Essays of an Information Scientist, Vol:6, p.430-442, 1983 Current Contents, #52, p.5-17, December 26, 1983

Current Comments®

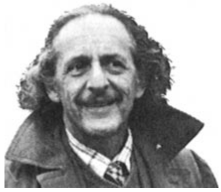

Artificial Intelligence: Using Computers to Think about Thinking. Part 2. Some Practical Applications of AI

Number 52 December 26, 1983

This is the second of a two-part essay on artificial intelligence (AI). Part 1 discussed "knowledge representations"—models of cognition that AI investigators have used in their attempts to build thinking machines.[1] One of the most ambitious goals of AI research is to develop programs which enable machines to perform commonsense reasoning. It will be many years, however, before this goal is achieved.

In the meantime, AI research has spawned numerous spin-offs that are just beginning to enter the commercial market. These include programs that enable robots to "see" and "feel" and machines which can follow instructions written in natural language (NL). But the most ambitious and successful AI applications thus far are "expert systems." These computer programs are designed to duplicate the problem-solving processes of experts in various fields. This essay will cover some of these expert systems and review a limited number of other applications of AI.

During the early years of AI research, investigators were intent upon discovering the general principles underlying intelligence. In developing the knowledge representations described in Part 1, AI researchers tried to use these principles to create an "inference engine." Ideally, such an engine would solve any type of problem, from winning at chess to diagnosing disease. But attempts at developing computer-based problem-solving strategies for any situation met with limited success. Techniques developed for solving problems in one domain were usually inadequate for other domains. However, programs equipped with a great deal of information about a single domain performed as well as, and sometimes better than, experts in that field.[2] The superior performance of these knowledge-based, or expert, systems convinced many AI researchers that problem solving demands huge banks of knowledge as well as reasoning procedures.[3]

Today, the goal of expert systems research is to transfer a specialist's knowledge into a program so the information can be efficiently accessed and used by the computer to solve problems. This includes "textbook learning"—the facts obtained from training and reading.[4] It also includes heuristic knowledge, or rules of thumb developed through years of experience and judgment. Heuristics are essentially educated guesses about which solutions to a problem are most likely to be successful. They don't guarantee correct answers. But they do save time by limiting the search for solutions to those most likely to be correct. Computer and other scientists called "knowledge engineers" sometimes spend years picking experts' brains for these facts and heuristics, and then structure them into computer programs.

Expert systems also include "inference procedures" that determine which heuristics and facts should be brought to bear on a problem. One such inference procedure is backward chaining, in which you suggest a possible solution to a problem and work backward to see if it's correct. MYCIN, an expert system that assists in medical diagnoses, reasons in this manner.[2] It advances a disease hypothesis based on a few known symptoms. Then it looks for other symptoms which support the hypothesis, requesting additional tests and information as needed. In most expert systems, the inference procedure for deciding which facts and heuristics to use is separate from the knowledge base of facts and rules. This makes it easier to add or modify facts and rules as new information becomes available.

The first, and probably best-known, AI programs that focused on a limited domain were chess programs. Donald Michie, University of Edinburgh, Scotland, explains that chess is ideal for modeling specialist knowledge because it is a very well-defined domain.[5] A large amount of formal information is available in the form of instructional works and commentaries. And numerical scales of performance are available in the national and international rating systems. Equally important, chess is a game that calls on a wide range of cognitive functions, from logical calculation to imaginative thinking. The numerous chess-playing programs developed in the 1950s and 1960s tested the proficiency with which various knowledge representations used facts and heuristics to solve problems.[5]

By the mid-1960s, AI researchers began to expand beyond chess and puzzle playing, or what Edward A. Feigenbaum, Stanford University, California, calls "toy problems,"[6] (p. 62) into practical problems. The first such project resulted in DENDRAL, an expert system that identifies the chemical structures of unknown compounds.[7,8] DENDRAL was launched at Stanford University in 1966 by Feigenbaum and Joshua Lederberg, now president of Rockefeller University, New York. Both were interested in modeling and assisting scientific thinking, and Lederberg, a geneticist/molecular biologist with expertise in chemistry, had developed a computer language for describing the structure of complex molecules.[9] As the project grew, Carl Djerassi, also at Stanford University, contributed his expertise.[10]

The original DENDRAL program performed three basic operations on the chemical formulas and mass spectral data with which it was provided. In the first phase, called "plan," it translated general and specific prior knowledge and heuristics into a specific repertoire of constraints. In the next phase, called "generation," it generated plausible structures based on such constraints as the number of rings, double bonds, and atoms of various types in each molecule. In the final phase, called "test," each plausible structure was tested. The computer first generated sets of instrument data that would be expected to describe each structure. Then it compared each set of data to actual data about the compound. The closest fits were then ranked for the user.

In the past few years, scientists working on the DENDRAL project have focused most of their attention on the planning and generation portion of the program, and on making this portion available to users. Called CONGEN, for constrained structure generation, this generator has been expanded to infer plausible structures using a wide variety of instrumental data. In 1982, CONGEN was made commercially available through the computer network Compuserve.[11]

Another spin-off of DENDRAL is Meta-DENDRAL, a program that generates its own rules from mass spectral data on chemical compounds. After receiving mass spectral data on a family of compounds, Meta-DENDRAL generates planning and test rules that describe how these compounds fragment when studied with mass spectrometry. Some of the rules generated by Meta-

DENDRAL duplicated those formulated by expert chemists, while others were entirely original.[12]

DENDRAL demonstrated that AI techniques could be used to solve real problems within a limited area of knowledge.[6] Paradoxically, it also demonstrated that it is easier to model the reasoning processes of specialists than to program the steps a child goes through in understanding language, or making commonsense inferences.[2] This is because the facts and judgments an expert uses in making a decision are easier to identify and categorize than are the reasoning processes used for general problem solving.

So far, the most successful expert systems have been programs that weigh and balance evidence about data to determine how they should be categorized. Differential diagnosis, for example, is "a classical medical example of such a problem,"[2] according to Richard O. Duda, Syntelligence, Menlo Park, California, and Edward H. Shortliffe, Stanford University School of Medicine. A physician arrives at a diagnosis by evaluating a variety of symptoms and test results. Although this is a fairly complicated procedure, it is based on identifiable facts and heuristics and, therefore, lends itself to computer modeling. For this reason, and because computers can consider many diseases that physicians might not encounter in everyday practice, numerous expert systems have been designed to assist doctors in diagnosing and treating disease.[2]

The knowledge representation used most widely in these expert or "consultation" systems is the "production rule" approach, according to William B. Gevarter, National Aeronautics and Space Administration, Washington, DC.[13] As mentioned earlier, each production rule is a heuristic, also called an "if-then" rule or "condition-action" pair. Each rule or set of rules includes facts about a domain which can be used to solve problems in that domain. Figure 1 shows one of the 500 such rules used by the infectious diseases expert system MYCIN. With MYCIN, a physician enters information about a patient into the computer. The computer then searches for the rules that can be applied to this information. If more information is needed, the computer will ask the physician to supply it. Then the rules for these additional data are applied. This process continues until a diagnosis, and treatment, can be recommended.

Since the acceptability of an expert system depends on the confidence with which physicians can accept its suggestions, MYCIN and several other consultation systems can provide explanations of their reasoning processes. At the physician's request, MYCIN will list the production rules it used in its diagnosis or treatment recommendation, and cite references to the literature that support these rules. MYCIN will also explain why it has requested additional tests or other information.[2] An example of this explanation feature is shown in Figure 2.

Figure 1: Sample MYCIN production rule.

Mycin

If: 1) The infection which requires therapy is meningitis,

And 2) The patient has evidence of a serious skin or soft tissue infection,

And 3) Organisms were not seen on the stain of the culture,

And 4) The type of infection is bacterial,

Then:

There is evidence that the organism (other than those seen on cultures or smears) which might be causing the infection is staphylococcus-coagpos (.75) or streptococcus (.5).

Source: **Gevarter W B.** Expert systems: limited but powerful. *IEEE Spectrum* 20(8):39-45, 1983.

Figure 2: Explanation features of MYCIN. This excerpt from the MYCIN run illustrates some of the explanation facilities provided by the program. At entry 38, MYCIN asks whether the patient has been seriously burned. Instead of answering, the user asks why this particular question is being asked. In response, the program states both the goal of the question and the rule it was attempting to apply. This explanation leads the user to wonder how MYCIN established that the infection might be bacterial; in response to the HOW command, the program lists the five rules that it applied, four of which supported this conclusion. This ability to inspect the program's reasoning methods provides the user with a firmer basis for understanding the final conclusions.

```
34) Is the meningitis a hospital-acquired infection?
** NO
35) Is Pt219's clinical history consistent with EPIGLOTTITIS?
** NO
36) Is Pt219's clinical history consistent with OTITIS-MEDIA?
** NO
37) Has the patient ever been to an area that is endemic for coccidiomycoses?
** UNKNOWN
38) Is Pt219 a burn patient?
** WHY

[i.e. WHY is it important to determine whether PATIENT-219 has been seriously burned?]

[4.0] This will aid in determining the organisms (other than those seen on cultures or smears) which
    might be causing the infection.
  It has already been established that
    [4.1] the infection which requires therapy is meningitis, and
    [4.2] organisms were not seen on the stain of the pending csf culture (CULTURE-1), and
    [4.3] the type of the infection may be bacterial
  Therefore, if
    [4.4] Pt219 has been seriously burned
  then
    there is suggestive evidence that Pseudomonas-aeruginosa is one of the organisms (other than those
    seen on cultures or smears) which might be causing the infection
    Author: Yu
    [RULE578]

    [back to question 38...]
    ** HOW 4.3

[i.e. HOW was it established that the type of the infection may be bacterial?]

The following rules concluded about: the likelihood that the type of the infection may be bacterial

                 Rule        Yes       No

    [4.5]        RULE500     (.4)
    [4.6]        RULE501               (.17)
    [4.7]        RULE502     (.1)
    [4.8]        RULE524     (.42)
    [4.9]        RULE504     (.28)

Thus, it has been established that the type of the infection may be bacterial.
```

Source: Duda R O & Shortliffe E H. Expert systems research. *Science* 220(4594):261-8, 1983.

MYCIN was developed in the mid-1970s at Stanford University by Bruce Buchanan and Shortliffe.[2] In the late 1970s, W.J. van Melle, Stanford University, transformed MYCIN's inference engine into an all-purpose program, called EMYCIN.[14] Since then, EMYCIN has been filled with production rules from a variety of fields, including structural engineering and computer repair. Another MYCIN spin-off, called GUIDON,[15] was developed by W.J.

Clancey and colleagues, Stanford University, to help medical students learn the methods and strategies used by experts in their field.

DENDRAL and MYCIN are two of the most influential expert systems. But they represent just a fraction of the work currently under way in this field. Dozens of systems have been, or are now being, developed at Edinburgh, Purdue, Rutgers, Yale, and Carnegie-Mellon Universities, Massachusetts Institute of Technology (MIT), and the University of Pittsburgh, to name only a few of the academic institutions involved in basic and applied AI research. The leading center for expert systems research in the US, however, is Stanford University, which houses the Stanford University Medical Experimental Computer for Artificial Intelligence in Medicine (SUMEX-AIM) network, a nationally shared computer network devoted to AI systems in biomedicine. Systems that have come out of this project include PUFF,[16] developed by respiratory specialists J. Osborn, R.J. Fallat, and B. Votteri, Pacific Medical Center, San Francisco, and Stanford University computer scientists L. Fagan, P. Nii, J.C. Kunz, J.S. Aikins, and D. McClung. PUFF diagnoses and recommends therapies for pulmonary dysfunction. PARRY, a program that simulates paranoid thought processes, was developed on SUMEX-AIM by K.M. Colby, University of California, Los Angeles,[17] and CASNET/Glaucoma was designed for ophthalmology by C.A. Kulikowski and S.M. Weiss, Rutgers University, New Brunswick, New Jersey, and A. Safir, Mount Sinai School of Medicine, New York.[18] INTERNIST, a pioneering program developed by H.E. Pople and J.D. Myers, University of Pittsburgh, Pennsylvania, diagnoses diseases in internal medicine.[19,20] It includes some 4,000 symptoms cross-referenced to about 500 diseases.[21]

Not all expert systems are biomedical. TEIRESIAS,[22] developed by R. Davis, now of MIT, while he was at Stanford University, is designed to assist knowledge engineers in creating and updating the knowledge bases used in expert systems. The success of TEIRESIAS, and other programs that "build" expert systems, is crucial to the commercial future of these systems. Knowledge engineers presently spend hundreds of hours collecting and structuring specialist knowledge before an expert system can be built. Other non-biomedical expert systems include PROSPECTOR,[23] a geology consultant developed by Duda and colleagues, while at SRI International, Menlo Park, California, and MACSYMA,[24] designed by Joel Moses, MIT, for solving algebra and calculus problems.

A number of private firms have also entered the AI arena. Digital Equipment Corp. and International Business Machines Corp. (IBM) are working on separate expert systems that diagnose "sick" computers.[3] Xerox Corp. and Texas Instruments are developing systems to assist in the design of computer chips. And Schlumberger Ltd., which employs many AI researchers in its three AI laboratories, is developing a system that analyzes data on geologic formations.[3]

Several companies have also been launched by established AI scientists. Feigenbaum and several of his colleagues at Stanford University started Teknowledge Inc. and Intelligenetics Inc., both in Palo Alto, California. Teknowledge markets expert systems, and offers consulting and training services to companies that want to build their own. Intelligenetics markets expert systems that assist in gene splicing experiments. Computers designed for AI research are sold by two companies, Lisp Machines, Inc., and Symbolics Inc., which were founded by MIT researchers.

Expert systems aren't the only AI applications entering the marketplace. Machine Intelligence Corp., Optical Recognition Systems Inc., and General Electric Co. are among a growing number of companies marketing software for recognizing visual images.[25] Artificial Intelligence Corp. and Cognitive Systems Inc. are marketing NL software that enables users to communicate with computers in natural languages such as written English.[3] The state of the art for these applications is not as advanced in dealing with technical domains as are expert systems.[13] But these programs can function impressively within the limited domains for which they've been designed.

Part 1 of this essay reviewed some of the problems AI researchers confronted while developing NL programs. At that time, I noted that language is a highly complex cognitive function, involving much more than the syntactical manipulation of words. Creating a program that can communicate effectively requires the inclusion of the assumptions, or large body of shared knowledge, speakers generally bring to a conversation. Colby points out that information about the context in which language is used and the speaker's intentions are also required.[26] Consequently, as with expert systems, the most successful NL programs have been those equipped with a great deal of information about a restricted topic. One of the first NL programs to operate in such a domain was SHRDLU, developed by Terry Winograd, Stanford University. It answered questions about an imaginary world of blocks, pyramids, and an arm that moved them.[27] A variety of systems have also been built for studying NL and for querying expert systems. Commercial interest in NL, however, is focused on generic systems which can be filled with knowledge about the data bases to which they are linked.

The knowledge in most NL systems consists of a pragmatic or semantic grammar.[28] These are the rules for parsing, or breaking apart, and interpreting sentences likely to be used to request information from a particular data base. LIFER,[29] an NL program developed by G.G. Hendrix while at SRI International, for example, interprets the sentence, "What is the age of Mary's sister," by matching it against a series of grammatical questions, such as, "What is the attribute of the object?" Most NL programs are also equipped with contextual information about the data base. This includes definitions of words likely to be used in queries, and information on relationships between objects in the data base. Such information helps the program deal with ambiguities such as incomplete sentences (ellipses) and pronouns. An example of a conversation in which LADDER,[30] also developed by Hendrix and colleagues while at SRI, responds to incomplete sentences is shown in Figure 3. Finally, several NL programs will paraphrase the user's request so it can be confirmed. Once confirmed, the system translates the request into a formal computer language, which is then used to retrieve information from the data base.

Figure 3: An actual conversation with the LADDER system. Boldface words represent user input.

1—Give me the length of the Kenedy.
spelling > KENNEDY
PARSED!
Give the length of the ship KENNEDY JF

LEN = 1072
2—Width and draft
PARSED!
Trying Ellipsis: GIVE ME WIDTH AND DRAFT OF THE KENEDY
Give the beam and draft of the ship KENNEDY JF
(BEAM 130 DFT 36)
3—Who is her commander?
Give the rank and name of the commander of the ship KENNEDY JF
(RANK CAPT NAME MOFFETT P)
4—What about the Fox?
PARSED!
Trying Ellipsis on: THE FOX
WHO IS THE FOX COMMANDER
Give the rank and name of the commander of the ship FOX
(RANK CAPT NAME EVERETT J)

Source: **Hendrix G G & Sacerdoti E D.** Natural-language processing: the field in perspective. *BYTE* 6:304-52, 1981.

Whereas NL systems match written statements against grammatical rules and contextual knowledge, the vision programs developed by AI investigators match images picked up by a camera, laser, or light emitting diode (LED) against stored representations of those images. These systems examine scenes, and operate by reducing their colors, shapes, and textures into "primal sketches," which are essentially simple line drawings. Then the program "blocks out" a rough approximation of the scene. This consists of simple shapes, such as lines, cylinders, and cones, that the computer has been programmed to identify. Finally, the computer attempts to match these "mental images" of objects against three-dimensional sketches of objects it has been programmed to identify.[31]

Although these vision programs are in use for industrial quality control and inspection, most are limited to recognizing objects from only one perspective. Only a few can identify moving objects, or distinguish objects when background lighting changes. However, in situations where these factors are controlled, computer vision systems can inspect objects with remarkable speed and accuracy. Machine Intelligence Corp. offers one system that can inspect stationary parts for small dimensional defects at 900 parts per minute. One of the most advanced systems, the Optomation II from General Electric Co., can inspect randomly oriented parts at the same rate.[25]

ACRONYM,[32] an experimental program developed by R.A. Brooks, Stanford University, can identify objects in different configurations and from perspectives it has never seen before. It can also "guess" the identity of an object based on a partial view of it. It matches the visible portion of the object with its stored model of the corresponding portion of that object. Then it makes assumptions about what it sees.

Vision programs will initially have their greatest application in robotics. AI researchers have also contributed to robotics by creating programs that recognize objects by their physical touch as well as programs that understand a limited number of spoken words.[33] Several of these speech recognition systems—including HARPY, which was designed for document retrieval—were discussed in an earlier essay.[34] A number of researchers, including pioneer John McCarthy, Stanford University, believe that one day AI may even make possible a "universal manufacturing machine."[35] Essentially, this would be a robot capable of tailoring products to each person's design.

So many AI applications are now entering the marketplace that it's impossible to name all of them in this brief essay. These include educational programs for teaching young children math and more sophisticated programs for teaching medical students to reason like specialists. AI is also used in developing automatic programming systems. Like high-level computer languages, they will relieve programmers from the drudgery of specifying in detailed machine language what the computer should do.[36] And researchers at University of California, Los Angeles, have developed a speech prosthesis that helps patients who have difficulty recalling words.[37] I've already discussed how AI has made it possible to query data bases in NL. But AI techniques are also used to make data base queries more explicit, and to improve the efficiency with which information is retrieved by these systems.

ISI® has been using AI techniques to improve its data bases for some time. Our programs for verifying bibliographic information perform tasks that ordinarily require intelligence. These programs, so to speak, "decide" how to edit citations containing incorrect information, such as the volume, year, or the spelling of the author's name. When a reader examines a group of citations to the same paper, he or she can quickly recognize their similarities and identify the errors.

Similarly, our "expert" programs can recognize which version of a multicited work is accurate and can correct errors.[38]

While a great deal of judgment is built into our computer programs for verifying references in *Science Citation Index*® (*SCI*®), they only scratch the surface of the problem of artificially intelligent indexing. At a 1964 symposium of the National Bureau of Standards, I presented a paper entitled, "Can citation indexing be automated?", that is, "the capability of the computer automatically to simulate human critical processes reflected in the act of citation."[39]

The paper pointed out that "a considerable standardization of document presentations will be necessary, and probably not achievable for many years if we are to achieve automatic referencing.... On the other hand, many citations, now fortuitously or otherwise omitted, might be supplied by computer analysis of text."[39]

When this paper was reprinted in *Current Contents*® in 1970, I explained that the original title was badly chosen—that a more appropriate title was, "Can criticism and documentation of research papers be automated?"[40] Even though the term "artificial intelligence" was in use by 1970, it was new enough so that it was not obvious to use it in the title of the essay, although the original paper does mention an "artificially intelligent machine." It is significant that a week earlier, in a tribute to Lederberg,[41] I referred to his work, "Applications of artificial intelligence for chemical inference."[42]

ISI's program for classifying documents in *ISI/BIOMED*®, *ISI/CompuMath*®, *ISI/GeoSciTech*™, and *Index to Scientific Reviews*™ also simulates human judgment.[43] Traditional indexing requires the use of human indexers to classify documents by subject. Our system establishes the classification system algorithmically by co-citation analysis, and then assigns each paper to one or more categories. These programs also assess the relevance of each new paper. A current paper that cites many of the core works is given a higher relevance weight than a paper that cites only one or two.

Researchers in Japan are also working on AI-based programs for classifying documents.[44] And the work of R.S. Marcus, MIT, is interesting in connection with expert systems.[45]

Part 1 of this essay included a list of the AI research fronts identified through our *ISI/CompuMath* clustering programs. The core documents for *ISI/CompuMath* research front #80-0191, "Retrieval processes, computational linguistics, and language processing," are listed in Table 1. To show you how the core papers in Table 1 are related, we've also included a multidimensionally scaled map in Figure 4. The paper by A.M. Collins and M. Ross Quillian, then of Bolt Beranek & Newman, Inc., Cambridge, Massachusetts, plays a central role in the research front. It discusses the way humans and computers store and retrieve information from long-term memory. If you decide to search this particular research front through the ISI Search Network, you'll retrieve about 80 papers published between 1976 and 1983.

W.W. Bledsoe, University of Texas, Austin, who wrote the three core papers in research front #80-0724,[46-48] has played an important role in nonresolution theorem proving, also called natural deduction. This type of theorem-proving program uses heuristics to speed up parts of the proof.

Research front #80-0726, "Computer-aided diagnosis and clinical judgment," includes the work of Howard L. Bleich, Beth Israel Hospital, Boston.[49] In an earlier essay,[50] I discussed the Paperchase system,[51] which he developed with Gary L. Horowitz. The other core paper is by J.E. Overall, Kansas State University, Manhattan, and C.M. Williams, University of Florida, Gainesville,

Figure 4: A cluster map of *ISI/CompuMath®* research front #80-0191 "Retrieval processes, computational linguistics, and language processing."

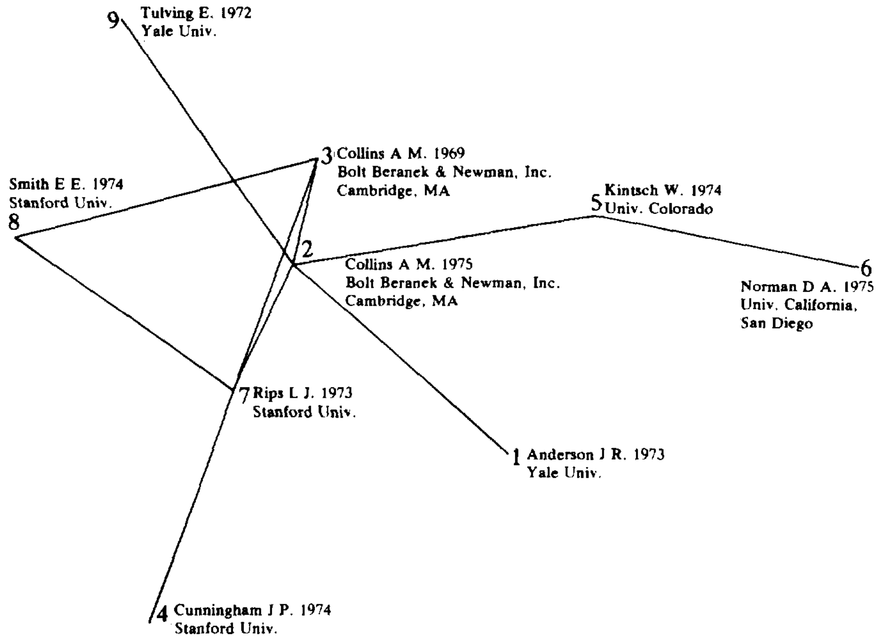

Table 1: Core papers of *ISI/CompuMath®* research front #80-0191 "Retrieval processes, computational linguistics, and language processing."

Anderson J R & Bower G H. *Human associative memory.* New York: Wiley, 1973. 524 p.
Collins A M & Loftus E F. A spreading-activation theory of semantic processing. *Psychol. Rev.* 82:407-28, 1975.
Collins A M & Quillian M R. Retrieval time from semantic memory. *J. Verb. Learn. Verb. Behav.* 8:240-7, 1969.
Cunningham J P & Shepard R N. Monotone mapping of similarities into a general metric space. *J. Math. Psychol.* 11:335-63, 1974.
Kintsch W. *The representation of meaning in memory.* New York: Wiley, 1974. 279 p.
Norman D A, Rumelhart D E & LNR Research Group. *Explorations in cognition.* San Francisco, CA: W.H. Freeman, 1975. 430 p.
Rips L J, Shoben E J & Smith E E. Semantic distance and the verification of semantic relations. *J. Verb. Learn. Verb. Behav.* 12:1-20, 1973.
Smith E E, Shoben E J & Rips L J. Structure and process in semantic memory: a featural model for semantic decisions. *Psychol. Rev.* 81:214-41, 1974.
Tulving E & Donaldson W, eds. *Organization of memory.* New York: Academic Press, 1972. 423 p.

Table 2: Core papers of *ISI/CompuMath®* research front #80-0739 "Nonrecursive grammars, natural languages, and inductive inference of formal languages."

Blum L & Blum M. Toward a mathematical theory of inductive inference. *Inform. Contr.* 28:125-55, 1975.
Feldman J. Some decidability results on grammatical inference and complexity. *Inform. Contr.* 20:244-62, 1972.
Gold E M. Language identification in the limit. *Inform. Contr.* 10:447-74, 1967.
Horning J J. *A study of grammatical inference.* PhD dissertation. Stanford, CA: Stanford University, Department of Computer Science, August 1969. No. CS-139; AI memo-98. NTIS/PC A08 MF A01.
Solomonoff R J. A formal theory of inductive inference. Part I. *Inform. Contr.* 7:1-22, 1964.
Solomonoff R J. A formal theory of inductive inference. Part II. *Inform. Contr.* 7:224-54, 1964.

on a program for diagnosing thyroid function.[52] About 18 current papers can be retrieved by searching in this research front.

Another research front, #80-1963, "Knowledge-engineering and computer-aided medical decision-making," zeros in on the work of G.A. Gorry, MIT, and G.O. Barnett, Massachusetts General Hospital, Boston,[53] and H.R. Warner and colleagues, University of Utah, Salt Lake City.[54] These papers focus on congenital heart disease and other disease diagnoses using the computer.

In Table 2, we've listed the six core papers that define research front #80-0739, "Nonrecursive grammars, natural languages, and inductive inference of formal languages." There were about 69 papers published on these themes. Research front #80-1155 identifies the field, "Cognition, psychological epistemology, and experiments in artificial intelligence." The two core works here include Noam Chomsky's *Language and Mind*[55] and a paper by R.C. Schank, Yale University, New Haven, Connecticut, on conceptual dependency,[56] both published in 1972.

Table 3 presents a selected list of highly cited books and articles on AI. This list also shows how often each of these publications was cited in *SCI* and *Social Sciences Citation Index*® from 1961 to 1983. Three of the publications on this list, Minsky's paper in *The Psychology of Computer Vision*, Quillian's "Semantic memory," and Schank and Abelson's *Scripts, Plans, Goals and Understanding*, focus on knowledge representations discussed in Part 1 of this essay. The paper by Waltz and the book by Winograd are excellent reviews of computer vision and NL, respectively, while Simon's book is a short, nontechnical discussion of AI and related topics. The most-cited publication is the 1965 book by N.J. Nilsson, SRI International, which reports early work on techniques that enable machines to learn by classifying and evaluating information. His more recent book, *Problem-Solving Methods in Artificial Intelligence*, is an influential text on theorem proving and

Table 3: A selected list of highly cited publications in artificial intelligence. A = number of citations from *SCI*®, 1961-1983, and *SSCI*®, 1966-1983. B = bibliographic data.

A	B
70	**Bobrow D G & Winograd T.** An overview of KRL, a knowledge representation language. *Cognitive Sci.* 1:3-46, 1977.
49	**Boden M A.** *Artificial intelligence & natural man.* New York: Basic, 1977. 537 p.
44	**Brown J S & Burton R R.** Diagnostic models for procedural bugs in basic mathematical skills. *Cognitive Sci.* 2:155-92, 1978.
56	**Clowes M B.** On seeing things. *Artif. Intell.* 2:79-116, 1971.
40	**Kosslyn S M & Shwartz S P.** A simulation of visual imagery. *Cognitive Sci.* 1:265-95, 1977.
45	**McCarthy J & Hayes P J.** Some philosophical problems from the standpoint of artificial intelligence. (Meltzer B & Michie D, eds.) *Machine intelligence. 4.* New York: American Elsevier, 1969. p. 463-502.
119	**Minsky M, ed.** *Semantic information processing.* Cambridge, MA: MIT Press, 1969. 440 p.
154	**Minsky M.** A framework for representing knowledge. (Winston P H, ed.) *The psychology of computer vision.* New York: McGraw-Hill, 1975. p. 211-77.
45	**Newell A, Shaw J C & Simon H A.** Report on a general problem-solving program. *Information processing. Proceedings of the International Conference on Information Processing,* 15-20 June 1959, Paris, France. Paris: UNESCO, 1960. p. 256-64.
462	**Nilsson N J.** *Learning machines.* New York: McGraw-Hill, 1965. 137 p.
263	**Nilsson N J.** *Problem-solving methods in artificial intelligence.* New York: McGraw-Hill, 1971. 255 p.
143	**Quillian M R.** Semantic memory. (Minsky M, ed.) *Semantic information processing.* Cambridge, MA: MIT Press, 1969. p. 227-70.
350	**Schank R C & Abelson R P.** *Scripts, plans, goals and understanding.* Hillsdale, NJ: Lawrence Erlbaum, 1977. 248 p.
119	**Shortliffe E H.** *Computer-based medical consultations, MYCIN.* New York: Elsevier, 1976. 264 p.
413	**Simon H A.** *The sciences of the artificial.* Cambridge, MA: MIT Press, 1969. 123 p.
75	**Waltz D.** Understanding line drawings of scenes with shadows. (Winston P H, ed.) *The psychology of computer vision.* New York: McGraw-Hill, 1975. p. 19-91.
186	**Winograd T.** *Understanding natural language.* New York: Academic Press, 1976. 195 p.
107	**Winston P H, ed.** *The psychology of computer vision.* New York: McGraw-Hill, 1975. 282 p.
65	**Winston P H, ed.** *Artificial intelligence.* Reading, MA: Addison-Wesley, 1977. 444 p.

AI search techniques. It includes a detailed discussion of heuristic search methods.

Although university researchers have been involved in AI for some 30 years, a number of governments are just now beginning to recognize the enormous impact AI may have on industrial progress. The UK, for example, recently committed itself to a five-year, $300 million investment in university and industrial research on advanced information technology, including AI.[57] A multinational collaboration, called the European Strategic Program for Research in Information Technology (ESPRIT), has also been proposed by the European Common Market.[6] And though little literature on AI has been coming out of the Soviet Union, where AI used to fall under the rubric of cybernetics, computer scientists there have been programming with List Processor (LISP), the primary programming language for AI, since the 1960s.[58]

At present, the US is the leader in AI research and development.[6] This dominant position is at least partially due to two decades of support from the US Defense Department's Advanced Research Projects Agency (Darpa). The National Institutes of Health Biotechnology Resource Program of the Division of Research Resources also funds AI research, most notably by supporting the SUMEX-AIM computer system. In 1982 alone, governmental agencies and private firms channeled some $50 million into AI research.[6] And Darpa has just launched a massive new AI initiative, called "Strategic Computing and Survivability." According to Feigenbaum,[21] Darpa will spend about $50 million on this project in the 1983-1984 fiscal year and increase funding to a level approaching $200 million during the tenth, and final, year of the program.

This represents a substantial investment in AI. But in their recently published book, *The Fifth Generation*, Feigenbaum and Pamela McCorduck warn that the US may soon lose its two- to three-year lead to the Japanese.[6] In 1982, that country embarked on an ambitious national ten-year plan to lead the world in computer technology, including AI. Japanese government and industry officials plan to spend about $800 million over the next ten years on this joint government-industry effort. Feigenbaum and McCorduck believe that unless the US launches a comparable nationwide, collaborative effort, the Japanese may ultimately become the dominant factor in the computer world.

The AI applications discussed in this essay demonstrate how basic research eventually has practical results.[59] Equipping computers with commonsense reasoning and general problem-solving techniques still remains a goal of basic AI research. Undoubtedly, there will be new spin-offs as we teach computers to "think." Possibly of greater importance, we may learn more about how humans use their brains to think and thereby create what we call knowledge or natural intelligence. It remains to be seen precisely where the boundary between artificial and natural intelligence begins or ends.

* * * * *

My thanks to Joan Lipinsky Cochran and Amy Stone for their help in the preparation of this essay.

REFERENCES

1. **Garfield E.** Artificial intelligence: using computers to think about thinking. Part 1. Representing knowledge. *Current Contents* (49):5-13, 5 December 1983.
2. **Duda R O & Shortliffe E H.** Expert systems research. *Science* 220(4594):261-8, 1983.
3. Artificial intelligence. The second computer age begins. *Bus. Week* (2729):66-75, 1982.
4. **Roberts S K.** Artificial intelligence. *BYTE* 6(9):164-78, 1981.

5. **Michie D.** Computer chess and the humanization of technology. *Nature* 299:391-4, 1982.
6. **Feigenbaum E A & McCorduck P.** *The fifth generation.* Reading, MA: Addison-Wesley, 1983. 275 p.
7. **Lederberg J & Feigenbaum E A.** Mechanization of inductive inference in organic chemistry. (Kleinmuntz B, ed.) *Formal representation of human judgment.* New York: Wiley, 1968. p. 187-218.
8. **Lindsay R K, Buchanan B G, Feigenbaum E A & Lederberg J.** *Applications of artificial intelligence for chemical inference: the DENDRAL project.* New York: McGraw-Hill, 1980. 256 p.
9. **Lederberg J.** Topological mapping of organic molecules. *Proc. Nat. Acad. Sci. US* 53:134-9, 1965.
10. **Garfield E.** A tribute to Carl Djerassi: reflections on a remarkable scientific entrepreneur. *Essays of an information scientist.* Philadelphia: ISI Press, 1983. Vol. 5. p. 721-30. (Reprinted from: *Current Contents* (42):5-14, 18 October 1982.)
11. **Alexander T.** Practical uses for a "useless" science. *Fortune* 105:138-45, 1982.
12. **Feigenbaum E A.** The art of artificial intelligence—themes and case studies of knowledge engineering. *AFIPS Conf. Proc.* 47:227-40, 1978.
13. **Gevarter W B.** Expert systems: limited but powerful. *IEEE Spectrum* 20(8):39-45, 1983.
14. **van Melle W J.** *A domain-independent system that aids in constructing knowledge-based consultation programs.* PhD dissertation, Stanford University, 1980.
15. **Clancey W J.** *Transfer of rule-based expertise through a tutorial dialogue.* PhD dissertation, Stanford University, 1979.
16. **Aikins J S, Kunz J C, Shortliffe E H & Fallat R J.** PUFF: an expert system for interpretation of pulmonary function data. *Comput. Biomed. Res.* 16(3):199-208, 1983.
17. **Colby K M.** Modeling a paranoid mind. *Behav. Brain Sci.* 4:515-60, 1981.
18. **Weiss S M, Kulikowski C A & Safir A.** A model-based consultation system for the long-term management of glaucoma. *IJCAI-77. Proceedings of the 5th International Joint Conference on Artificial Intelligence*, 22-25 August 1977, Cambridge, MA. Pittsburgh, PA: Carnegie-Mellon University, 1977. Vol. 2. p. 826-32.
19. **Pople H.** The formation of composite hypotheses in diagnostic problem solving: an exercise in synthetic reasoning. *IJCAI-77. Proceedings of the 5th International Joint Conference on Artificial Intelligence*, 22-25 August 1977, Cambridge, MA. Pittsburgh, PA: Carnegie-Mellon University, 1977. Vol. 2. p. 1030-7.
20. **Miller R A, Pople H E & Myers J D.** *Internist-I*, an experimental computer-based diagnostic consultant for general internal medicine. *New Engl. J. Med.* 307:468-76, 1982.
21. **Feigenbaum E A.** Telephone communication. 28 November 1983.
22. **Davis R.** *Use of meta level knowledge in the construction and maintenance of large knowledge bases.* PhD dissertation, Stanford University, 1976. (Reprinted in: Davis R & Lenat D, eds. *Knowledge-based systems in artificial intelligence.* New York: McGraw-Hill, 1982. p. 229-490.)
23. **Duda R, Gaschnig J & Hart P.** Model design in the PROSPECTOR consultant system for mineral exploration. (Michie D, ed.) *Expert systems in the micro-electronic age. Proceedings of the 1979 AISB Summer School*, July 1979, Edinburgh, Scotland. Edinburgh: Edinburgh University Press, 1979. p. 153-67.
24. **NASA.** *Proceedings of the 1977 MACSYMA Users' Conference*, 27-29 July 1977, Berkeley, CA. Washington DC: National Aeronautics and Space Administration, 1977. 501 p.
25. **Allan R.** Robotics comes into its own with improved vision, sense perception, teaching languages. *Electron. Des.* 30(9):87-100, 1982.
26. **Colby K M.** Personal communication. 10 November 1983.
27. **Winograd T.** *Understanding natural language.* New York: Academic Press, 1972. 195 p.
28. **Hendrix G G & Sacerdoti E D.** Natural-language processing: the field in perspective. *BYTE* 6:304-52, 1981.
29. **Hendrix G G.** Human engineering for applied natural language processing. *IJCAI-77. Proceedings of the 5th International Joint Conference on Artificial Intelligence*, 22-25 August 1977, Cambridge, MA. Pittsburgh, PA: Carnegie-Mellon University, 1977. Vol. 1. p. 183-91.
30. **Hendrix G G, Sacerdoti E G, Sagalowicz D & Slocum J.** Developing a natural language interface to complex data. *ACM Trans. Database Syst.* 3(2):105-47, 1978.
31. **Alexander T.** Teaching computers the art of reason. *Fortune* 105(10):82-92, 1982.
32. **Brooks R A.** Model-based three dimensional interpretations of two dimensional images. (Drinan A, ed.) *IJCAI-81. Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, 24-28 August 1981, Vancouver, BC, Canada. Vancouver, BC: University of British Columbia, 1981. Vol. II. p. 619-24.
33. **Struble W T.** The year of the robot. *Technol. Rev.* 85(2):87, 1982.
34. **Garfield E.** From Vocoder to Vocalock—speech recognition machines still have a long way to go. *Essays of an information scientist.* Philadelphia: ISI Press, 1981. Vol. 4. p. 579-85. (Reprinted from: *Current Contents* (34):5-11, 25 August 1980.)
35. **Hilts P J.** Interview: John McCarthy. *OMNI* 5(7):101-6; 122-6, 1983.

36. Automatic programming. (Barr A & Feigenbaum E A, eds.) *The handbook of artificial intelligence.* Los Altos, CA: William Kaufmann, 1982. Vol. II. p. 295-379.
37. **Colby K M, Christinaz D, Parkison R C, Graham S & Karpf C.** A word-finding computer program with a dynamic lexical-semantic memory for patients with anomia using an intelligent speech prosthesis. *Brain Lang.* 14:272-81, 1981.
38. **Garfield E.** Quality control at ISI: a piece of your mind can help us in our quest for error-free bibliographic information. *Current Contents* (19):5-12, 9 May 1983.
39. --------------. Can citation indexing be automated? (Stevens M E, Giuliano V E & Heilprin L B, eds.) *Statistical association methods for mechanized documentation. Symposium proceedings,* Washington, 1964. Washington, DC: NBS, 15 December 1965. Misc. Publ. No. 269. p. 189-92.
40. --------------. Can criticism and documentation of research papers be automated? *Essays of an information scientist.* Philadelphia: ISI Press, 1977. Vol. 1. p. 83. (Reprinted from: *Current Contents* (9):4, 4 March 1970.)
41. --------------. Joshua Lederberg—multidisciplinarian *extraordinaire. Essays of an information scientist.* Philadelphia: ISI Press, 1977. Vol. 1. p. 81-2. (Reprinted from: *Current Contents* (8):4-5, 25 February 1970.)
42. **Lederberg J, Sutherland G L, Buchanan B G, Feigenbaum E A, Robertson A V, Duffield A M & Djerassi C.** Applications of artificial intelligence for chemical inference. I. The number of possible organic compounds. Acyclic structures containing C, H, O, and N. *J. Amer. Chem. Soc.* 91:2973-6, 1969.
43. **Garfield E.** ABCs of cluster mapping. Parts 1 & 2. Most active fields in the life and physical sciences in 1978. *Essays of an information scientist.* Philadelphia: ISI Press, 1981. Vol. 4. p. 634-49. (Reprinted from: *Current Contents* (40):5-12, 6 October 1980 and (41):5-12, 13 October 1980.)
44. **Miyamoto S, Miyake T & Nakayama K.** Generation of a pseudothesaurus for information retrieval based on cooccurrences and fuzzy set operations. *IEEE Trans. Syst. Man Cybern.* SMC-13(1):62-70, 1983.
45. **Marcus R S.** An experimental comparison of the effectiveness of computers and humans as search intermediaries. *J. Amer. Soc. Inform. Sci.* 34(6):381-404, 1983.
46. **Bledsoe W W.** Splitting and reduction heuristics in automatic theorem proving. *Artif. Intell.* 2:55-77, 1971.
47. **Bledsoe W W, Boyer R S & Henneman W H.** Computer proofs of limit theorems. *Artif. Intell.* 3:27-60, 1972.
48. **Bledsoe W W & Bruell P.** A man-machine theorem-proving system. *Artif. Intell.* 5:51-72, 1974.
49. **Bleich H L.** Computer evaluation of acid-base disorders. *J. Clin. Invest.* 48:1689-96, 1969.
50. **Garfield E.** The impact of hospital libraries on the quality and cost of health care delivery. *Current Contents* (8):5-10, 21 February 1983.
51. **Horowitz G L & Bleich H L.** Paperchase: a computer program to search the medical literature. *N. Engl. J. Med.* 305:924-30, 1981.
52. **Overall J E & Williams C M.** Conditional probability program for diagnosis of thyroid function. *J. Amer. Med. Assn.* 183:307-13, 1963.
53. **Gorry G A & Barnett G O.** Experience with a model of sequential diagnosis. *Comput. Biomed. Res.* 1:490-507, 1968.
54. **Warner H R, Toronto A F & Veasy L G.** Experience with Baye's theorem for computer diagnosis of congenital heart disease. *Ann. NY Acad. Sci.* 115:558-67, 1964.
55. **Chomsky N.** *Language and mind.* New York: Harcourt Brace Jovanovich, 1972. 194 p.
56. **Schank R C.** Conceptual dependency: a theory of natural language understanding. *Cognitive Psychol.* 3:552-631, 1972.
57. **Dickson D.** Britain rises to Japan's computer challenge. *Science* 220(4599):799-800, 1983.
58. **Glushkov V M.** USSR, computing in. (Belzer J, Holzman A G & Kent A, eds.) *Encyclopedia of computer science and technology.* New York: Dekker, 1979. Vol. 13. p. 498-507.
59. **Garfield E.** How can we prove the value of basic research? *Essays of an information scientist.* Philadelphia: ISI Press, 1981. Vol. 4. p. 285-9. (Reprinted from: *Current Contents* (40):5-9, 1 October 1979.)