7.0 Artificial Intelligence

Introduction

The Panel's projection of the future of the software technology called Artificial Intelligence (AI) was greatly helped by the timely appearance of a paper on the subject, prepared by a distinguished group of AI scientists, co-chaired by SAB and panel member Davis and organized by the American Association for Artificial Intelligence. The first order of business for the reader who wants to better understand the future of AI is to read this paper ("A Report to ARPA on Twenty-First Century Intelligent Systems", AI Magazine, Fall 1994). We quote liberally from this paper in the material below.

Advances in computers and telecommunications have made a vast quantity of data available to us, and given us computational power that puts the equivalents of mainframes on the desktop. However, raw information processing power alone, like brute strength, is useful but insufficient. To achieve their full impact, systems must have more than processing power—they must have intelligence. They need to be able to assimilate and utilize large bodies of information, to collaborate with people and to help them 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. The concepts, techniques, and technology of the IT area called Artificial Intelligence offer 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 intelligence in software.

Below we describe AI applications and underlying technology that will enable intelligent systems to meet Air Force needs in the next five to twenty years. We will refer to these applications as "high-impact application systems."

7.1 Intelligent Simulation Systems

Elsewhere in this volume (See Modeling and Simulation, Chapter 9) we offer a view of the future of modeling and simulation technology. A new generation of *intelligent* simulation capabilities will support the construction of programs that model complex situations, involving both complicated devices and significant numbers of intelligent simulated people. The simulated worlds that can be generated today have limited physical realism and severely lack realism in their simulations of people.

7.1.1 Artificial Simulation Actors

The systems of the future will differ in both scale and function from those that exist today. In the next generation or two of simulations, thousands of "actors" will play roles. It might be economical to use actual people for only a few of these roles; the rest could be simulated using AI techniques.

A key challenge is constructing realistic humanlike actors. In the future, these actors will be able to coordinate perception, planning, and action (discussed later), learn, understand and interact with their world, deal with other actors, and use natural language. 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.

7.1.2 Simulation in Engineering

Another quite different use for intelligent simulation will be for an advanced form of engineering design and evaluation. AI programs can be given the knowledge and reasoning power to use the physics and engineering principles underlying the design of artifacts. Such programs will speed up and make more accurate the formulation of design models that can be tested by conventional numerical simulation. An evaluation environment for new products, such as vehicles or airplanes, could use simulations of people to test the feasibility of a product's construction, use, and maintenance before it has been built. A new product design could be "used" by simulated people while it has only a virtual existence. Potential customers could try out the product in a simulation.

7.2 Intelligent Information Resources

Information-resource specialist systems will support effective use of the vast resources of the national information infrastructure. These systems will 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 will adapt to changes in users' needs and abilities as well as changes in information resources. They will be able to communicate in human terms in order to assist those with limited computer training. These systems constitute an important class of *intelligent agents*, discussed in Chapter 5.

7.3 Intelligent Associate Systems

Software designed to act as an intelligent, long-term team member could help to design and to operate complex systems. An *intelligent associate 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. Associate 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.

The Boeing 777 aircraft illustrates that some major advances in design technology have already taken place. New tools enabled designers to check spacing and clearance so accurately that a physical mock-up version of the plane was not needed. But these tools still had limitations. They did not incorporate, for example, vast volumes 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, has now been lost, and will be sorely missed when the design is revised in the future (as all designs are).

7.3.1 Intelligent Help for Information Overload

Sensor and communication systems provide the warfighter with a wealth of data for decision making. In the future this wealth threatens to be overwhelming. The clearest uses of Intelligent Associates will be to assist individual users and teams to gather, cull, organize, and interpret data relevant to a situation. The Information Applications Panel discusses the future of Information Fusion. The AI technology to make this vision a reality will mature in the next 10-20 years.

A recent report by the Office of Science and Technology Policy noted that in 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. The translation to a warfighter's workstation is obvious.

7.3.2 Intelligent Help for Ease of Use and Communications

To realize this enormous potential, Intelligent Associates 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. Associates will allow the use of diagrams and gestures, combining media and modalities in whatever mix is best for getting the message across. The commands that users issue will be general and often vague; nevertheless the Associate must accurately determine how to perform such commands. The information that a user needs will often not be stored at any one site; thus the Associate will need to be able to access multiple sites and recognize common information (see Chapter 6 on mediators). To actively and continuously seek out useful information, an Associate will need to learn which topics are of long- and short-term interest to each user.

7.3.3 Intelligent Help for Organizational Processes

The Associate will remember and recall the rationale of previous decisions, and, in times of crisis, explain the methods and reasoning previously used to handle that situation. Intelligent Associates will incorporate intelligent simulation and information resources systems as components.

For example, an Intelligent Associate for aircraft design will enhance collaboration by keeping communication flowing among the large, distributed design staff, the program managers, the customer, and the subcontractors. It will 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.

7.3.4 Intelligent Help for Software Development

One critical area in which Intelligent Associates will assist is software development: keeping track of specifications, design proposals, and implementations for a software project throughout its life cycle; recording the design decisions of a constantly changing team; and being a repository of solutions and components for new projects. For the new architecture-based software development, interactive AI methods will be used to instantiate requirements and specifications as a bridge to an automatic coding process, and bring many "special case exceptions" into the code (from a case library).

7.3.5 Intelligent Help for Finding Analogous Situations and Cases

The Intelligent Associate will use methods for reasoning by analogy. Analogy techniques could be used to look for existing specifications, components, or implementations that match some new requirement.

7.3.6 Intelligent Help for Managing Complexity

Intelligent Associates will assist with many of the problems that arise in using our ever-more-complex systems, including diagnosis, planning, and operational tasks. For example, they will add significant value to the operational control of air vehicles and weapons systems. During both normal operations and emergencies, the Associate will monitor information derived from sensors in the control arena or cockpit, providing guidance and advice based on previous experience to the warfighter.

7.4 The AI Technology Underlying High-Impact Applications

A common core of capabilities is needed to construct AI applications. These include:

a. Abilities to reason about the task being performed with the knowledge that is appropriate to the task.

b. To reason about the collaborative process and the knowledge and capabilities of other systems and people participating in an interaction.

c. To communicate with users in human terms, producing and understanding combinations of spoken and written language, drawings, images, and gestures.

d. To perceive the world.

e. To coordinate perception, planning, and action.

f. 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 by industry and the Armed Services, as well as many other applications in planning, learning, perception, and language processing.

7.4.1 Learning, Automatic Adaptation

Virtually all high impact application systems can be more powerful if they can learn from experience. For example, Intelligent Associates 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 Associate to deal with new types of problems, for example, drawing on its experience in the design of one type of UAV and applying it to the design of another.

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. There have been significant advances. These include:

a. Goal-directed learning, in which programs make decisions about what, when, and how to learn.

b. Practical methods for learning in the presence of a significant number of irrelevant features.

c. The use of knowledge the system already has to improve the quality of learning.

d. Use of machine-learning techniques for scientific discovery and other kinds of data mining.

e. The integration of learning with planning, language processing, and perception-action.

f. Active learning, in which programs design experiments and other information-gathering activities that supplement the analysis of presented data.

In the future, neural networks will develop to be powerful pattern classifiers and will be used as "front-ends" for symbolic reasoning programs. Genetic algorithm methods will slowly develop as relatively weak search methods to aid machine learning.

7.4.2 The Plan-Decide-Act-Monitor Cycle

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.

7.4.2.1 Planning

Basic research in planning has provided a substantial base on which to develop intelligent planning capabilities. 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 constraints can automatically develop plans consisting of thousands of actions, both sequential and parallel, in domains such as logistics and battle planning. New capabilities will manage the trade-offs among acting, planning, and acquiring further information to reduce uncertainty.

Since the technology for planning from AI and from Operations Research is highly developed, we can forecast that over the next 10-20 years most human decision making involving complex sequences of actions and parallel courses of actions will be assisted by (at least) semi-automatic computer planning

7.4.2.2 Perception and Language

The ability of computer systems to perceive and communicate has evolved dramatically over the past decade. 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 will be rapidly commercialized. We can expect to see in the coming decade, highly effective continuous speech understanding systems, with tens of thousands of words, with error rates not exceeding 1-2%. Such systems will ride the wave of increasing computer power available cheaply.

The systems will 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. Such applications for automatic or partially-automatic natural language understanding will become commonplace in 10-20 years.

Growth in the coverage (scope) of natural language understanding, and its reliability, will track the growth of the large knowledge base described later, with high performance capability in specialty areas occurring in five years. However, for unconstrained (but unsophisticated) human discourse, the time frame is more like 20 years.

7.4.2.3 Perception

Significant technical progress will 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 hand-crafted 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 will 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, will improve the robustness of the interpretation process and decrease the time necessary to adapt a perceptual system to a new domain.

Finally, 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 the technology will occur that will allow the combining of the best features of both approaches.

7.4.2.4 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 professor'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, an Intelligent Associate helping a warfighter might provide information using a combination of maps, diagrams, text, and spoken descriptions. These various media will be combined so that information is communicated in the manner most appropriate to the particular user and task at hand.

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.

Broad-band human-computer interaction has the potential for large payoffs. Techniques for fusing multimodal input will serve as the basis for simpler interfaces that allow the user to combine pictures, speech, mouse, and keyboard input, using each where it is most convenient. It is either ironic or amusing that we will be using one of our most powerful (software) technologies to simplify the use of our complex IT systems.

7.4.2.5 Finding Something by Its Content

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 and the DII (Defense Information Infrastructure) become realities. Intelligent systems will 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. Recognition of moving images of objects in video material, a substantial benefit to analysis, will be required in the future and will be available.

7.4.2.6 The Power to Reason

In any realistic problem, 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 so 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, i.e., cannot be done efficiently no matter how much faster we make our computers. Consequently, we must develop fast methods for plausible reasoning that can be shown to lead to good—if not necessarily optimal—solutions.

Sensor data providing the system with recent information may be imprecise and, from time to time, unreliable because of sensor failures, drifts, or extreme operating conditions. This incomplete and vague data must be reconciled, integrated with available statistical information, and analyzed to identify trends and situations that require corrective actions. 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:

a. "Anytime" reasoning, techniques for enabling a system to reach the best possible conclusion within the time available.

b. Nonmonotonic reasoning, techniques for leaping to conclusion based on partial information in a justifiable way that allows conclusions to be withdrawn if necessary as new information comes in.

c. Case-based reasoning, techniques for using previously acquired solutions to old problems as the basis for new solutions to new problems.

d. Bayesian networks, techniques for using causal and probabilistic information efficiently.

Techniques based on probability or inexactitude fall into a class that the Japanese now call "soft" (as opposed to the purely logical "hard"). Also in this class are "plausible" reasoning methods based on heuristic knowledge, and reasoning based on "fuzzy" set definitions with membership ranges. These methods will probably be the most widely used methods of reasoning in 10-20 years (vs the "hard" methods). For the larger world of computer applications, these "soft" reasoning methods may become more important than calculation (numeric computation).

7.4.2.7 Representation

A variety of representations that capture information at multiple levels of abstraction and in different degrees of detail will allow programs to reason effectively about complex systems. For instance, the most abstract level will represent the core conceptualization, providing information about the way an artifact accomplishes its goals. Programs will be able to reason quickly, but only imprecisely with representations at this level. More specific representations will encode more details and enable more precise reasoning, but with greater computational cost and increased difficulty in interpretation.

7.5 The Scaling Up to Large AI Systems

The most important limit to the intelligence of current AI systems is their narrow scope and shallow depth. These systems have not been given (by us) enough knowledge to assist us in the great variety of our tasks. Nor have the systems yet learned large bodies of knowledge by machine learning methods. In the next decade, major advances will be made in producing an international distributed knowledge base of the "widely-shared knowledge" of our society and our science. (Think of this as the "Knowledge Web" in analogy to the World Wide Web.) The knowledge web will be the backdrop to all AI systems. Networks will be the medium by which this knowledge is accumulated and distributed. The commonly held view that "you have to tell a program everything" in order for it to perform properly will become obsolete. The "widely shared knowledge" base coupled to reasoning programs will supply necessary but missing detail.

Brittleness has been a perennial problem with the thousands of expert 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 the previously mentioned widely shared knowledge about the world. The large knowledge base will solve the brittleness problem for most AI applications by providing this "fall-back" knowledge.

In the 10-20 year future, we believe the knowledge web will contain tens of millions of objects, rules, and logic formulas (perhaps hundreds of millions). A variety of reasoning packages will be available to plug-and-play with the knowledge web; and will be easily customizable to the users' needs.

The combination of natural language understanding methods, machine learning methods, and the knowledge web may give us a powerful surprise. Knowledge-based learning methods may grow the knowledge web in a "bootstrapping" way by understanding natural language text (similar to how we come to know things by reading about them). For specific domains of application, such a combination will reduce the domain knowledge acquisition time by factors of ten to one hundred.

7.6 AI Technology Points of Interest

The technical requirements of significant Air Force applications are considerably broader than AI technology alone can provide. However, AI capabilities will be key to making Air Force systems intelligent, adaptable, far more accessible to the relative unskilled user, and, thus, dramatically more effective.

Providing AI capabilities is not an all-or-nothing proposition. Although the development of systems with very sophisticated capabilities will require long-term effort, in each category of application more restricted but still usefully intelligent systems can and will be developed.

7.7 Interdisciplinary R&D

High-impact AI applications require coordinated efforts of research and development across the several 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.

7.8 Summary

Artificial Intelligence technology 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. AI 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.