions or goal states. Another area is the design of techniques for task and resource allocation in different organizational structures, for example, negotiation protocols, contracting methods, and other communication languages and standards that enable agents that were developed separately to collaborate as part of one system. A key question is how to take account of the reliability of the information a system receives from other agents and the confidence it can have that another agent will keep its commitments. Other investigations in the area of collaboration consider the cost-benefit tradeoffs in communication. Communication, although often valuable, is not always cost free: in dynamic situations, such as realistic simulated worlds, the time cost of communication must be considered.⁵ A final research issue concerns remote and distributed control of physical agents, for example, in teams of robots. The unpredictable delays in networks that can be tolerated for many text and vision applications are unacceptable when motion-control information is transmitted.

Perception

Most of the high-impact application systems require an ability to handle several types of perceptual information. For example, teams of robots need vision, language, and touch capabilities to function realistically. Design associates will need to interact with their environments and other team members; operational associates will need to monitor the behavior of complex systems using a variety of sensing devices. Humanlike communication that will make computers more accessible to everyone will require advances in perceptual capabilities, such as image interpretation, gesture recognition, and spoken-language understanding.

The ability of computer systems to perceive and communicate has evolved dramatically over the past

decade as a result of research in AI and related disciplines that address issues of human and machine perception. Large-vocabulary, discrete-phrase speech recognition is commercially available; several laboratories have developed speaker-independent real-time continuous speech-recognition systems for tasks requiring several thousand word vocabularies. These systems complement advanced natural language-processing techniques, which now support automated clipping services for categorizing newspaper stories, as well as partially automated translation of technical manuals into foreign languages. Automatic vision systems are used commercially for inspection of manufactured parts, and semiautomatic systems for various image-analysis tasks should be available soon. Critical capabilities such as real-time stereo analysis are also within reach.

One central challenge in all areas of perception is to increase the range of signals that can be interpreted, for example, understanding unrestricted outdoor scenes as opposed to known industrial parts, or naturally occurring speech as opposed to *read* speech. A second challenge is to increase the accuracy of the interpretation process. A third is to enable real-time perception with acceptable accuracy. The methods being investigated in the perception community include using more sources of information, and designing automatic training methods that work alone or in combination with handcrafted rules and models. Image understanding techniques are being developed to interpret multiple views of the same scene or event, for example, in a video of an object in motion. Symbolic rules and models can be augmented by methods that learn automatically from data the likelihood that a rule or model component will be applicable in a given situation. These techniques, which take advantage of informative statistical patterns that humans cannot reliably detect, improve the robustness of the interpretation process and decrease the time necessary to adapt a perceptual system to a new domain. A final research challenge that is central to all the perceptual modalities is how

to coordinate symbolic methods with nonsymbolic ones (for example, stochastic methods or neural networks). Research has reached the stage where significant advances in combining the best features of both approaches look to be within reach.

Human-Computer Communication in Multiple Modalities

Communication among people is marked by its flexibility, from the casual nod of a passerby conveying a greeting, to a high-school teacher's math lecture with its complex interaction of lecturing, drawing diagrams on a chalkboard, and answering questions. People use a number of different media to communicate, including spoken, signed, and written language; gestures; sounds; drawings, diagrams, and maps. The high-impact application systems must also be able to understand the full range of communication media.

For example, the editor of the six o'clock news or an intelligence analyst might request a video clip containing the man to the left of Mandela, while pointing to a photograph in an on-line magazine. The photograph's caption might identify the individual as deKlerk; alternatively, his face might need to be compared to a library of prominent South Africans. Similar challenges arise in teams of robots. To process the statement, "That's where the repair kit is kept," said with a finger pointing to a cabinet, a robot must combine understanding of the utterances of the (human) trainer with interpretation of the images it is currently receiving from its cameras. For humans to participate in activities using simulated environments, capabilities such as gesture and expression recognition will be required.

Intelligent systems must also be able to convey information using all the communication media. For example, an IRSS helping a user find a new home might provide information using a combination of maps, diagrams, text, and spoken descriptions. These various media need to be combined so that information is communicated in the manner most appropriate

ate to the particular user and task at hand.

Articulate intelligent systems require integrating and using multiple modalities. Interpretation and synthesis processes in individual modalities are subject to a certain degree of error; even humans misunderstand each other. The joint use of multiple modalities permits one modality to compensate for interpretation errors of another. For example, a speech recognition module might do poorly on infrequent proper names, but these names are easy to write with a pen.

The ambitious systems envisioned here are, of course, a long way off, but even limited progress in broad-band human-computer interaction has the potential for large payoffs. Techniques for fusing multimodal input could serve as the basis for simpler interfaces that allow the user to combine speech, mouse, and keyboard input, using each where it is most convenient. Advances in media coordination could allow for more efficient interactions between computer and user in the standard computer applications of today, for example providing guidance to a spread-sheet user through coordinated audio and highlighting of cells. Simple models of how people perceive maps and diagrams could lead to automatic mapping systems or graphic design tools that would allow a normal computer user to adequately perform such tasks in the absence of a highly trained graphic designer.

Research in this area can start from a solid base in vision, facial gesture modeling, speech recognition and synthesis, natural-language processing, and automated design of informational graphics and animations. Central challenges for interpretation include developing representations that enable combining information from different modalities and developing techniques for synchronizing different interpretation processes. Central challenges from the production side include managing the content to be conveyed so that it is appropriate to the media available, apportioning it correctly for conveyance by the appropriate medium,

and synchronizing the multiple media components.

Content-Based Retrieval

The Internet is already populated with enormous amounts of multimodal information, from pages containing images, text, and graphics to video with sound track. This wealth of information will grow ever more extensive when the NII becomes a reality. IRSSs will need to provide access to a wide variety of information, including visual and audio data, in addition to commonplace structured databases.

Any access to these materials beyond the simple keyword and hypertext browsers now available will require automatic indexing schemes that work across multiple modalities and will require capabilities for content-based retrieval. The Mandela-deKlerk photo query in the previous section provides a simple example. A successful reply requires that clips in the video library be scanned using both symbolic content (such as captions) and visual information (face recognition). Recognition of moving images presents a yet more impressive challenge, and substantial benefit. Consider for example, a medical student who wants to see an example of a particular kind of suturing technique. The relevant video clip may be filed under a different category, such as the overall procedure being performed (such as ulcer surgery). The student needs an intelligent system capable of looking through the video library to find a clip that offers a good illustration. Finding even potentially relevant clips might require significant reasoning, such as first narrowing the search down to clips of surgeries, then to particular types of surgeries, and then perhaps to particular physicians. Then the candidate clips would need to be reviewed to determine which ones show the type of suturing in question, and finally selecting the best example.

Existing language-processing and vision techniques, mentioned in earlier sections, can provide a starting point for intelligent, content-based retrieval of information. Significant research challenges include determin-

ing the kinds of image annotations, video, and audio data that are needed to enable efficient and effective access; developing techniques for automatically processing raw data to produce these annotations; providing a means of representing multimodal queries, whether in query languages for users or as target translation languages for sophisticated human-computer communication systems; developing capabilities for performing these tasks quickly enough so that users can afford to search many images or videos; and integrating multiple access techniques.

This research thrust, possibly even more than the others, requires coordination of AI researchers with those in other areas of computer science, notably database and network experts. AI can provide a significant body of experience in the process of task and domain modeling and in interpretation methods. This research thrust can also benefit from the results of work in learning and perception. Progress in this area will likely impact, and benefit from, other areas of AI concerned with interpretation from various sources, such as diagnosis.

Reasoning and Representation

Research in reasoning and representation is needed to support the full range of high-impact applications systems. For example, in software engineering, AI representation and reasoning techniques can be used to describe the interfaces of software components and to find how to connect components from different sources to achieve a complex goal. In computer engineering, logical reasoning techniques are crucial in specifying and verifying complex digital systems such as telecommunication protocol chips, while advanced simulation systems require qualitative reasoning techniques to avoid the computational bottlenecks of solving large systems of equations in simulations of complex physical systems.

In any realistic problem, including those that arise in the high-impact applications systems, reasoning must be done under less than perfect conditions. Intelligent information systems must deal with data that is

imprecise, incomplete, uncertain, and time varying. They must be able to manage with domain knowledge that is incomplete, and they must do as they meet pressing real-time performance requirements. Finding a solution that is guaranteed to be optimal—under any reasonable interpretation of *optimal*—can be shown to be computationally intractable: it cannot be done efficiently no matter how much faster we make our computers. Consequently, we must develop fast heuristics that can be shown to lead to good—if not necessarily optimal—solutions.

The IPC for transportation illustrates the typical challenges faced by a reasoning system. Sensor data providing the system with recent information about traffic are bound to be imprecise and, from time to time, unreliable because of sensor failures, drifts, or extreme operating conditions such as overheating. This incomplete and vague data must be reconciled, integrated with available statistical information, and analyzed to identify trends and situations that require corrective actions such as changing the timing of traffic lights. Decisions must be made quickly and in a way that can be justified to the end-user.

AI research to date has partially addressed these issues by developing many specialized reasoning techniques, including *anytime reasoning*, techniques for enabling a system to reach the best possible conclusion within the time available; *nonmonotonic reasoning*, techniques for leaping to conclusions based on partial information in a justifiable way that allows conclusions to be withdrawn if necessary as new information comes in; *case-based reasoning*, techniques for using previously acquired solutions to old problems as the basis for new solutions to new problems; and *Bayesian networks*, a technique for using causal and probabilistic information efficiently.

Each of these techniques—and others developed by the AI community—works only in limited domains. Further research is needed to extend the scope and efficiency of these techniques and to integrate them. More-

over, given that finding the optimal solutions is beyond our capabilities (or any capabilities we can hope to develop), we need to better understand the degree to which the solutions these techniques provide approximate the optimal solutions and the conditions under which a technique can be used safely. Progress in these areas should not only help us build better systems, but should also increase our understanding of how humans manage to deal with complexities as well as they do.

A variety of representations that capture information at multiple levels of abstraction and in different degrees of detail will be needed to deal effectively with complex systems. For instance, the most abstract level will represent the core conceptualization, providing information about the way an artifact accomplishes its goal. Systems will be able to reason quickly, but only imprecisely, with representations at this level. More specific representations will encode additional detail and enable more precise reasoning, but at greater computational cost and with increased difficulty in interpretation. AI research to date has produced a substantial repertoire of representation techniques. Research is needed to identify the levels of representation appropriate for modeling different types of complex systems, and to develop reasoning techniques that support combining representations of different devices at different layers into effective models of a complex system.

Building Large Systems

Achieving the goals set forth in the preceding sections requires integrating multiple capabilities. The high-impact applications require solutions that retain efficiency and robustness in large-scale, demanding environments. This section describes several needs that pervade both the research thrust areas and the development of these applications.

Better programming tools: For many years AI researchers developed their own programming environments, which typically were years ahead of their time. However, this is no longer

the case. The computer industry, having taken over the role of tool developer, has largely ignored the needs of advanced computing researchers, instead focusing on less ambitious, but more profitable, markets. Better programming tools are of enormous importance if we are to build large-scale systems that integrate multiple capabilities.

Sharable resources: The advances in AI technology necessary for large-scale applications cannot be achieved by individual researchers working alone. Instead, researchers must build on the work of others, working collaboratively. We need to develop and maintain large-scale knowledge-bases and program libraries and to create knowledge representation capabilities that will allow these shared resources to be used successfully.

Common-sense knowledge: Brittleness has been a perennial problem with intelligent systems constructed to date: they are good at their task but their performance falls off drastically as they move away from that task. Human expertise is far more flexible; it rests on a large stock of common-sense knowledge about the world, a very large collection of basic facts and inferences. A substantial common-sense knowledge base would lend an important improvement to the performance of many systems.

Significant progress in this area depends on the development of improved knowledge representation and reasoning techniques (see the section Reasoning and Representation), and is likely to have pronounced influence on research in this area in return, because “plausible” reasoning methods in several domains—especially temporal and spatial reasoning—are likely to incorporate aspects of common-sense knowledge. Moreover, attempts to encode large-scale common-sense knowledge effectively will continue to provide demanding tests of the representational power and effectiveness of formalisms.

Semantics of module composition: Large systems cannot be built simply by composing fragmented capabilities for individual technical problems. Constructing high-impact application systems will require integrating many

capabilities from each of the research thrust areas we have described. Because each of these capabilities is itself typically realized in a complex software system, the integration of multiple intelligent capabilities into yet larger systems will stretch the limits of current technology. To succeed, we need to better understand how to model the capabilities and limitations of each module, the ways groups of modules can be combined, and how they interact when so combined both through their interfaces and through the more subtle constraints the modules impose on each other's internal workings.

Support for integrated systems: Working on integrated systems typically demands a team approach. Although past efforts have been located in a single geographical location, increased network capabilities make it important to consider the tradeoffs in having distributed research teams and the support needed to make them.

Experimental techniques for large systems: The experimental evaluation of large systems requires new approaches to enable designers to evaluate the abilities and limitations of individual components as parts of larger systems. Overall performance figures for complete systems are useful indicators of progress, but in general are not reliable predictors of the performance of individual components in new systems or tasks.

Interdisciplinary research: High-impact applications require coordinated efforts of research and development across areas of computer science. Building these systems will require combining AI methods with non-AI approaches and embedding AI technology within larger systems. In addition, many of the fundamental scientific challenges require collaborative, interdisciplinary efforts in the cognitive sciences and engineering.

Education: Much of the transfer of AI techniques to applications that answer societal needs will occur through students, trained in universities and research centers, who join projects in the computer industry and start new companies. Sustaining the kinds of projects needed for high-impact application systems will

require a strong community of AI researchers and practitioners. University and other research laboratories and centers of excellence are a vital part of the infrastructure. Support for education at the masters, Ph.D., and postdoctoral level is crucial.

Many important advances in computer science—including the development of time-sharing, compilers, massively parallel computers, and object-oriented programming—evolved from efforts to support AI research. Creating the infrastructure required to produce integrated AI systems has acted as a forcing function for advanced computer systems development. As AI is applied in more complex and demanding areas of software development, such as the NII and large-scale commercial and military systems like DART, this impact seems likely to continue, producing breakthroughs that benefit the entire computer industry.

AI is also becoming an enabling technology for software applications that are not traditionally thought of as involving intelligence. Televisions that learn their owners daily viewing habits illustrate this trend, as do more important invisible uses of AI techniques, such as those in the Apple Newton. Large software systems can often benefit markedly from the integration of a little AI, a process which will continue rapidly as its profitability becomes more widely perceived. Such integration increasingly requires a deeper understanding of many basic issues in the foundations of AI.

Conclusion

National competitiveness depends increasingly on capacities for information analysis, decision making, and flexible design and manufacturing. Strength in these areas was once limited by insufficient data, lack of computational power, or inadequate control mechanisms. Many critical limitations, however, can now be overcome only by adding intelligence to systems.

Basic research in AI will, in the long run, contribute not only to our scientific knowledge but also to our technological base and to a wide variety of

applications. It will provide the foundation for systems that can search large bodies of data for relevant information; help users to evaluate the effects of complex courses of action; and work with users to develop, share, and effectively use knowledge about complex systems and processes. It will make it possible to build a wide range of application systems that assist decision makers in adapting and reacting appropriately to rapidly changing world situations.

Notes

1. This article contains the full text of the report to ARPA on twenty-first century intelligent systems. Additional copies of this report are available from the AAI office.
2. The general role of many computing areas in addressing national needs has been described elsewhere [1, 2].
3. Extensive examples of working systems can be found in the annual *Proceedings of the Innovative Applications of Artificial Intelligence Conference*, sponsored by AAI.
4. ARPA won a Gold Nugget award based on that success, and Vic Reis, the then-current ARPA Director was subsequently quoted as saying that DART justified ARPA's entire investment in artificial-intelligence technology.
5. Many of these problems arise in computer systems research as well. Although the constraints are typically different, this is an area in which interdisciplinary research and cross-fertilization are likely to be beneficial.

References

- [1] Committee on Physical, Mathematical, and Engineering Sciences (CPMES) and the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET). *High Performance Computing and Communications: Toward a National Information Infrastructure*. Washington, D.C.: Office of Science and Technology Policy, 1994.
- [2] IITA Task Group. *Information Infrastructure Technology and Applications*. Washington, D.C.: Office of Science and Technology Policy, February 1994.