

Artificial Intelligence

Darrel Patrick Wash
1989

Do you want an omelet or pancakes for breakfast ? Toast or a muffin ? Just place your order with Robobutler and then go take your shower. When you are ready to eat, your personal robot will serve your meal, clean up the kitchen, and remind you of your day's business appointments. As this helpful little fellow hands you your briefcase at the door, he awaits your instructions for the day; you tell him to vacuum the floor, make the beds, do the laundry, and prepare dinner. Then you get in your car, tell it to drive you to work, sit back, and enjoy the morning paper.

If you think these machines sound too good to be real, you're right. Before Robobutler and other mechanical helpers can work the way we imagine, they will need to be able to understand human language, see and react to what they see the way a person would, and perform many other human-like actions, all of which are examples of artificial intelligence. But just what is artificial intelligence ? How much of this technology already exists ? And who is working to develop it further ?

Artificial intelligence (AI) is not defined in the same way by all who use the term because of disagreement over what intelligent behavior is. One faction within the AI community defines intelligence as the ability to cope with change and to incorporate new information in order to improve performance. Existing technologies don't appear capable of this. The broader view, however, is that artificial intelligence is that which mimics human reasoning or sensing. We already see examples of this capability in expert systems, industrial robots, machine vision, parallel processing, and neural networks, all of which are discussed below.

Because AI is hard to define, it is hard to state who works in it and to describe their occupations. Workers include researchers with advanced degrees and software designers who have a bachelor's degree in computer science with an emphasis on AI. Currently, relatively few people are engaged in developing AI products, no more than 8,000 according to knowledgeable sources. Rapid growth in demand for workers to develop this software is expected, but the number of new jobs in this area will still be relatively small.

Expert Systems

The most commercially successful AI application is the expert system. An expert, or knowledge-based, system is a computer program that acts as a consultant to decisionmakers. The program contains information on a particular subject, known as the knowledge base, that is most frequently represented as a series of "if-then" rules, although in some systems the knowledge is represented as frames, objects, or semantic networks. When applied to a problem, the system searches for solutions in the same logical patterns that human experts would use. The largest and most complex expert systems have over a thousand rules and may take 2 years or more to create. Development takes so long because, for the system to be usable, the expertise to be captured must be clearly defined, and all the steps a human expert would take to draw a conclusion must be spelled out beforehand. This is the job of the knowledge engineer, who is the link between the software developer and the end user.

Knowledge engineers are somewhat similar to systems analysts. They must not only know the capabilities of the computer hardware and software but also understand the user's operation and what he or she wants to accomplish through automation. They must also be able to deal effectively with people from many different backgrounds -- for example, medical researchers, engineers, financial analysts, and industrial machinery mechanics.

The knowledge engineer interviews one or more experts and then distills the information relevant to the application into a set of rules or some other form of knowledge representation that reflects the behavior of human experts. It can be difficult to get people to part with knowledge they have spent a lifetime acquiring. Even when the experts are eager participants, they can miss the importance of some subtle steps in their decisionmaking process.

When the knowledge representations are complete, a programmer creates the knowledge base by coding the information in an AI language. In some situations, the knowledge engineer does the actual programming. Finally, there is debugging, testing with trial problems, and refining.

Expert systems have been around since the early 1970's. One of the classic systems is MYCIN, which diagnoses blood infections. MYCIN represents the knowledge of physicians who are experts in this field, and it enables other, less knowledgeable, physicians to consult with these experts for help in diagnosis and treatment.

EMPTY MYCIN, a later development, represented a major advance for AI. It separated the information about blood infections from the inference engine, the rules that apply the knowledge. This breakthrough allowed MYCIN'S inference engine to be used with other knowledge bases and led to the development of expert system shells.

An expert system shell is, in effect, an expert system without the expertise. More technically, it is a set of utilities that programmers can integrate into their existing computer system. Users can then fill the shell with the knowledge needed for a specific application. Shell programs permit programmers to create expert systems on desktop computers using conventional languages.

Because shells are much less expensive and more flexible than customized expert systems, their development has caused the overall use of expert systems to skyrocket. No one knows precisely how many expert systems are currently in operation, partly because companies now recognize these systems as competitive weapons and some companies are unwilling to divulge information about them. Nevertheless, knowledgeable sources estimate that several thousand are in use, and the number is growing rapidly. Between one-half and three-fourths of the Fortune 500 corporations now use expert systems; many large corporations have hundreds of systems in use.

The popularity of expert systems is growing because more and more organizations are finding that they improve quality, raise productivity, cut costs, and increase profits by helping their employees "work smarter". These tools enable computers to deal with ambiguity and questions of judgment that are too subtle for conventional data processing techniques. They also help workers handle much larger quantities of data when making decisions. Regardless of the setting, these systems give the user alternatives ranked according to their probability of success. These probabilities are based on the data available; as more data are added, the reliability of the results improves. Thus, the system appears to learn from the program as it adapts new knowledge into the database.

It also allows users to ask both "What if I do this ?" and "What is the best solution for this problem ?" Users also can ask the computer the question "Why is this the best solution ? " before applying the results. Expert systems help workers at many levels make better decisions by providing them with additional information and a

structured way to use that information. These systems are being used in many different types of organizations throughout the economy. For example, public utilities use them to monitor and improve the performance of their coal-fired boilers. About one-fourth of the Nation's largest insurance carriers currently use expert systems, primarily to analyze insurance applications. Manufacturers use them to design products, control processing, and to serve customers who have a problem with a product. Sitting at a computer, employees enter the symptoms and other data provided by the caller, and the program helps them diagnose the problem and prescribe corrective measures.

Related Technologies

Expert systems are only one of a cluster of AI technologies with commercial significance. Others involve creating equipment with a human-like ability to move, see, and communicate. The problems encountered so far have driven home to scientists just how complex we humans are and how difficult it is to simulate even our most basic reasoning processes. Getting a robot to do something that a baby does naturally requires hundreds, perhaps thousands, of detailed instructions.

Most researchers now concede that it could be a long time before the necessary breakthroughs occur that will lead to the production of machines that think or reason in any fully human sense of the word, or that act autonomously, or that speak and understand human speech in all its complexity. Nevertheless, researchers keep making advances in these areas.

Robotics

Although a far cry from Robobutler, robots are increasingly useful. Industrial robots have been around for about 25 years, doing simple, repetitive tasks requiring no decisionmaking, and doing them with superhuman speed and precision. They now perform a remarkably wide variety of tasks, from assembling computers, artillery shells, and vacuum cleaners to popping frozen dinners into their trays, inspecting different kinds of products, and drilling holes for brain surgery. They are also used in hazardous environments, such as mines and nuclear powerplants, and for handling toxic waste.

Significant improvements have been made in robotics technology in recent years. For example, to reprogram a new robot arm, an operator simply selects from among choices that the robot itself offers; a worker can teach the robot's eyes to recognize a new part in less than 10 minutes. The number of robots in use has grown to over 32,000. As prices fall, they will become more attractive to potential users, especially as employers continue to experience difficulty attracting workers. Robot technology may be especially desirable in fast-food preparation, some areas within health care, building maintenance, and security, industries that traditionally have relied on young people.

Machine vision

Advances in computer capabilities and in the understanding of how image processing should be performed have led to the development of software that enables vision systems to better simulate the human visual process. These systems can store a digitized photograph of an object or scene and recognize a good bit of what is there.

Machine vision is preferable to human vision in many manufacturing and inspecting processes. Machine systems can measure items more precisely and store a standard frame of reference for judging an object over a longer period. In addition, machine vision tends to offer faster inspection speeds. These systems are primarily

used to make robots more effective in industrial settings. However, they are being used more often in medical research, such as analyzing blood samples. Natural language processing. AI research has led to the development of computers that can understand simple written instructions with limited vocabulary, such as "List all widgets sold in July". This is an improvement over conventional programming languages, but it is limited to the small number of English words that the computer can digest.

Artificial intelligence software also is being incorporated into speech recognition systems that make computers more user friendly. Systems are available that permit executives to access their data bases without entering commands on the keyboard and that enable disabled persons to control computers. Other applications include automatic dialers for cellular car phones, thus enabling the driver to keep at least one hand on the wheel, and systems that facilitate inventory control in factories and baggage handling at airports.

Parallel processing

Because speech recognition and machine vision applications require computers to sort through a vast amount of data, a major research effort has focused on ways to increase the processing speed of computers. Parallel processing -- linking large numbers of smaller computers together and assigning each a portion of the overall job -- yields much faster speed and greater power at a fraction of the cost of supercomputers or even mainframes.

Neural networks

Another new development is neural network computing, which attempts to duplicate the way the human brain processes information. Conventional computer circuits are arranged in series, with each transistor linked to only two or three others. In neural networks, each transistor is hooked up to most, if not all, of the others. A signal entering the system quickly fans out across the entire network, and all of the transistors process it. This new type of processing can be achieved in two different ways, with specially designed neural network computers or with advanced software used on standard computers.

Neural networks are not only faster than conventional computers, they can also develop new if-then rules from the data they receive. They also have superb pattern-recognition capability. For example, once a network has been shown three or four views of a particular face, it will instantly recognize that face from any other angle. This pattern matching has tremendous implications for military applications, product assembly and inspection, and natural language processing.

Employment Trends

The demand for AI specialists has changed over the years, reflecting the different stages of development the field has undergone. For years, activity in AI was restricted to college campuses and corporate research centers. Researchers came from a wide variety of academic backgrounds, including biomechanics, computer engineering, computer science, economics, electrical engineering, linguistics, neurobiology, optics engineering, physics, and psychology. Most had advanced degrees, usually doctorates. When the technology became commercially viable, many of the scientists involved in research left their laboratories to start their own companies.

At this point, the competition for experienced AI specialists intensified, with salaries increasing significantly as employers competed for available workers. This development, along with the publicity surrounding the success of specific AI products, stimulated more interest in the field. Both the availability of training programs and the

number of students taking these courses increased.

Although there are still spot shortages for experienced professionals in certain specialties, the widespread shortage seems to have eased. Salary increases for AI professionals have moderated.

Skill requirements are changing. Requirements are coming down as the field moves out of the research and development phase and into product refinement and implementation. Maintaining software based on AI principles is less demanding than developing the original technology. The use of expert system shells has also played a major role in this change. Developing a complex, standalone expert system requires much more programming skill and experience than does adapting a shell to the needs of a user. Because of this evolution, a bachelor's degree in computer science, with an emphasis on AI, is acceptable for a wide array of jobs, although persons involved in the design of AI software still need advanced degrees. Some employers even let people without advanced degrees perform knowledge engineering, generally with shells.

Demand for computer professionals with AI skills will continue to rise. Even if further advances in basic technology aren't made -- and this is highly unlikely -- the development, integration, implementation, and maintenance of products based on existing technologies will require many additional skilled workers. Demand for workers will also increase as more organizations use AI products.

Growth will occur in software houses and hardware developers producing AI products and in large corporations who are developing their own AI capabilities. In these organizations, demand will be strongest for programmers who can program in LISP or other AI languages. Demand is also growing for knowledge engineers, who can work closely between the programmer and the user to design expert systems.

The primary area of growth, however, will be in user organizations. More programmers and systems analysts with AI skills will be needed to integrate shell programs into existing systems and to develop and maintain in-house systems, much as in-house computer professionals develop accounting or database applications today. Artificial intelligence is clearly merging with the field in which most computer professionals are employed, the development and maintenance of management information systems (MIS). In this environment, the strongest demand will be for programmers and systems analysts with experience in regular systems and a good working knowledge of AI.

Where To Get the Training

Those with AI training will be better able to deal with changes coming down the road. The trend to incorporate AI principles and techniques into software is accelerating, and many observers predict that these advanced languages will become the programming standard. If this does occur, those already knowledgeable about AI will have a distinct advantage. Between 80 and 100 colleges and universities offer AI courses in the computer science department, and the number is growing. At these schools, a student can earn a degree in computer science with an emphasis in AI. Students who get extensive training in computer science, especially applications development and programming -- coupled with courses in LISP or other advanced languages, expert system development, as well as machine vision, natural language, and other applications -- should have very good job prospects. Jobseekers with this background are prepared for both the old and the new, enabling them to work in a traditional MIS environment as well as in an AI environment. The training situation for knowledge engineers is similar to that for systems analysts a decade ago. There are no formal programs for these people, or job specifications for that matter. Currently only about 20 schools, mostly on the west coast, offer courses in knowledge engineering. However, this number is expected to rise. Training in knowledge engineering and other aspects of artificial intelligence is also available through AI consulting firms and professional development

companies.