ARTIFICIAL INTELLIGENCE: STATE OF THE ART

Technical Note 198

October 1979

By: Harry G. Barrow, Senior Computer Scientist

Artificial Intelligence Center
Computer Science and Technology Division

To be published in Proceedings of Conference on the Factory of the Future (presented by Infotech International in England).

ARTIFICIAL INTELLIGENCE

Harry G. Barrow
SRI International
Menlo Park, California

I INTRODUCTION

The birth date of the field of Artificial Intelligence (AI) is open to debate. Its name, however, was bestowed upon it in 1956 by John McCarthy [23]. From time to time other names have been suggested, such as "Machine Intelligence," "Cognology," "Theoretical Psychology," or "Experimental Philosophy," but the original name endures. As the diversity of the proposed alternatives might imply, AI encompasses a broad range of activities and motivations.

Defining the scope of a scientific field is always difficult, usually resulting in something vague, broad, and uninformative; for example, try looking up "Physics," "Chemistry," or "Biology" in a dictionary such as Webster's. What is more, there are always areas, like molecular biology, that cannot be safely assigned to one field. For the present discussion, I will offer the following description of AI:

> Artificial Intelligence is concerned with understanding the principles and building working models of intelligent behavior.

For some research workers, the first part of this description is paramount, and the goal is that of understanding the human intellect. They operate in the classical scientific manner, developing a theory and then building a computer model to test it experimentally, in much the same way as a physicist or a psychologist. For others, the second part is the driving force, and they seek to design machines that can relieve people of burdensome intellectual tasks. They do so by building computer systems that perform at an acceptable level of competence, and then they attempt to further understand and extend the methods and principles involved. This approach is more akin to aeronautical engineering and the development of aerodynamics. As in other fields, science and engineering, principles and applications, go hand in hand. These two main operational approaches share the very practical step of the working model. For the time being at least, the model takes the form of a computer program, because the digital computer is currently the most powerful and versatile modeling medium available (as evidenced by the decline of random logic and the rise of programmable logic in many areas of electronics). Because of the centrality of the digital computer, AI is often considered to be a branch of computer science.

Whatever the motivation, the problems tackled by research workers in AI are typically large, complex, and involve ambiguity and uncertainty. AI attempts to come to grips with the real world, with all its confusion and richness, and to interact with people, with all their irrationalities and idiosyncrasies.

II WHAT DOES AI ENCOMPASS?

The field of AI is comprised of a broad range of research activities which overlap to greater or lesser degrees. Part of the current research focuses upon a particular area of intelligent activity (like vision or automatic programming), and part upon general techniques that underly such activities (like search or inference). We may broadly categorize research as in the following table.

Table I: Some Areas of Research in Artificial Intelligence

Activities	Techniques
Visual Perception	Representation and Modeling
Robotics	Control Structures
Speech Understanding	Heuristic Search
Language Understanding	Planning
Expert Systems	Perception
Data Management	Deduction
Automatic Programming	Induction
Game Playing	Learning
	AI Languages and Systems

It is probably safe to say that AI seeks to understand and model virtually all human intellectual activity. In doing so, it has drawn from and contributed to many other disciplines, particularly psychology, mathematical linguistics, mathematical logic, operations research, decision theory, pattern recognition, and computer science. It has stimulated important developments in software technology, especially concerning advanced programming languages and systems.

III WHO DOES IT?

Five or ten years ago, AI research was almost limited to five main centers: Carnegie-Mellon University, MIT, Stanford University, Stanford Research Institute*, and Edinburgh University. Now, however, although these places remain large (50 people in the SRI group) and preeminent, many others now exist. Many universities, including Yale, UC Berkeley, Texas, and Massachusetts, now have small groups (usually in computer science departments); and many companies, including BBN, RAND, Xerox, General Motors, Lockheed, and Texas Instruments, have established groups who develop and apply AI techniques. The field is no longer the private domain of the USA and the UK; valuable work is now carried out in

----

* Now known as SRI International.

Canada, France, West Germany, Italy, and Sweden. Especially noteworthy are the concentrated efforts being made in Japan and the Soviet Union.

Funding for AI research in the USA was originally provided mainly by the Advanced Research Projects Agency (ARPA) of the Department of Defense. AI is now also supported by the research offices of the three armed services (AFOSR, ARO, and ONR), the National Science Foundation, the National Institutes of Health, and NASA, as well as by in-house efforts of various companies.

Publications in AI appear in a number of places. Many of the highest quality papers are published in the journal Artificial Intelligence. Papers also appear in IEEE transactions on Computers, Systems, Man, and Cybernetics, and the relatively new Pattern Recognition and Machine Intelligence. Other sources are the journals Cognition and Computer Graphics and Image Processing. In addition, the International Joint Conference on Artificial Intelligence is held every two years, alternating with the rather smaller Artificial Intelligence and Simulation of Behavior Conference in Europe. The proceedings of these conferences contain many important papers.

IV A BRIEF HISTORY

Prior to the development of large computers, work on theory and practice of intelligent systems was sporadic with a few highlights, such as the work of Alan Turing in the 1930s and 1940s [38].

On the early machines of the late 1950s and early 1960s, programs were written that processed symbolic information to solve simple puzzles, prove theorems in logic and geometry, perform integration, and play games, such as checkers or chess. In these efforts, the concern was to do these things at all, rather than to do them efficiently or powerfully. Nevertheless, many basic ideas, such as heuristic search, were formulated, and special tools, such as the programming language LISP, were developed.

In the middle 1960s the main AI research groups were established, and they began systematic exploration of problems in natural language, automatic problem solving, and visual perception. At this time, ARPA became a major patron of the field.

In the late 1960s efforts at the major centers were directed toward development of integrated robot systems with capacities for vision, planning, and manipulation or navigation. At this time, work also began on "expert systems"--systems that made use of a large amount of knowledge about a limited subject.

In the early 1970s a period began of increasing emphasis on applications and upon specialization. ARPA instigated a major attempt to develop systems that understand continuous speech. Systems began to emerge that demonstrated significant levels of competence, and even expertise, in dealing with restricted classes of problems.

V THE STATE OF THE ART

In this paper I shall not attempt a comprehensive survey and evaluation of the field. Instead, I shall attempt to give a feeling for the current situation by highlighting a few AI systems that display an interesting level of competence in real intellectual tasks. A great deal of fascinating and valuable work exists at theoretical and exploratory levels, but for now I will concentrate upon practical issues.

If you are interested in obtaining a more extensive background of the subject, you will find an excellent survey by Nils Nilsson in the Proceedings of the 1974 IFIP Conference [25], and a shorter assessment of the state of the art in 1977 by Duda, Nilsson, and Raphael in Research Directions in Software Technology [10].

In recent years a number of informative textbooks on AI have been published and are some of the leading exponents of the field [7, 24, 30, 42, and 43]. They give much more detailed descriptions of the methodology, principles, and philosophy of AI than is possible in a short review article.

The best way to appreciate the current state of the art of AI is to look at some recent programs, assess their capabilities, and examine their workings. In the following sections, I shall consider a few major research areas within AI and give an outline description of some of their more interesting working programs, going through one of them in some detail. (Programs I shall consider in depth are those with which I am most familiar. The example programs, therefore, will often originate from SRI.) The areas I have chosen are: Expert Systems, Natural Language Understanding, Question Answering, Vision, and Robotics.

VI EXPERT SYSTEMS

During the last decade, a number of AI systems have been developed that specialize in some narrow domain, which might lie within algebra, or geology, for example. Each system encompasses a quantity of highly specialized knowledge appropriate for the range of problems it is intended to meet, and exhibits performance that compares favorably with that of a human expert. The domains typically involve qualitative knowledge and value judgements, rather than algorithmic solutions. Thus knowledge is often fuzzy and may be in terms of rules of thumb, and the expert systems contain machinery for making inferences from such uncertain rules.

An early example is the MACSYMA system [22]. MACSYMA is a program to aid applied mathematicians in such tasks as solving equations like

$$Y^{2x} - 3Y^{x} + 2 = 0 \quad \text{(for x)} \quad ,$$

symbolic integration like

$$\int (x + e^{x})^{2} \, dx \quad ,$$

or factorization of polynomials. These tasks can be tedious and error-prone when expressions become large. The user commands MACSYMA to perform various operations on equations and expressions, such as substitution, solving, or integrating; it carries out the action and does the bookkeeping. The system is now used and maintained by a consortium of the Energy Research and Development Administration, NASA, NIH, and the Navy laboratories.

A famous expert system is DENDRAL [6], which infers chemical structure from organic mass spectrogram and nuclear magnetic resonance data. For some families of molecules, particularly estrogenic steroids, it operates more accurately and much more quickly than the best human mass-spectrum analysts. DENDRAL has built into it a considerable amount of knowledge of chemical structures and how they break up in mass spectroscopy, provided by a number of collaborating human experts. It is now routinely used by a national community of chemists.

There have been several expert systems developed in the area of medical diagnosis and treatment. One of the earliest and best known is MYCIN [35], a consultation system intended to assist physicians by diagnosing bacterial infections and suggesting therapies. Its expertise is comparable to that of a general practitioner on these problems. PUFF [12] is a similar system for diagnosing pulmonary-function disorders, given case histories and results of various lab tests. It too performs at the expert level, having been tested on 150 cases in comparison with human experts, and having achieved an extremely high degree of agreement with them. INTERNIST [28] is a program that specializes in internal medicine. It is still under development, but it already rivals human performance in some areas.

These medical diagnosis systems all use similar mechanisms, though the specific knowledge involved is different. I will briefly outline how such systems work, using yet another expert system, PROSPECTOR, as an example [9].

PROSPECTOR is an interactive aid for geologists involved in mineral exploration. It attempts to determine what mineral deposits may be present in an area and with what certainty. There are about 35 main types of mineral deposits that are of economic and geological interest; as it presently stands, PROSPECTOR has knowledge of nine of them, relating to various classes of lead, zinc, copper, and uranium deposits. The user holds a conversation with the system in simple English. He first supplies information that he thinks might be relevant about the prospective site, such as "There is galena," or "There might be sphalerite." PROSPECTOR then begins to ask him questions, such as "To what degree do you believe that there is evidence of metamorphism?", to which the user responds with an estimate of confidence ranging from -5 (certainly not) to +5 (certainly so). On the basis of the user's input, the system estimates the likelihood of various alternative hypotheses, determines the best one, and asks a question designed to verify or refute it. At any point the user can ask PROSPECTOR why it is pursuing its current line of reasoning. Typing "WHY" to the previous question would result in the reply:

> "The evidence of high-temperature mineralization is discouraging for the prospective ore body being an MVTD. However, if the high temperatures were due to subsequent metamorphism, then this discouraging evidence can be discounted."

How does PROSPECTOR work? Before the program is written, much time is spent with geologists, asking them how they go about their prospecting activities: what signs they look for, what evidence they seek to justify hypotheses, what do they believe are the processes that led to formation of ore deposits, and so forth. This interview process is not as straightforward as you might think. An expert is often unaware of precisely how he comes to his conclusions. He will not, in general, give hard and fast rules, and he may even contradict himself. Thus, a great deal of interaction is required between AI scientists and the experts, and many iterations are necessary before both parties are satisfied.

When sufficient information has been gathered, it is organized into a set of hypotheses, like "There are abundant quartz sulfide veinlets with no apparent alteration halos," or "Alteration is favorable for the potassic zone." The hypotheses are linked by rules of the form:

IF "There are abundant quartz sulfide veinlets with
no apparent alteration halos"

THEN "Alteration is favorable for the potassic zone"
(LS,LN) ,

where LS and LN are numbers that reflect the reliability of this rule. The hypotheses are not represented in the computer by English text, of course, but by elements of a data structure. The text is stored, however, so that it can be typed out for the user's benefit.

When a question is asked of the user, it comes from the IF part of some rule. The user's confidence estimate is combined with the rule's reliability measures and the current likelihood of the hypothesis, in the IF and THEN parts of the rule, to derive a new likelihood for the THEN hypothesis. Since other rules may have the same hypothesis in their IF part, they too can have the likelihoods of their THEN hypotheses updated; consequences thus propagate through a network of rules, a fragment of which is shown in Figure 1. The system then selects, according to an appropriate strategy, what is now the best hypothesis to pursue, and asks another question. The machinery that does all this is independent of the specific problem area (geology or medicine); only the rules themselves are problem-specific.

If the user asks "WHY" to a question, the system can respond by printing the text corresponding to the THEN part of the rule--the hypothesis it was attempting to establish.

In addition to the rules concerned with determining whether and what type of ore body exists, there is a set of rules concerned with locating its likely position. Information from a geological map, including outcrops, faults, mineral concentrations in the soil, and so forth, are input via a digitizing table; and the rules are applied to

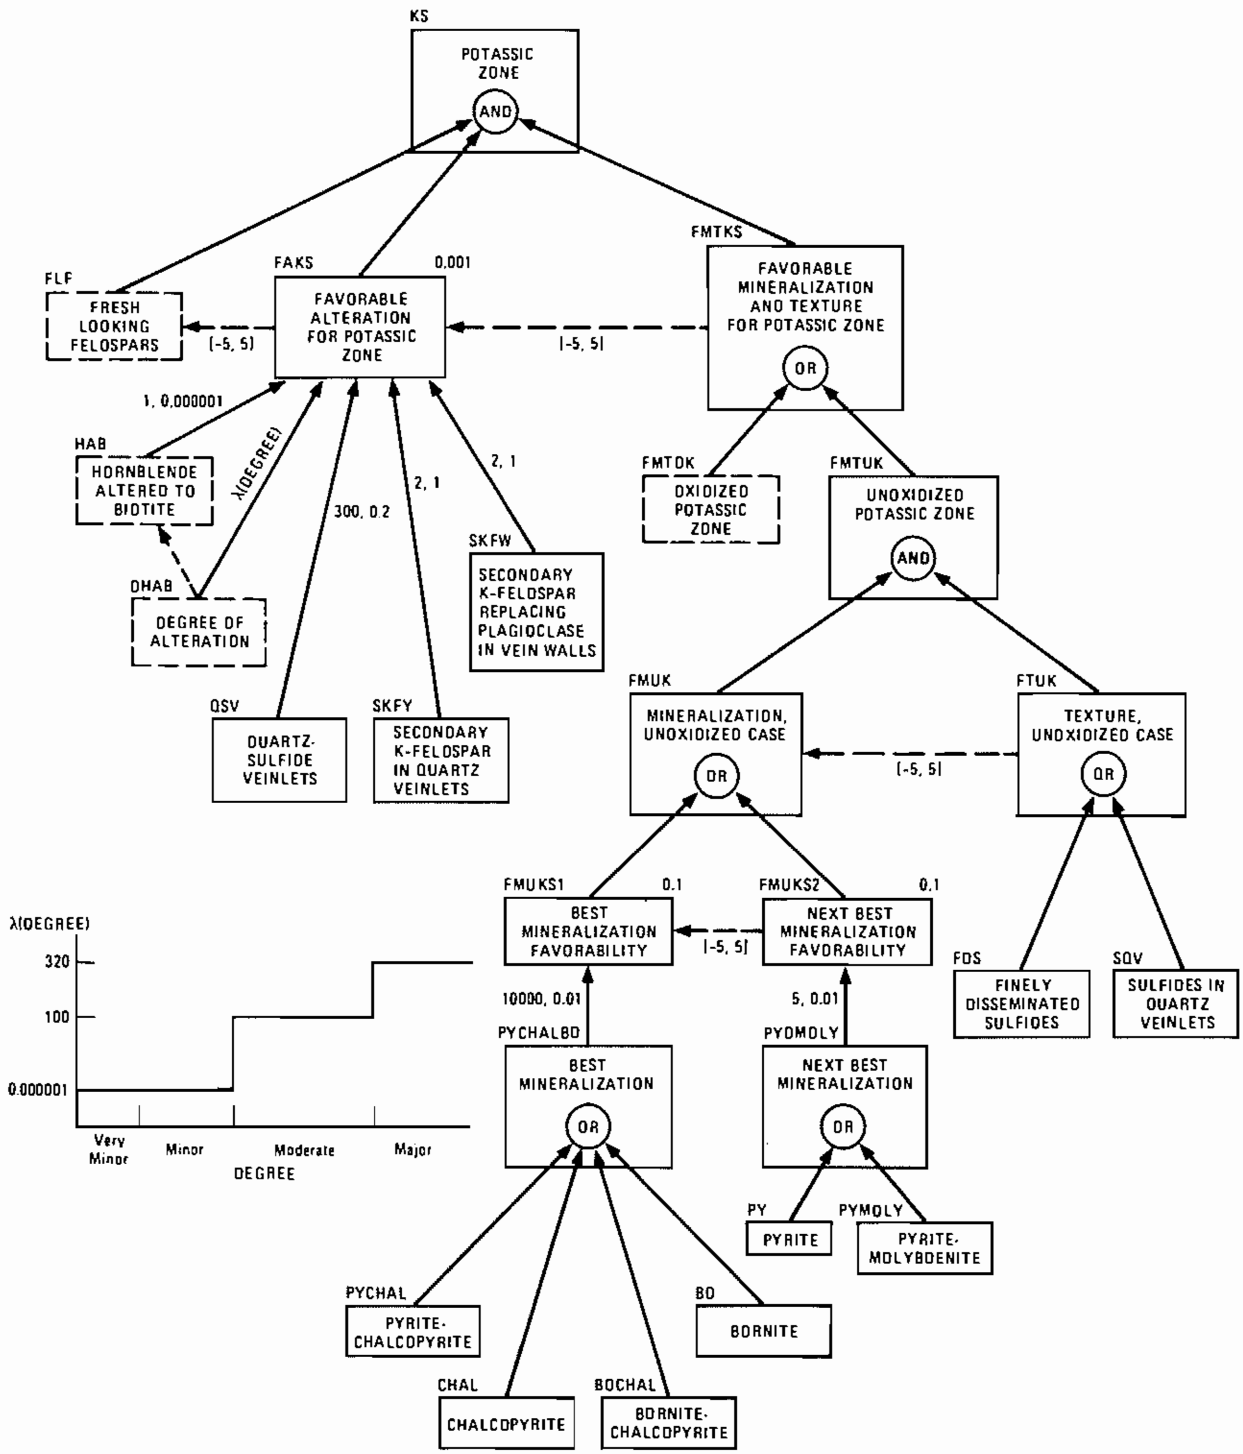

FIGURE 1 A FRAGMENT OF PROSPECTOR'S RULE NETWORK

determine for every point the likelihood of the presence of ore. In a real prospecting situation, test drill holes would normally be made at the most promising location. So far, the program has only been given data from sites at which the exact location and dimensions of ore deposits are known; the estimates it produced agreed well with what was actually found, and hypothetical drill holes would have struck ore.

PROSPECTOR has reached a level of performance that appears satisfactory to the geologists involved in the project; that is, it asks appropriate questions and arrives at conclusions very similar to those of a human consultant. It currently contains over 900 rules; but, as mentioned earlier, they do not cover all the major types of deposit, and so the program is not yet of commercial interest. PROSPECTOR has been tested on data from about 15 sites with known ore deposits, though it has not yet been used to find a new one. The economics of mining are such that, should it ever do so, it would pay for all the research and time involved perhaps thousands of times over!

VII EMPLOYMENT OF EXPERT SYSTEMS

Clearly, the scope of expert systems built upon these lines is very great, and we are only beginning to explore their possibilities. In industrial environments, for example, one can imagine various management aids for everything from planning budgets to organizing production--anywhere, in fact, where value judgments and rules of thumb are used in complex tasks. Expert systems could also aid engineers during the design process by helping them keep track of the many constraints and design rules they have to work with. More directly analogous to the medical-diagnosis programs might be interactive programs to aid troubleshooting and repair of manufactured or capital equipment. Such programs already exist for certain areas of electronics; Sussman is developing a rule-based program to aid in design [36], and Brown has developed a program to teach troubleshooting of power supplies [5].

Expert systems can be developed for some specific industrial applications, although, to my knowledge, no one has yet put up the funds to do so. It is of vital importance, however, to select the area of expertise very carefully. Each area has its own characteristics that may mean the difference between possibility and impossibility under the current state of the art. What is more, developing an expert system can take several man-years of effort, so the economics must be carefully weighed.

VIII NATURAL LANGUAGE UNDERSTANDING

One of the major goals of AI research has been to develop the means of interacting with machines in natural language (in contrast to a computer language), either spoken or written. Such a means of communication would permit untrained persons to use complex computer systems or equipment without having to learn special techniques to do so. The earliest work in the 1950s was on automatic translation from one language to another and was based upon grammar and syntax, without involving subtleties of meaning. Perhaps you have heard the story of the translation program that was asked to translate "out of sight, out of mind" into Russian and back, and produced "invisible idiot." Whether the story is true, I do not know, but I am prepared to believe it is!

In current attempts to handle natural language, the need to use knowledge about the subject matter of the conversation, and not just grammatical niceties, is recognized--it is now believed that reliable translation is not possible without such knowledge. It is essential to find the best interpretation of what is uttered that is consistent with all sources of knowledge--lexical, grammatical, semantic (meaning), topical, and contextual. Following this approach of "understanding" the sentence, a few systems have been developed that display interesting levels of performance.

In 1971 Winograd produced a simple demonstration system that made use of knowledge about what it was discussing [40]. Part of the program displayed a simple world on a screen, showing various blocks, pyramids, a box, and a simulated robot arm (see Figure 2). The user could type commands, like "Put the cube that is smaller than the green pyramid into the box." When the input had been successfully interpreted, the simulated robot would carry out the action on the screen. The user could also ask questions like "Is there a large block behind a pyramid?", getting the printed reply "Yes, three of them: a large red one, a large green cube, and the blue one." He could even ask questions about the robot's actions, like "Why did you pick up the small cube?", getting the reply, "To clear the top of the red cube to pick it up." While not powerful enough to handle complex, real-world situations, the system was a landmark on the way to doing so.

Another program that has a limited domain of discourse is the LSNLIS system, developed by Woods [44]. LSNLIS has a data base of information about the moon rock samples returned by the Apollo astronauts. It successfully answers a wide range of unconstrained, typed, English questions about the rocks and their properties.

To date the emphasis in natural-language-understanding research has been upon building into the system enough information about the world or topic of conversation to enable it to sift out the interpretations of the input that make sense; grammatical correctness is not enough to resolve the ambiguities that occur in normal conversation (can you find the four reasonable interpretations of "Time flies like an arrow"?), nor to adequately constrain the search for the correct parse. Obviously, knowledge of the subject matter is even more important when dealing with a dialogue, rather than a single, self-contained utterance. For example,

> "Put the pulley on the spindle."
> "OK, I've done that."
> "Now pass the belt over it."

-- does "it" refer to the pulley, the spindle, or the putting?

Does the shortest thing the tallest pyramid's support supports support anything green?

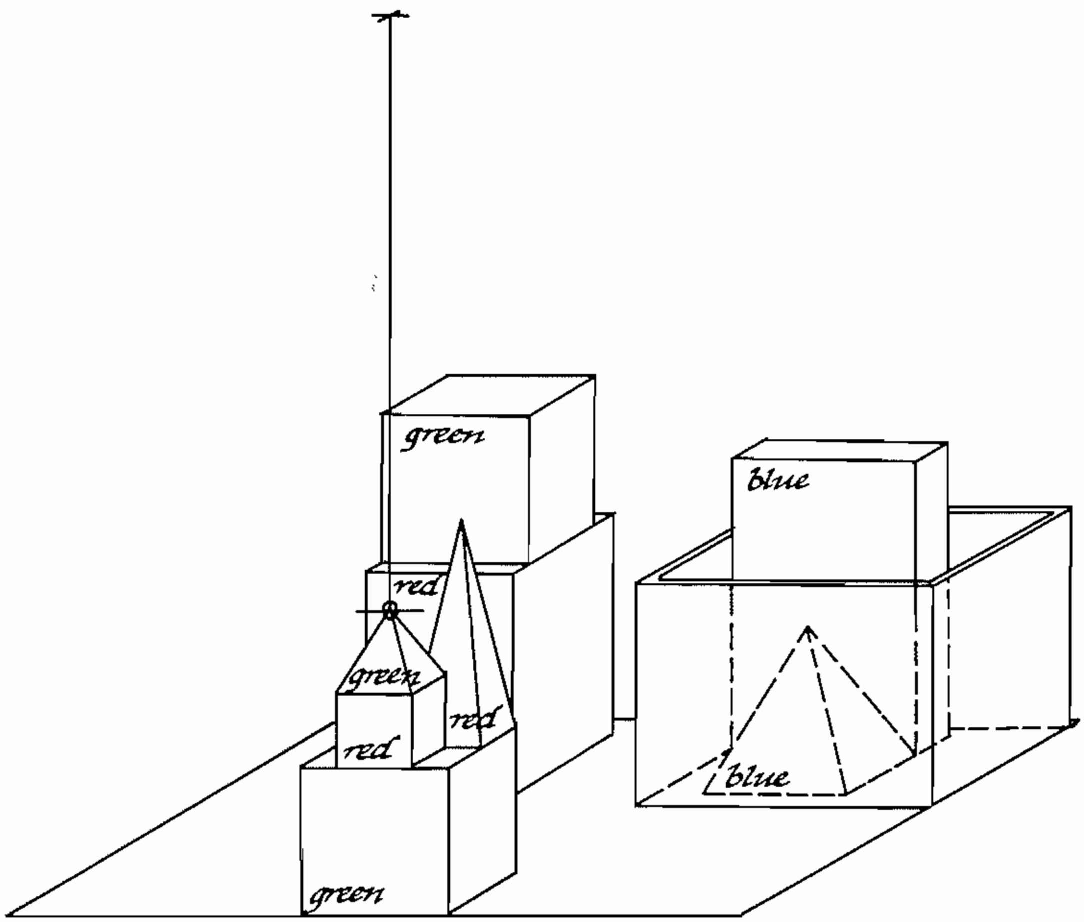

FIGURE 2 WINOGRAD'S SIMPLE WORLD FOR UNDERSTANDING NATURAL LANGUAGE

As a consequence of these observations, most of the recently developed systems handle discourse about a very constrained task or topic; a significant part of the problem is structuring the knowledge. This is not so restrictive as it might appear; some programs have been produced that may be of real, practical value, despite not being capable of handling the full richness of natural language.

One recent program by Marcus [20] is perhaps worthy of note, because it attempts to simplify later stages of interpretation by producing the best parse first, based upon a somewhat different variety of grammar than others have used. It appears to be quite efficient, not needing to rescue itself from false tracks, except for the types of sentence that humans find confusing, such as "The horse raced past the barn fell" (which is perfectly correct English). There is considerable interest in this program because it is possible that it may operate in a way which is quite close to that used by people. It may represent a promising line of development, but it is not ready for the real world yet.

While research continues to grapple with the full richness of human language, there have been some more pragmatic approaches that attempt to make available simple natural language communication for constrained situations. These systems capitalize upon the well-understood ideas in grammatical parsing and bring in the necessary semantics by structuring the grammar. For example, in parsing "How long is the Enterprise?", the grammar would not simply represent "Enterprise" as a noun, but as a ship's name. This semantic structuring is important for systems that must actually DO something in response to the input. The grammar will be augmented with information about how to respond, and the response must depend upon the meaning of words. For example, answering "How long is ...?" may require different actions, depending on whether the object is a ship (look it up in a data base) or a file (go and count the characters).

In accordance with these ideas, a system called ROBOT has been developed [15] as a simple interface between a user and a data-base management system. The user can ask simple questions in English and have answers retrieved from the data base and printed out. The ROBOT program is now being marketed in the USA.

LIFER, developed by Hendrix [16], is a general-purpose natural language system which can be added to a complex computer program to provide a painless means of communicating questions or instructions to it. The appropriate grammar and semantic routines must be provided for each new application, but that is a comparatively straightforward task. LIFER has been used as the means of communication with the PROSPECTOR, LADDER, and HAWKEYE systems discussed in other sections (in English, Swedish, and French).

With such packages, restricted natural-language-text access to computer systems is possible now. It is not difficult to find input utterances that they cannot interpret; but, despite their limitations, their promise is such that I am sure we shall see their widespread application in the very near future.

IX SPEECH UNDERSTANDING

So far I have been discussing communication via a typewriter keyboard. The situation for spoken language is not quite so advanced. The difficulty is that all of the problems involved in interpreting text occur with an added dimension of complexity--that of breaking the input sound into phonemes (the basic sounds of language) and words. It is hard, if not impossible, to do this out of context, even for people. To do an adequate job, it is necessary to bring in grammar and meaning, and to make searches; the computation required is apparently great.

In 1971 ARPA established a research effort to develop systems capable of "understanding" natural speech at a useful level of performance. The goal was to develop a prototype system that would be able (in the context of a limited domain of discourse) to understand (in a few times, slower than real time) an American (whose dialect is not extremely regional) speaking (in a "natural" manner) ordinary (although perhaps somewhat simple) English sentences constructed from a 1,000-word vocabulary. These goals were very ambitious for the state of the art at the time; the project brought together workers in acoustics and speech research, as well as in AI, to meet them. In 1976 the goals were essentially met by a system developed at Carnegie-Mellon University, called HARPY [19]. HARPY was tested on 184 sentences, spoken by five different speakers using an inexpensive close-talking microphone, after it had learned the individual speaker's characteristics by training on over 20 sentences from him. It correctly interpreted 95% of the sentences, taking about ten times as long to interpret the sentence as it took to speak it.

[Note: If you are willing to keep to very limited inputs, say about 100 predefined, 3-second sentences, then there is already equipment on the market (for example, Nippon Electric's $80,000 connected speech recognizer or Threshold Technology's simpler $10,000 recognizer). Such equipment does not deal with the meaning of the utterance, only its sound, so the user must know clearly which predetermined phrases are recognizable by the machine. It is, however, useful enough to provide input to machines in constrained situations; at SRI we have experimented with using speech input to train a computer-controlled Unimate, saying "Up ..., left ..., down ..., two ..., inches ..., etc."]

The generation of speech output is a much easier problem. If we have a piece of text, preferably in phonetic notations, the generation of appropriate sounds is a straightforward process, governed by a collection of rules. There are plenty of devices on the market (such as the Votrax from Federal Screw Works) that do quite a good job of producing artificial speech. They may sound a little too artificial at times; but that does not cause any severe practical difficulties, since humans, considerably experienced with variations in dialect, are on the receiving end.

The more interesting, from an AI point of view, stage of speech generation is that of deciding what to say. A program might have information stored as a data structure with cross links, and the problem is to find a way of reexpressing the structure in a linear, sequential form. Once more, this problem is not as difficult as its inverse (inferring the data structure from the sequence), and fairly simple techniques suffice in practical situations.

The conclusion is clear--adequate techniques are currently available for speech output, but we are still a long way from fully general speech input to machines. For practical situations with limited topics of conversation, however, we are within a few years (almost certainly less than five) of useful systems.

X QUESTION ANSWERING

A long-standing area of interest in AI has been that of questioning a program or person about the knowledge it has. Originally concerned with problems of representing facts and making inferences from them, this work has led into the area of intelligent retrieval of information from data bases.

Nowadays, data bases are very common. In a manufacturing company, data bases might exist containing information about orders, accounts, personnel, components, and stock, among other categories. Typically, the person who needs to access the information is kept at a distance. He is usually provided with a standard summary of data at regular intervals, and if he needs something a little different, he must have a computer programmer write a special program for him. The programmer has to know a great deal to do so; he must know about the structure of fields in data base records, their names and the names of the various data files, what is in each file, which computer has which files, the query language of the data-base management system, the appropriate programming language, the job command language, logging in protocol, and interfacing conventions. He does not necessarily know, however, what sorts of questions the user will want to ask.

There have been some attempts in AI to remedy the situation. Ideally, the user would like to ask his question directly, in his own natural language, without worrying about the details of what is involved. LADDER is one system, developed at SRI [34], that lets him do so. It provides access to information about the ships of the entire US, USSR, and UK navies, including their type, location, armaments, personnel, destination, capabilities, and so forth. (The data base does not contain current or secret information, I should add.) The user may type into a keyboard complex questions in English, such as:

"How far is the Constellation from Charleston?",
"Where is the longest ship carrying vanadium ore?",
or "When will the Philadelphia reach port?".

If he has a terminal that supports graphics, the user can have LADDER draw maps for him as well:

"Display a map showing all ships with less than 40% fuel aboard that are within 2 hours' steaming time of an oiler."

There are various other nice features of the system. It can handle ellipses (partial sentences). For example, after asking "What is the length of the Skipjack?" and getting his answer, he can simply type "the beam" and get the beam of the Skipjack, then type "of the Enterprise" and get the beam of the Enterprise. LADDER also lets the user define his own abbreviated notations or paraphrases. For example:

"Define 'Q length Kennedy' to be like 'what is the length of the Kennedy'"

will permit him subsequently to say "Q length Skipjack" or "Q beam Sunfish" and get what he expects.

The LADDER system is composed of several stages (see Figure 3). First, a simple natural language interface program reads the input question and converts it into a structured internal representation. This stage is actually built upon LIFER, described earlier, and uses information in the form of a grammar for part of English, plus various routines to check on the meaningfulness of alternative interpretations of the sentence. The next stage of processing has information about how the data base is structured, the files involved, and so forth. It takes the initial query representation and plans how to retrieve the answer in an efficient way. Retrieval may involve several data-base accesses in sequence and computation to sift out what is needed. This stage essentially writes a simple program in data-base query language. Figure 4 shows a simple question and the corresponding data-base query program (the complexity of which is due to the weakness of the data base management system, not to LADDER).

The final stage is concerned with the process of actually carrying out the plan. LADDER can access files that are not on the same computer but may be distributed over various sites linked via a communications network. The third stage has information about the geographical location of files and how to reach them. It sends off the query program to the appropriate computer, where it is run and the results returned. It is sufficiently robust to function even when the remote computer is unavailable or crashes, by detecting the fact and rerouting the query to another computer that has the relevant information.

When the required data is returned from the remote computer, it is passed back through the various stages, which reformat it and perform any additional computations. At last, the data is presented to the user in an appropriate form--as tables, maps, or English statements.

LADDER has been operating for about two years now, being used by the U.S. Navy on an experimental basis. In that time its English grammar and the types of questions it can handle have been increased to the extent that it can now reliably respond to any reasonable query. Grammars have been written so that the user may make his inquiries in English, French, or Swedish; the corresponding program that LADDER produces can be in DATALANGUAGE for the Datacomputer relational data

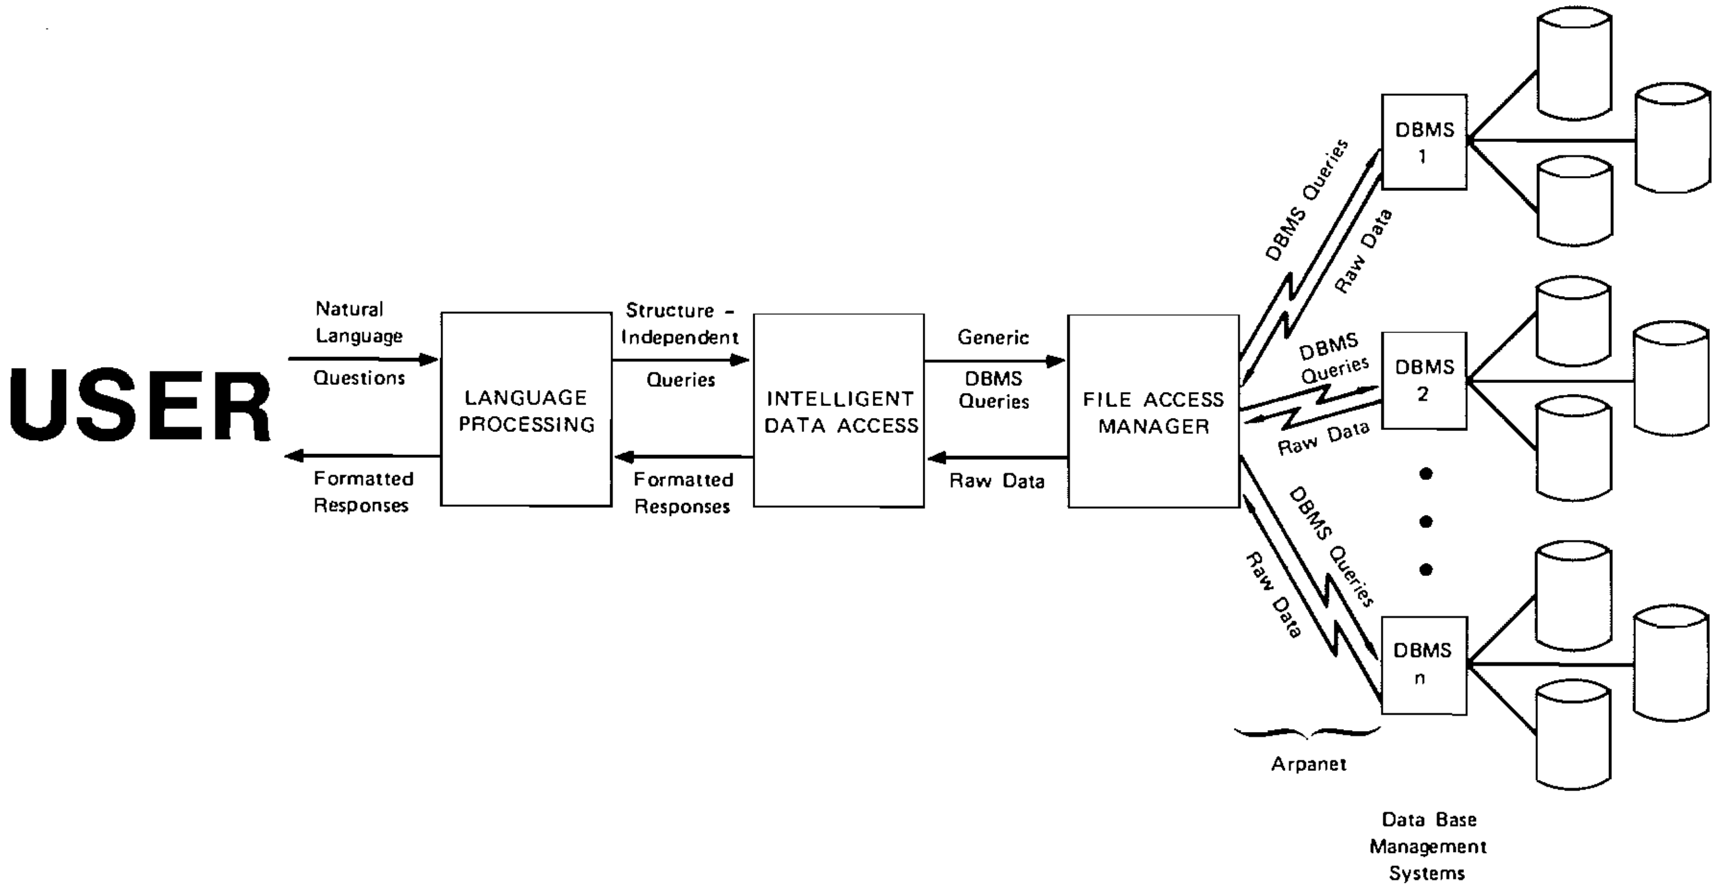

FIGURE 3 STRUCTURE OF THE LADDER SYSTEM (Language Access To Distributed Data With Error Recovery)

```
1_To what country does the fastest sub belong?

     COMPUTE XSTRX11 = '00.0' $
     COMPUTE XY10 = 0 $
     FIND FIRST SHIPCLASCHAR RECORD OF BLUEAREA AREA $
13   IF ERROR-STATUS = 307 GO TO 14 $
     COMPUTE XAND16 = 0 $
     IF SHIPCLASCHAR-TYPE2 NE 'S' GO TO 17 $
     COMPUTE XAND16 = 1 $
17   IF XAND16 = 0 GO TO 15 $
     COMPUTE XAND18 = 0 $
     IF SHIPCLASCHAR-TYPE1 NE 'S' GO TO 19 $
     COMPUTE XAND18 = 1 $
19   IF XAND18 = 0 GO TO 15 $
     COMPUTE XSTRZ12 = SHIPCLASCHAR-MCS $
     IF XSTRZ12 LT '00.0' OR XSTRX11 LE XSTRZ12 GO TO 15 $
     COMPUTE XSTRX11 = XSTRZ12 $
     COMPUTE XY10 = 1 $
     COMPUTE XSTR29 = SHIPCLASCHAR-MCS $
15   FIND NEXT SHIPCLASCHAR RECORD OF BLUEAREA AREA $
     GO TO 13 $
14   * $
     IF XY10 = 0 GO TO XT $
     FIND FIRST SHIPCLASCHAR RECORD OF BLUEAREA AREA $
20   IF ERROR-STATUS = 307 GO TO 21 $
     COMPUTE XAND23 = 0 $
     IF SHIPCLASCHAR-TYPE2 NE 'S' GO TO 24 $
     COMPUTE XAND23 = 1 $
24   IF XAND23 = 0 GO TO 22 $
     COMPUTE XAND25 = 0 $
     IF SHIPCLASCHAR-TYPE1 NE 'S' GO TO 26 $
     COMPUTE XAND25 = 1 $
26   IF XAND25 = 0 GO TO 22 $
     COMPUTE XAND27 = 0 $
     IF SHIPCLASCHAR-MCS NE XSTR29 GO TO 28 $
     COMPUTE XAND27 = 1 $
28   IF XAND27 = 0 GO TO 22 $
     FIND FIRST SHIPCLASDIR RECORD OF BLUEAREA AREA $
30   IF ERROR-STATUS = 307 GO TO 31 $
     COMPUTE XAND33 = 0 $
     IF SHIPCLASDIR-SHIPCLAS NE SHIPCLASCHAR-SHIPCLAS GO TO 34 $
     COMPUTE XAND33 = 1 $
34   IF XAND33 = 0 GO TO 32 $
     SET SHIP-UICVCN TO SHIPCLASDIR-UICVCN $
     FIND SHIP RECORD $
35   IF ERROR-STATUS = 326 GO TO 36 $
     PRINT SHIP-NAM SHIP-NAT $
37   FIND DUPLICATE SHIP RECORD $
     GO TO 35 $
36   * $
32   FIND NEXT SHIPCLASDIR RECORD OF BLUEAREA AREA $
     GO TO 30 $
31   * $
22   FIND NEXT SHIPCLASCHAR RECORD OF BLUEAREA AREA $
     GO TO 20 $
21   * $
     GO TO XT $
     END
```

FIGURE 4 A QUESTION TO LADDER AND THE CORRESPONDING DATA BASE QUERY LANGUAGE PROGRAM base or in IQL for the DBMS-20 CODASYL-style data base. Experiments with the system have shown that between 90 and 95 percent of user queries are successfully handled (the failures being due to queries beyond the scope of the available data or analyzing capabilities, or to their formulation in unusual language). LADDER is also quite efficient--the AI part of the system, which involves everything up to transmission of the query program over the network, is actually faster than the conventional data-base management system it calls upon.

XI VISION

Computer vision is a topic that has received a great deal of attention in AI research. The goal of providing a machine with its own means of perceiving the world around it is an important one. Such operations as inspecting assemblies, checking aerial photos for changes, or visually screening tissue cultures or blood samples are tedious and hard for humans to endure for long periods. Nevertheless today they are performed by people simply because they involve seeing.

Early work on vision was concentrated on the so-called "blocks world," a simplified domain comprised of polyhedral (flat-surfaced) objects on a table top, largely because it is easy to model the objects and their appearances.

A landmark in machine perception was a program developed by Roberts [31]. Roberts used an image-dissector camera to look at blocks-world scenes involving bricks, wedges, hexagonal prisms, or objects formed by sticking these together. His program could determine the dimensions, location, and orientation of the objects; and, as a demonstration of its "understanding," it could display a drawing of the scene observed from any other viewpoint.

Roberts' program operated by first finding the places in the image where brightness or shading changed abruptly, corresponding to points on the edges of the object. Then it grouped the points together into lines and formed a line drawing of the scene. It interpreted the line drawing by finding triangles, quadrilaterals, and hexagons, which suggested possible objects (triangles suggest wedges, etc.) and eventually accounted for all the lines and junctions as edges and corners of objects. From the appearance of the object in the image, its dimensions, location, and orientation could be computed.

Following Roberts' work, effort was spent on finding better ways to interpret line drawings of polyhedra, with the aim of thoroughly understanding a limited area before tackling more complex ones. Of particular interest were techniques that did not depend upon knowing what specific objects (cubes, wedges, etc.) could appear in the scene. Waltz [39] developed a successful approach that could analyze a line drawing as complex as Figure 5 and produce a detailed description, stating whether each line corresponded to a convex or concave object edge or to a shadow or crack, what the approximate orientation of the surface corresponding to each region was, and whether it was illuminated or shadowed.

Subsequent work in the blocks world attacked the various components of the problem, from finding edges and lines to interpreting them. No system has been put together using the best components; but I believe if it were done, the system would be very competent in its limited domain.

Development of competent general vision systems has turned out to be very difficult, mainly for the reason that in dealing with images of real scenes, one is forced to confront very many complicating issues simultaneously. Real objects do not necessarily have simple shapes and smooth surfaces, spatial structure is not known precisely, surfaces do not have trivial reflecting characteristics and have visible texture, and lighting is complex because of reflections from other objects. If simplified assumptions are made, there may be common situations in which they are violated. In consequence, research in computer vision has been advancing along several major fronts, resulting in development of component techniques, specialist programs for constrained situations, and interactive approaches.

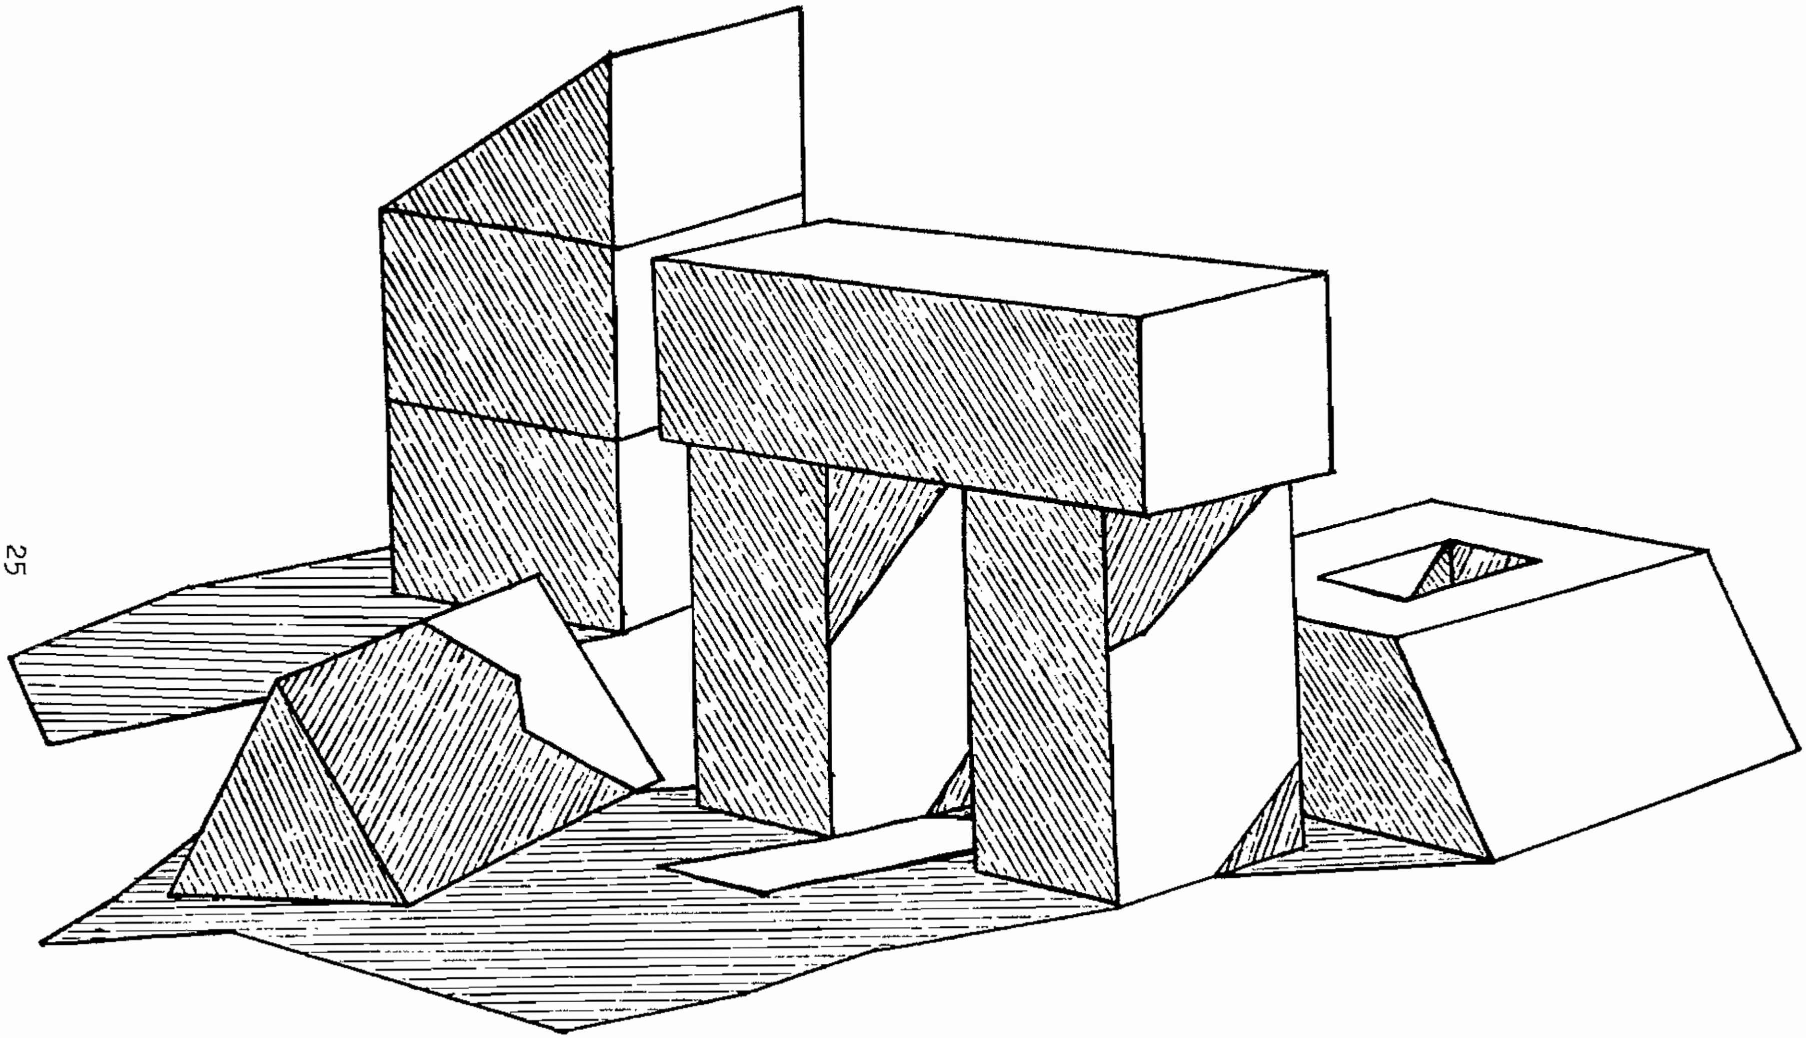

FIGURE 5 A LINE DRAWING CORRECTLY INTERPRETED BY WALTY'S PROGRAM

The scope of computer vision has expanded to include more natural environments (such as offices, landscapes, and aerial photos), with a broader range of sensory devices, including color TV, x-rays, synthetic-aperture radar, and active-ranging devices. This has spurred development of component techniques for finding features in images, recognizing objects, determining spatial structure, and so forth. The exploration of more natural domains has resulted in roads being made into problems of shading, texture, stereo, and motion in images. For example, one program, given an image of a smooth matt surface illuminated by a known light source, can determine reasonably accurately the three-dimensional shape of the surface from the visible shading [17].

As in other areas of AI, it was found that in vision, the more a priori knowledge the program had about what could be in the scene, the better it could perform. Many specialist programs have been written that perform some visual functions reasonably well within constraints. For example, one program detects lung spots due to pneumoconiosis in chest x-rays [18]. It has not been tested extensively, but apparently compares favorably with human radiologists, who show considerable variation in their ability to perform this task.

If the task is sufficiently well-constrained and controlled, it becomes possible to use rather simple techniques to extract all the desired information from an image. In the area of industrial automation, practical vision systems are beginning to emerge, based largely upon techniques developed a decade ago, which are adequate for a wide range of sensor-guided applications. While we await the development of more general computer vision, we can exploit the tremendous potential of what we already have in hand. This will be discussed in more detail in the following section.

An alternative interim approach, before totally automatic image interpretation becomes available, is an interactive system. The HAWKEYE system, developed at SRI, is a concept demonstration of interactive aids for cartographers or photo interpreters. The system has a library of images and a data base containing geometric and topological data, as found in conventional maps, together with more abstract information about a geographical region. It also has a simple natural language interface built with LIFER (described earlier). The user can ask what images are available of a given site and display them. He can calibrate images interactively (that is, determine the viewpoint, viewing direction, and characteristics of the camera), pointing to features in the displayed image and to their locations on a map. The system can then overlay known features from its digital map on the image, such as a network of roads, or can respond to questions like "Which building is this?" by referring to its map to determine the object the user indicates.

Making maps is still very much a manual process; features such as roads or rivers are traced with pencil and paper from photographs. The HAWKEYE system contains subroutines to aid this process. With a digital image on the display, the user simply indicates with cross hairs a few points on or near the road. The program then uses this information as a guideline with which to trace the road in detail. The resultant trace is more accurate than if it were made automatically (given the current state of the art) and much less tedious and potentially faster than if it were done purely manually.

A particularly interesting and significant trend in vision is that research in AI is beginning to contribute to the study of human vision. At at least two major centers (MIT and SRI), recent work has been concerned with understanding more about the physical basis of image formation--how characteristics of points in the scene, reflectance, illumination, surface orientation, and so on, determine the appearance in the image. The significant question is how (and if) the process can be inverted, deriving the scene characteristics from the image without requiring prior recognition of objects [3]. This branch of research is leading to suggestions as to how human visual systems may operate and how psychological experiments may test the theories [21]. Although there may not be practical payoffs in the near future, this work may ultimately be of the most value.

XII ROBOTICS

From the earliest, AI research has been concerned with the development of systems that integrate many diverse components. In the late 1960s and early 1970s, work was carried out at the major centers on robot systems that could perceive, reason, plan, and manipulate.

At MIT in 1972 a program called COPY could look (using an image-dissector camera) at a simple structure built of children's building blocks and understand it in terms of what shapes the blocks had and how they were put together. It could then look at a collection of blocks spread out on the table, select the ones it needed, and use a manipulator to physically build a mirror image of the structure [41].

At Hitachi, in Japan, the HIVIP system could perform a similar task by looking at a simple engineering drawing of the desired structure instead of the structure itself [11].

In Edinburgh the FREDDY robot system, with TV cameras and a manipulator, could be given a heap of parts that comprised a simple model, could sort out the individual pieces and identify them, and then assemble the model. The assembly was accomplished using touch and force feedback, and it made use of a vise to hold partial assemblies. Recognition of parts was accomplished by showing FREDDY the parts and how to grasp them [2].

At Stanford University a system with a TV camera and manipulator could look at the blocks of an "Instant Insanity" puzzle, determine how to solve the puzzle, and then stack the blocks up, demonstrating the solution [13].

At SRI an ambitious research program was carried out, using a mobile robot--a cart with a TV camera and rangefinder, but no manipulators. SHAKEY, as the system was known, could form and execute plans for navigating through a simple environment of rooms, doorways, and large blocks; its visual system could recognize and locate doorways, floor-wall boundaries, and the large blocks. It had sophisticated techniques to allow it to recover from unforeseen circumstances, and it could store (learn) generalized versions of the plans it produced for future use [29].

It is probably fair to say that, despite some efforts to make them capable of recovery from errors and unforeseen situations, none of the systems was robust or versatile enough to be of practical use. These exercises were nevertheless very valuable in determining which were the significant problems for robots perceiving and acting in the real world. Research has since focused upon the system components, perception, planning, etc. Indeed, the field has become sufficiently complex that no individual can hope to be expert in all areas, and instead he specializes.

Currently, work on robotics is again on the increase, but this time the goals are more practically oriented and the domains more constrained. At a number of centers, advanced industrial automation is being studied. The emphasis is now on practicality and applicability within the next few years. The techniques that are being employed are those that have been around for some time and are well understood; they are being adapted to work fast and reliably on small computers.

An early demonstration of what could be done was the WAVE system developed at Stanford University in 1973. This program used TV and tactile sensing to assemble an automobile water pump, using a power bolt driver to do so [4]. The manipulator control routines were capable of complying with external constraints, a capability that is essential for such operations as turning a crank or opening a door.

At SRI, over the last few years, in work supported by NSF and a consortium of interested companies, a number of realistic demonstrations have been mounted. For example, parts moving on a conveyor belt are observed with a TV camera to recognize them and determine their orientation so that a Unimate arm can then pick them up and pack them in boxes. In another demonstration, the Unimate places and bolts on the cylinder head of a small engine, using vision to locate the bolt holes and verify bolt insertion and correct assembly. A simple approach to bin picking has been shown, lifting castings from a bin with electromagnets, placing them on a table, and then using TV to isolate and locate parts so that they may be picked up individually. Visually guided (simulated) spot welding has also been performed, following lines, grooves, and corners on objects, both stationary and on a moving line [26 and 33].

Various inspection tasks have also been tackled, such as inspecting lamp bases for correct number and location of contacts, inspecting washing machine water pumps to verify presence and determine orientation of an actuating lever, and inspecting electrical boxes on a moving conveyor for deformation and missing knockouts [1].

The visual processing for these robotics and inspection tasks is performed by a self-contained subsystem, developed at SRI, known as the Vision Module [14]. The Vision Module consists of a solid-state TV camera, special preprocessing hardware, an LSI-11 minicomputer, and a software package. It operates by thresholding the input image brightness to produce a binary (black and white only) image. The hardware encodes the binary image efficiently; and the software can form a description of the image in terms of the individual regions and holes in it, with basic statistical measurements of shape, size, etc. From the description of the outlines of regions, objects can be recognized and their positions and orientations determined. The recognition process is effected by showing examples to the camera and naming them.

The Vision Module can form a basis for many industrial applications, from inspection to assembly. So far it has not been used on a factory floor, but copies of it are operating in a number of industrial laboratories.

It seems certain that current work on advanced automation in AI laboratories will find its way into factories within a relatively few years; any doubts on this score should be dispelled by perusing the proceedings of international conferences on industrial robotics [32]. Some products of AI research have been directly transplanted to industrial situations. For example, the PUMA arm being developed by

Unimation is an industrial version of an arm developed by Scheinmann for use in AI research laboratories. Many enlightened companies are helping to support the work in the USA, and some are establishing their own groups to develop and apply approaches pioneered in AI. At General Motors Research Laboratories in Warren, Michigan, for example, a strong research team has demonstrated visually guided mounting of automobile wheels, the automatic inspection of transistor chips for ignition systems, and an initial attempt at simple bin picking [27].

The present blind generation of industrial robots will certainly be superseded by a generation with simple visual capabilities; and, indeed, such robots are already being used experimentally in factory situations, at least in Japan. The basic techniques currently available are cheap, fast, and reliable enough; and there are many applications in which the environment is not sufficiently rigidly controlled to permit "pick and place" dead-reckoning approaches and in which imposing that control would be expensive. In such situations humans are wasted because their perceptual, reasoning, and manipulative capacities far exceed the minimal levels required. Industrial robots cannot yet be made as capable as people, but we have the expertise now to develop them beyond the rudimentary levels of their current existence.

XIII PAST, PRESENT, AND FUTURE

The field of artificial intelligence has matured significantly in the past few years. In the early stages of the 1950s and 1960s, AI research was the domain of a rather small band of enthusiasts, who were regarded rather suspiciously by scientists in more respectable fields. The problems they tackled were small-scale, involving simplified worlds and elementary (though novel) techniques. It was felt to be an achievement when a program managed to carry out an intellectual task at all, even if poorly by human standards. New ground was being broken and the initial successes came comparatively easily, leading to considerable optimism and some rash predictions of future progress.

During the late 1960s and early 1970s, the nature of AI research changed as the field moved into its adolescence. The scopes of experimental domains were broadened, and more realistic situations were investigated. It became evident how difficult it is to perform intellectual tasks well and how truly complex even the simplest task can be, involving many types and levels of knowledge. Programs were judged by the level of competence they displayed, and a more conservative attitude about achievements emerged. Motivations and methodologies of research workers became more clearly identifiable as "science" or "engineering."

The maturation process has continued throughout the 1970s. AI has become an accepted science, capable of providing valuable concepts and practical applications. It is no longer studied at only a handful of laboratories, but at many additional institutions, academic and commercial, in many countries. Already, in the USA, as many PhDs are conferred in AI as in such areas of computer science as programming languages, numerical analysis, or theory of computation [37].

In this paper, we have been concerned primarily with actual accomplishments in AI, particularly with the fruits of the past few years. We have seen that useful systems have been developed or are under development: expert systems exist with capabilities comparable with those of human experts in limited domains; packages exist for providing limited natural language interfaces; data-base query systems can now make information available to nonprogrammers; and simple visual capabilities can be used in industrial situations for inspection and robot control. Enough has been demonstrated that psychologists, linguists, and computer scientists are happy to collaborate with AI scientists and manufacturers, and the Department of Defense with AI engineers.

What, then, can we infer about the future of AI? In general terms, we can expect a steady increase in the speed, scope, and competence of AI systems as the state of the art advances. We can also expect a steady flow of practical systems into the real world.

In the near term, say the next two to five years, instead of a few demonstration programs, we shall see routine application of what already exists. In the USA it would seem that manufacturing and commerce will be among the first areas to benefit. The momentum that exists in producing more sophisticated industrial robots will continue, and we shall see more dexterous manipulators and widespread use of simple visual control. There will be steady progress toward more integrated factories, tying together design, production scheduling, and manufacturing; here computer-aided design, operations research, and AI should blend very nicely.

We can also expect AI to contribute to productivity in offices. I already make regular use of programs that help me create and edit documents, format them, and correct spelling errors. I also rely on the computer to remind me of appointments and help me handle my electronic mail. Most of these facilities originated from centers of AI research, although at present one would not classify them as AI programs. In the next few years we should see considerable expansion of such facilities, with much more knowledge being incorporated into them, and the emergence of systems that can act as personal assistants. Some laboratory experimentation with assistant programs has already been made. It is quite within the state of the art for such programs to help schedule the day, perhaps negotiating with other such assistants to make appointments, remind the user about meetings and deadlines, retrieve information about production or performance, handle accounts, perhaps prescan electronic mail, help draft and format documents, correcting grammar and spelling, and file and cross-reference them. The computer can reliably perform the trivial tasks that are so burdensome to people, thus freeing them for the more human aspects of their work. We shall certainly see more and more applications of AI in industry and commerce in the next few years.

Looking somewhat further (and less reliably) ahead, we can imagine AI firmly established in many areas of society. Expert systems are likely to prove a major area of payoff for AI--they will be able to provide skilled advice in business, law, agriculture, taxation, and design, as well as medicine and mineral exploration. Access via natural language, including speech, to complex systems will make them available to nonspecialists and the public at large, perhaps for the cost of a telephone call. Vision and deduction will liberate machines from their currently primitive level of activity; they will be able to cope with ill-structured situations and accidents and will even be able to roam freely, delivering goods, sweeping up spills, or checking for leaks.

To have made such predictions a few years ago would have required much more faith and might have been an exercise in whistling in the dark. However, the clear evidence of maturation of the field, the performance of application programs, and the state of the ongoing research all give me considerable confidence that we shall see these accomplishments in the next decade, or two at the most.

In addition to the evidence from within the field of AI, there is a very significant external factor, namely our rapid progress up the technological curve that doubles our computing power every two years. Computer processors that would have been considered adequate for a whole nation in the early 1950s are now built into typewriters and milking machines. The free availability of information-processing power will generate the need, the vision, and the motivation for increased efforts in AI. We are at the threshold of an era in which, in our everyday lives, we will use and interact with gadgets and systems that will be much more helpful and versatile, simply because they will contain elementary decision-making capabilities. It is especially in the areas of interaction between people and machines that artificial intelligence will make its greatest contributions, hopefully reversing the pressures for us to conform to the machines' world, and instead enabling them to conform to ours.

ACKNOWLEDGEMENTS

Suggestions and material for this paper came from R. O. Duda, P. E. Hart, B. Moore, N. J. Nilsson, W. T. Park, E. D. Sacerdoti, and J. M. Tenenbaum, all of the SRI Artificial Intelligence group.

REFERENCES

1. G. J. Agin, "An Experimental Vision System for Industrial Application," Proceedings Fifth International Symposium on Industrial Robots, IIT Research Institute, Chicago, Illinois (September 1975).

2. A.P.A. Ambler et al., "A Versatile System for Computer-Controlled Assembly," Artificial Intelligence, Vol. 6, No. 2, pp. 129-156 (Summer 1975).

3. H. G. Barrow and J. M. Tenenbaum, "Recovering Intrinsic Scene Characteristics from Images," in Computer Visual Systems, A. Hansen & E. Riseman, eds. (Academic Press, 1978). Also published as Technical Note 157, Artificial Intelligence Center, SRI International, Menlo Park, California (April 1978).

4. R. C. Bolles and R. Paul, "The Use of Sensory Feedback in a Programmable Assembly System," Memo CS-396, Computer Science Department, Stanford University, Stanford, California (October 1973).

5. J. S. Brown et al., "SOPHIE: A Sophisticated Instructional Environment for Teaching Electronic Troubleshooting (an Example of CAI in AI)," BBN Report No. 2790, Bolt, Beranek, and Newman, Inc., Cambridge, Massachusetts (March 1974).

6. B. G. Buchanan and J. Lederberg, "The Heuristic DENDRAL Program for Explaining Empirical Data," Proc. IFIP Congress, Vol. 71, Ljubljana, Yugoslavia (1971).

7. A. Bundy, Artificial Intelligence: An Introductory Course, (North Holland, New York, 1978).

8. R. O. Duda and P. Hart, Pattern Classification and Scene Analysis, (John Wiley and Sons, New York, 1973).

9. R. O. Duda et al., "Development of the PROSPECTOR Consultation System for Mineral Exploration," Final Report, SRI Projects 5821 and 6415, SRI International, Menlo Park, California (October 1978).

10. R. O. Duda, N. J. Nilsson, and B. Raphael, "State of Technology in Artificial Intelligence," in Research Directions in Software Technology, P. Wegner, ed., pp. 729-749 (MIT Press, Cambridge, Massachusetts, 1979).

11. M. Ejiri et al., "An Intelligent Robot with Cognition and Decision-Making Ability," Proceedings of the Second International Joint Conference on Artificial Intelligence, London, England, pp. 350-358 (1971).

12. E. A. Feigenbaum, "The Art of Artificial Intelligence--Themes and Case Studies of Knowledge Engineering," Proc. Conference FIPS, Vol. 47, National Computer Conference, Anaheim, California (June 5-8, 1978). Also published in Proceedings of Fifth International Joint Conference on Artificial Intelligence, MIT, Cambridge, Massachusetts, pp.1014-1029 (1977).

13. J. Feldman et al., "The Use of Vision and Manipulation to Solve the Instant Insanity Puzzle," Proceedings Second International Joint Conference on Artificial Intelligence, London, England, pp. 359-364 (1971).

14. G. J. Gleason and G. J. Agin, "A Modular Vision System for Sensor-Controlled Manipulation and Inspection," Technical Note 178, Artificial Intelligence Center, SRI International, Menlo Park, California (1977). Also published in Proceedings of the Ninth International Symposium on Industrial Robots, Washington, D.C. (March 1979).

15. L. R. Harris, "ROBOT: A High Performance Natural Language Processor for Data Base Query," Proceedings of the Fifth International Joint Conference on Artificial Intelligence, MIT, Cambridge, Massachusetts, pp. 903-904 (1977).

16. G. G. Hendrix, "LIFER: A Natural Language Interface Facility," Technical Note 135, Artificial Intelligence Center, SRI International, Menlo Park, California (1977).

17. B.K.P. Horn, "Obtaining Shape from Shading Information," in The Psychology of Computer Vision, P. H. Winston, ed., 115-15 (McGraw-Hill, New York, New York, 1975).

18. C. C. Li, A. M. Savol, and C. P. Fong, "An Improved Algorithm for the Detection of Small, Rounded Pneumoconiosis Opacities in Chest X-rays," Proceedings IEEE Conference on Pattern Recognition and Image Processing, Chicago, Illinois, pp. 390-395 (May 1978).

19. B. T. Lowerre, "The HARPY Speech Recognition System," PhD Thesis and Technical Report, Computer Science Department, Carnegie-Mellon University, Pittsburgh, Pennsylvania (1976).

20. M. Marcus, "A Design for a Parser for English," Proceedings A.C.M. National Conference, Houston, Texas, pp. 62-67 (1976).

21. D. Marr, "Representing Visual Information," in Computer Visual Systems, A. Hansen & E. Riseman, eds. (Academic Press, 1978). Also published as MIT AI Lab Memo No. 415, Cambridge, Massachusetts (May 1977).

22. W. A. Martin and R. J. Fateman, "The MACSYMA System," Proc. ACM Second Symposium on Symbolic and Algebraic Manipulation, S. R. Petrick, ed., pp. 23-25 (Los Angeles, California, March 1971).

23. P. McCorduck, "History of Artificial Intelligence," Proceedings of the Fifth International Joint Conference on Artificial Intelligence, Cambridge, Massachusetts, pp. 951-954 (August 1977).

24. N. J. Nilsson, Problem-Solving Methods in Artificial Intelligence, (McGraw-Hill Book Company, New York, New York, 1971).

25. N. J. Nilsson, "Artificial Intelligence," Proc. IFIP Congress '74, Stockholm, Sweden (1974).

26. D. Nitzan, "Flexible Automation Program at SRI," Proceedings of the 1979 Joint Automatic Control Conference, Denver, Colorado, pp. 754-759 (June 1979).

27. W. Perkins, "Model-Based Vision System for Scenes Containing Multiple Parts," Proceedings of the Fifth International Joint Conference on Artificial Intelligence, Cambridge, Massachusetts, pp. 678-684 (1977).

28. H. E. Pople, Jr., "The Formation of Composite Hypotheses in Diagnostic Problem Solving: An Exercise in Synthetic Reasoning," Proceedings Fifth International Joint Conference on Artificial Intelligence, Cambridge, Massachusetts, pp. 1030-1037 (1977).

29. B. Raphael et al., "Research and Applications--Artificial Intelligence," Final Report on SRI Project 8973, Artificial Intelligence Center, SRI International, Menlo Park, California (December 1971).

30. B. Raphael, The Thinking Computer: Mind Inside Matter (W. H. Freeman and Co., San Francisco, California, 1975).

31. L. G. Roberts, "Machine Perception of Three-Dimensional Solids," MIT Lincoln Labs., Lexington, Massachusetts, Technical Report No. 315 (May 1963). Also published in Optical and Electro-Optical Information Processing, J. T. Tippett et al., eds., pp. 159-197 (MIT Press, Cambridge, Massachusetts, 1965).

32. Robots, Proceedings of the Ninth International Symposium on Industrial Robots, Washington, D.C., Dearborn, MI, Society of Manufacturing Engineers (September 1978).

33. C. A. Rosen, "Exploratory Research in Advanced Automation," Fifth Report to the National Science Foundation (Grant GI-38100X1), SRI International, Menlo Park, California (January 1976).

34. E. D. Sacerdoti, "Language Access to Distributed Data with Error Recovery," Technical Note 140, Artificial Intelligence Center, SRI International, Menlo Park, California (February 1977).

35. E. H. Shortliffe, "Computer-Based Medical Consultations: MYCIN," Elsevier, New York (1976).

36. G. J. Sussman, "Electrical Design: A Problem for Artificial Intelligence Research," Proceedings Fifth International Joint Conference on Artificial Intelligence, Cambridge, Massachusetts, pp. 894-900 (1977).

37. O. E. Taulbee and S. D. Conte, "Production and Employment of PhDs in Computer Science," Communications of the ACM, Vol. 19, No. 6, pp. 311-313 (June 1976).

38. A. M. Turing, "Intelligent Machinery," in Machine Intelligence, Vol. 5, B. Meltzer and D. Michie, eds. (Edinburgh University Press, Edinburgh, and American Elsevier, New York, 1970). (Written in 1947, but previously unpublished.)

39. D. C. Waltz, "Generating Semantic Descriptions from Drawings of Scenes with Shadows," AI TR-271, Artificial Intelligence Laboratory, MIT, Cambridge, Massachusetts (August 1972). Also published in The Psychology of Computer Vision, P. H. Winston, ed., pp. 1-91 (McGraw-Hill Book Co., New York, New York, 1975).

40. T. Winograd, "Procedures as Representation for Data in a Computer Program for Understanding Natural Language," Technical Report AI TR-17, MIT, Cambridge, Massachusetts (1971). Also published as Understanding Natural Language (Academic Press, New York, New York, 1972).

41. P. H. Winston, "The MIT Robot," in Machine Intelligence, Vol. 7, B. Meltzer and D. Michie, eds., (Edinburgh University Press, Edinburgh, and American Elsevier Publishing Co., New York, 1972).

42. P. H. Winston, Artificial Intelligence (Addison-Wesley, Reading, 1977).

43. P. H. Winston and R. H. Brown, Artificial Intelligence: An MIT Perspective, (MIT Press, Cambridge, Massachusetts, 1979).

44. W. A. Woods, R. M. Kaplan, and B. Nash-Webber, "The Lunar Sciences Natural Language System: Final Report," Report No. 2378, Bolt, Beranek, and Newman, Inc., Cambridge, Massachusetts (1972).