

SRI International



ARTIFICIAL INTELLIGENCE

Technical Note 89

July 1973

By: Nils J. Nilsson
Staff Scientist

Artificial Intelligence Center

333 Ravenswood Ave. • Menlo Park, CA 94025
(415) 859-6200 • TWX: 910-373-2046 • Telex: 334 486



STANFORD RESEARCH INSTITUTE
Menlo Park California 94025 USA

Abstract

ARTIFICIAL INTELLIGENCE

by

Nils J. Nilsson

July 1973

This paper describes current progress in Artificial Intelligence (AI). Particular emphasis is given to describing AI as an independent field of study having both applied-technological and scientific-theoretical branches. The major problems faced by AI research are discussed, and progress toward their solutions is evaluated. It is argued that practical applications of AI research will most likely be concentrated in computer systems that understand natural language and perform advanced-automation tasks. It is also claimed that AI research will, at long last, make theoretical psychology a possibility.

ARTIFICIAL INTELLIGENCE

Nils J. NILSSON
Artificial Intelligence Center
Stanford Research Institute
Menlo Park, California 94025, USA

Technical Note 89

This paper is a survey of Artificial Intelligence (AI). It divides the field into four core topics (embodying the base for a science of intelligence) and eight applications topics (in which research has been contributing to core ideas). The paper discusses the history, the major landmarks, and some of the controversies in each of these twelve topics. Each topic is represented by a chart citing the major references. These references are contained in an extensive bibliography. The paper concludes with a discussion of some of the criticisms of AI and with some predictions about the course of future research.

1. INTRODUCTION

Can we ever hope to understand the nature of intelligence in the same sense that we understand, say, the nature of flight? Will our understanding of intelligence ever be sufficient to help us build working models—machines that think and perceive—in the same way that our understanding of aerodynamics helps us build airplanes? Intelligence seems so varied. We see it when a chemist discovers the structure of a complex molecule, when a computer plays chess, when a mathematician finds a proof, and even when a child walks home from school. Are there basic mechanisms or processes that are common to all of these activities and to all others commonly thought to require intelligence?

The field of Artificial Intelligence (AI) has as its main tenet that there are indeed common processes that underlie thinking and perceiving, and furthermore that those processes can be understood and studied scientifically. The processes themselves do not depend on whether the subject being thought about or perceived is chemistry, chess, mathematics, or childhood navigation. In addition, it is completely unimportant to the theory of AI who is doing the thinking or perceiving—man or computer. This is an implementational detail.

These are the emerging beliefs of a group of computer scientists claiming to be founding a new science of intelligence. While attempting to discover and understand the basic mechanisms of intelligence, these researchers have produced working models in the form of computer programs capable of some rather impressive feats: playing competent chess, engaging in limited dialogs with humans in English, proving reasonably difficult mathematical theorems in set theory, analysis, and topology, guessing (correctly) the structure of complex organic molecules from mass-spectrogram data, assembling mechanical equipment with a robot hand, and proving the correctness of small computer programs.

Whether the activities of these workers constitute a new scientific field or not, at the very least AI is a major campaign to produce some truly remarkable computer abilities. Like going to the moon or creating life, it is one of man's grandest enterprises. As with all grand enterprises, it will have profound influences on man's way of life and on the

way in which he views himself. In this paper, I will try to describe the AI campaign, how it seems to be organized into subcampaigns, who is doing what, some of the current internal controversies, and the main achievements. There is the usual word of caution: I've made some rather large simplifications in attempting to stand aside from the field and look at it with perspective. Not all workers would necessarily agree with what follows.

Before beginning we must discuss an important characteristic of AI as a field, namely, that it does not long retain within it any of its successful applications. Computer aides to mathematicians, such as differential equation solvers, that originated (at least partly) from AI research, ultimately become part of applied mathematics. A system, named DENDRAL, that hypothesizes chemical structures of organic molecules based on mass-spectrogram data is slowly escaping its AI birthplace and will likely become one of the standard tools of chemists. This phenomenon is well-recognized by AI researchers and has led one of them to state that AI is known as the "no-win" field. It exports all of its winning ideas.

On reflection, this is not surprising. When a field takes as its subject matter all of thinking, and then when particular brands of that thinking are applied to chemistry, mathematics, physics, or whatever, these applications become parts of chemistry, mathematics, physics, etc. When people think about chemistry, we call it part of chemistry—not an application of psychology. The more successful AI becomes, the more its applications will become part of the application area.

Destined apparently to lack an applied branch, is there a central core or basic science of AI that will continue to grow and contribute needed ideas to applications in other areas? I think the answer is yes. Just what form these central ideas will ultimately take is difficult to discern now. Will AI be something like biology—diverse but still united by the common structure of DNA? What will be the DNA of AI?

Or will the science of AI be more like the whole of science itself—united by little more than some vague general principles such as the scientific method? It is probably too early to tell. The

present central ideas seem more specific than does the scientific method but less concrete than DNA.

2. WHAT IS HAPPENING IN AI?

2.1 The structure of the field

As a tactic in attempting to discover the basic principles of intelligence, AI researchers have set themselves the preliminary goal of building computer programs that can perform various intellectual tasks that humans can perform. There are major projects currently under way whose goals are to understand natural language (both written and spoken), play master chess, prove non-trivial mathematical theorems, write computer programs, and so forth. These projects serve two purposes. First, they provide the appropriate settings in which the basic mechanisms of intelligence can be discovered and clarified. Second, they provide non-trivial opportunities for the application and testing of such mechanisms that are already known. I am calling these projects the first-level applications of AI.

I have grouped these first-level applications (somewhat arbitrarily) into eight topics shown spread along the periphery of Figure 1. These are the eight that I think have contributed the most to our basic understanding of intelligence. Each has strong ties to other (non-AI) fields, as well as to each other; the major external ties are indicated by arrows in Figure 1.

Basic mechanisms of intelligence and implementational techniques that are common to several applications, I call core topics. It seems to me that there are four major parts to this central core:

- * Techniques for modeling and representation of knowledge.
- * Techniques for common sense reasoning, deduction, and problem solving.
- * Techniques for heuristic search.
- * AI systems and languages.

These four parts are shown at the center of Figure 1. Again, we have indicated ties to other fields by arrows. It must be stressed that most AI research takes place in the first-level applications areas even though the primary goal may be to contribute to the more abstract core topics.

If an application is particularly successful, it might be noticed by specialists in the application area and developed by them as a useful and economically viable product. Such applications we might call second-level applications to distinguish them from the first-level applications projects undertaken by the AI researchers themselves. Thus, when AI researchers work on a project to develop a prototype system to understand speech, I call it a first-level application. If General Motors were to develop and install in their assembly plants a system to interpret television images of automobile parts on a conveyor belt, I would call it a second-level application. (We should humbly note that perhaps several second-level applications will emerge without benefit of obvious AI parentage. In fact, these may contribute mightily to AI science itself.)

Thus, even though I agree that AI is a field that cannot retain its applications, it is the second-level applications that it lacks. These belong to

the applications areas themselves. Until all of the principles of intelligence are uncovered, AI researchers will continue to search for them in various first-level applications areas.

Figure 1, then, divides work in AI into twelve major topics. I have attempted to show the major papers, projects, and results in each of these topics in Charts 1 through 12, each containing references to an extensive bibliography at the end of this paper. These charts help organize the literature as well as indicate something about the structure of work in the field. By arrows linking boxes within the charts we attempt to indicate how work has built on (or has been provoked by) previous work. The items in the bibliography are coded to indicate the subheading to which they belong. I think that the charts (taken as a whole) fairly represent the important work even though there may be many differences of opinion among workers about some of the entries* (and especially about how work has built on previous work).

Obviously, a short paper cannot be exhaustive. But in this section I will summarize what is going on in AI research by discussing the major accomplishments and status of research in each of the twelve sub-headings.

2.2 The core topics

Fundamentally, AI is the science of knowledge--how to represent knowledge and how to obtain and use knowledge. Our core topics deal with these fundamentals. The four topics are highly interdependent, and the reader should be warned that it is probably wrong to attempt to think of them separately even though we are forced to write about them separately.

2.2.1 Common-sense reasoning, deduction, and problem-solving (Chart 1)

By reasoning, etc., we mean the major processes involved in using knowledge: Using it to make inferences and predictions, to make plans, to answer questions, and to obtain additional knowledge. As a core topic, we are concerned mainly with reasoning about everyday, common domains (hence, common sense) because such reasoning is fundamental, and we want also to avoid the possible trap of developing techniques applicable only to some specialized domain. Nevertheless, contributions to our ideas about the use of knowledge have come from all of the applications areas.

There have been three major themes evident in this core topic. We might label these puzzle-solving, question-answering, and common-sense reasoning.

Puzzle-solving. Early work on reasoning concentrated on writing computer programs that could solve simple puzzles (tower of Hanoi, missionaries and cannibals, logic problems, etc.). The Logic Theorist and GPS (see Chart 1) are typical examples. From this work certain problem-solving concepts were developed and clarified in an uncluttered atmosphere. Among these were the concepts of heuristic search, problem spaces and states, operators (that transformed one problem state into another), goal and subgoal states, means-ends analysis, and reasoning backwards. The fact

* In particular, some might reasonably claim machine vision (or more generally, perception) and language understanding to be core topics.

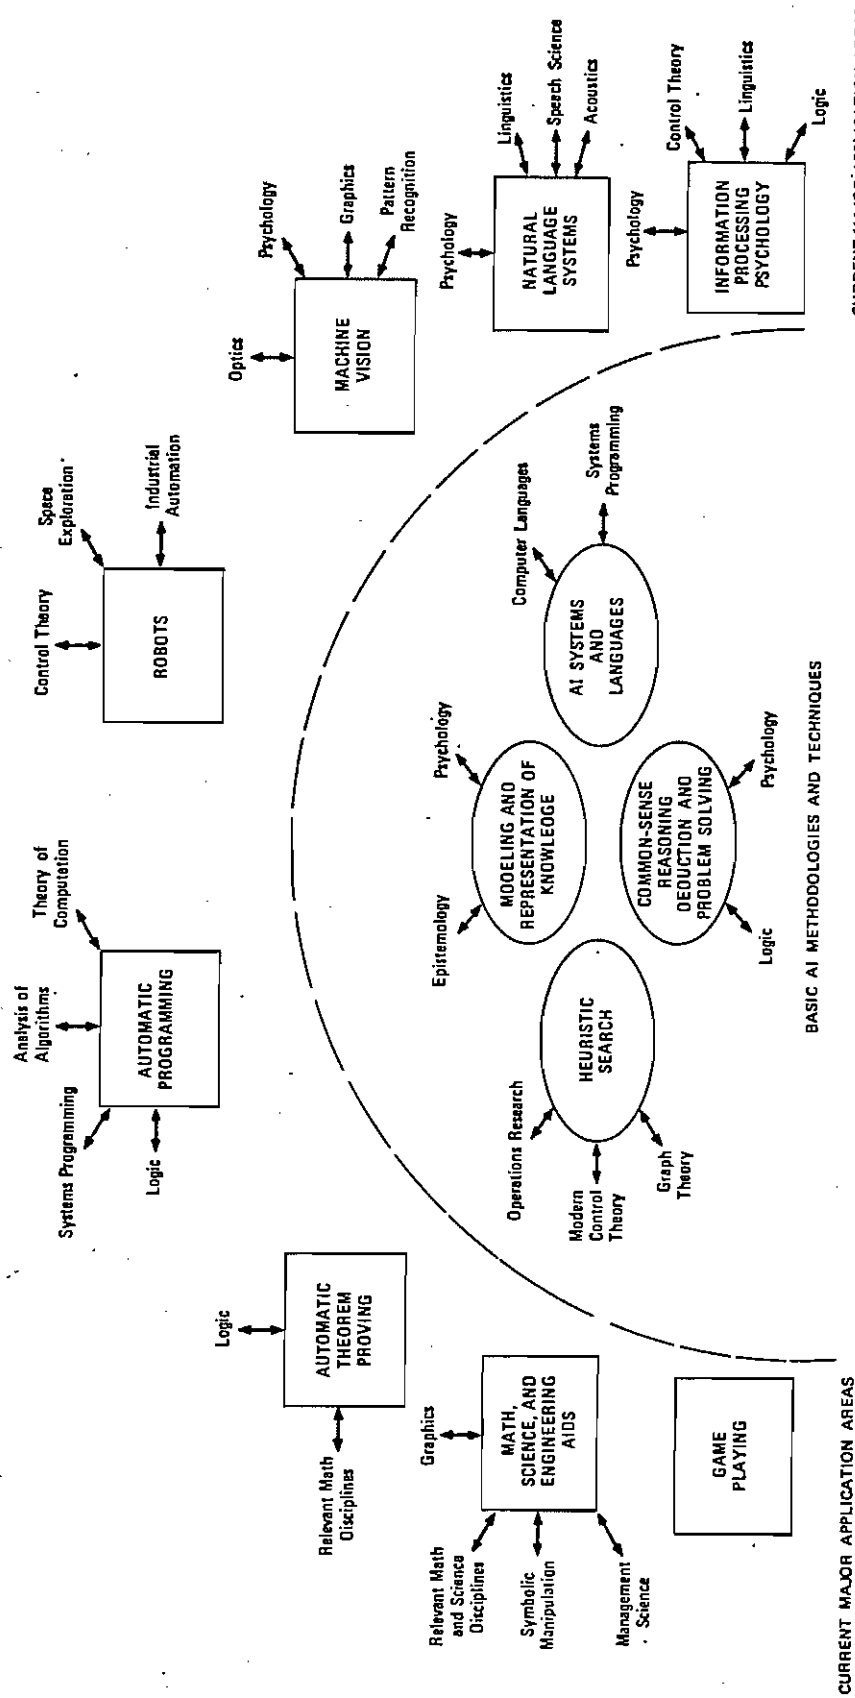


FIGURE 1 MAJOR SUB-PARTS OF AI SHOWING TIES TO OTHER FIELDS

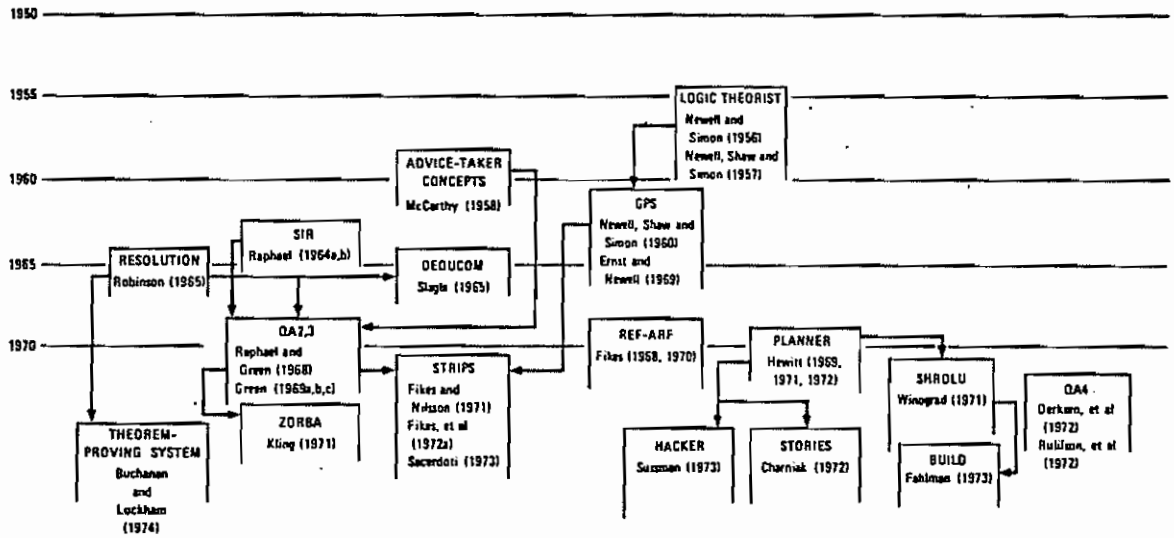


CHART 1: COMMON-SENSE REASONING, DEDUCTION AND PROBLEM SOLVING

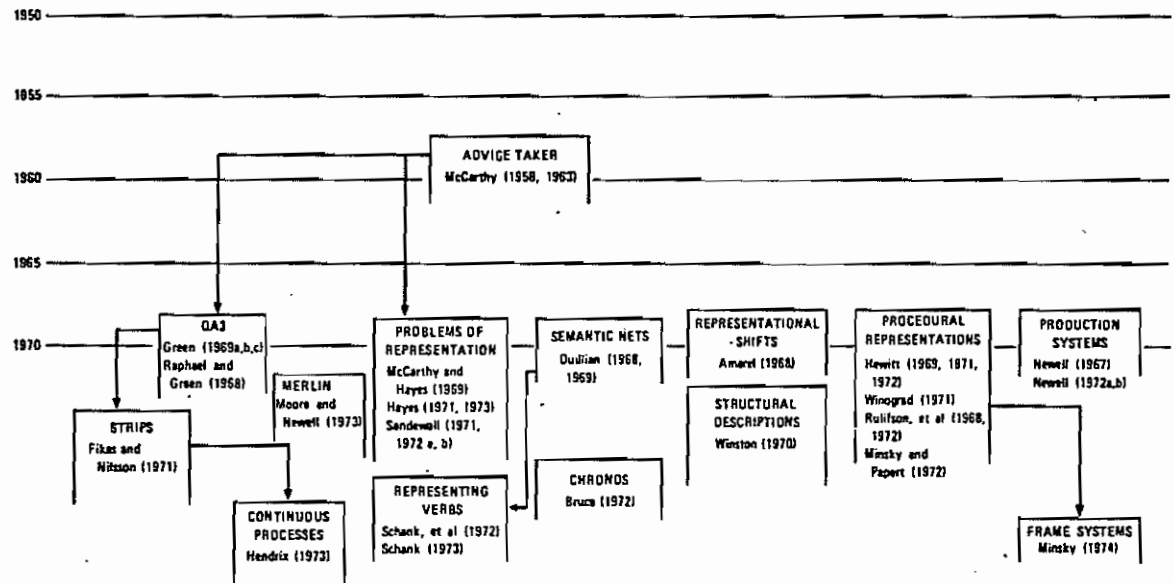


CHART 2: MODELING AND REPRESENTATION OF KNOWLEDGE

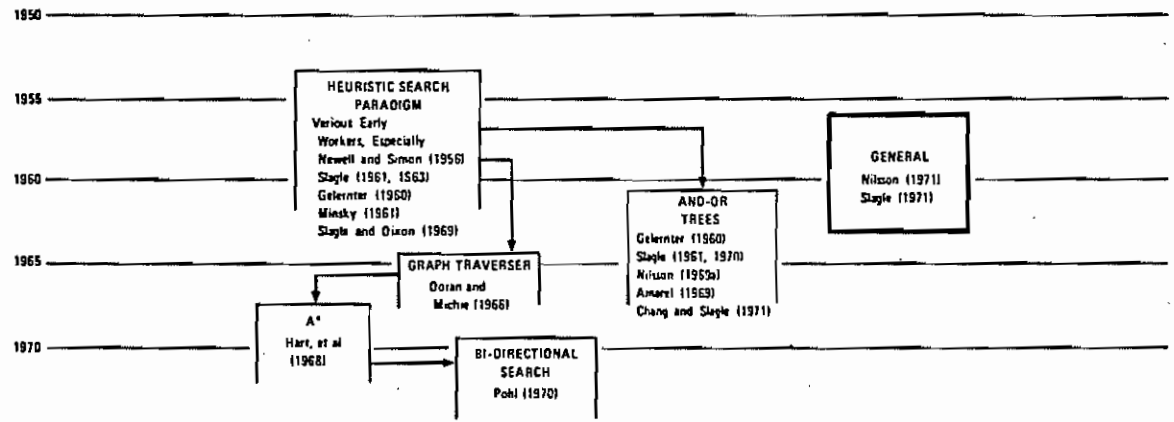


CHART 3: HEURISTIC SEARCH

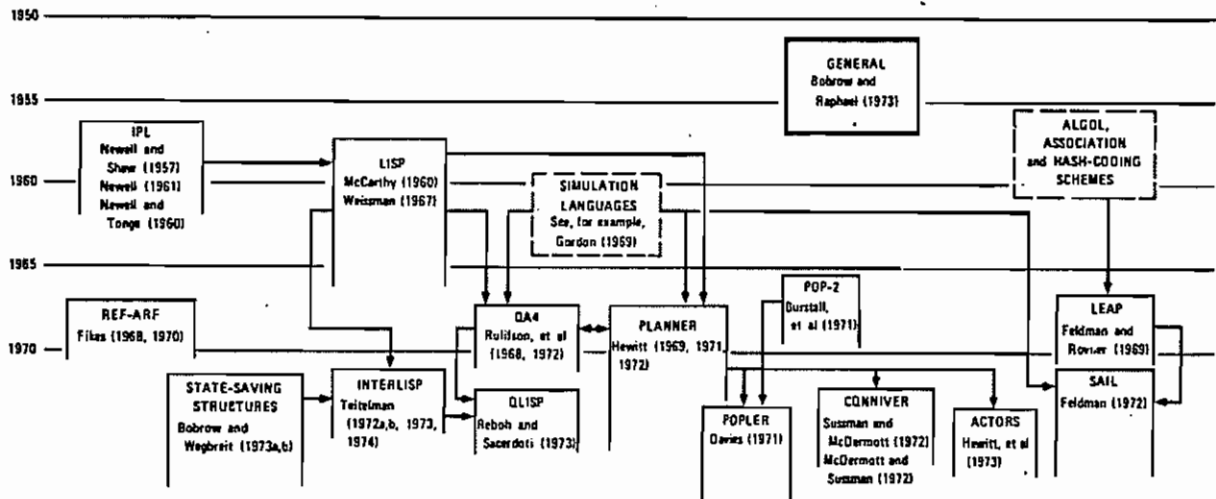


CHART 4: AI SYSTEMS AND LANGUAGES

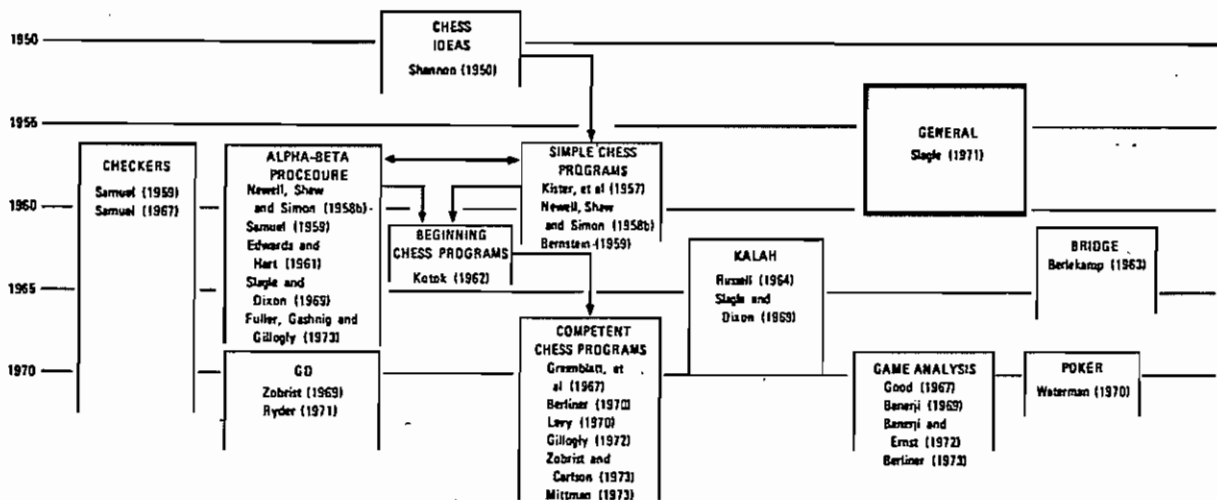


CHART 5: GAME PLAYING

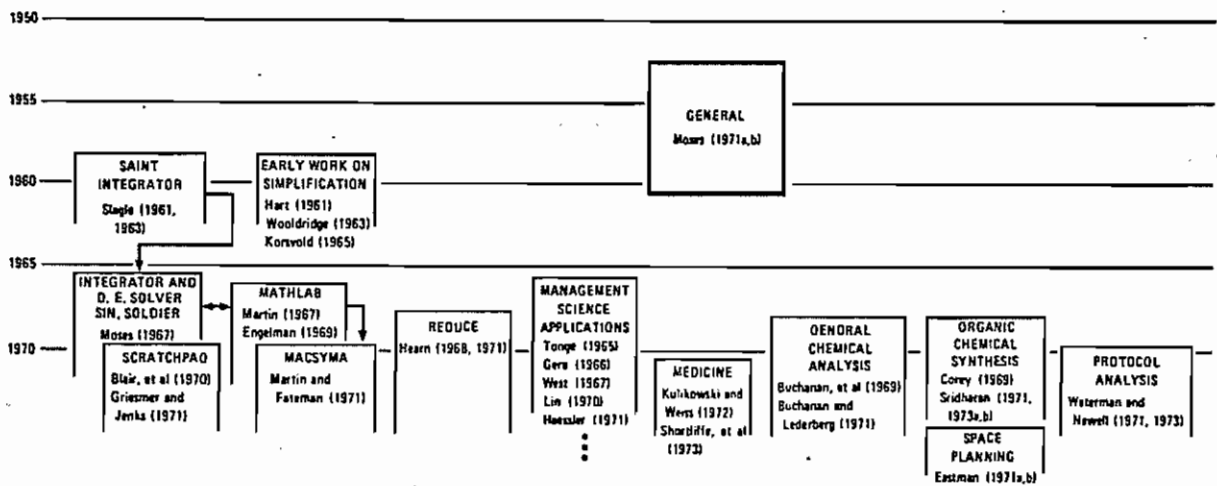


CHART 6: MATH, SCIENCE AND ENGINEERING AIDS

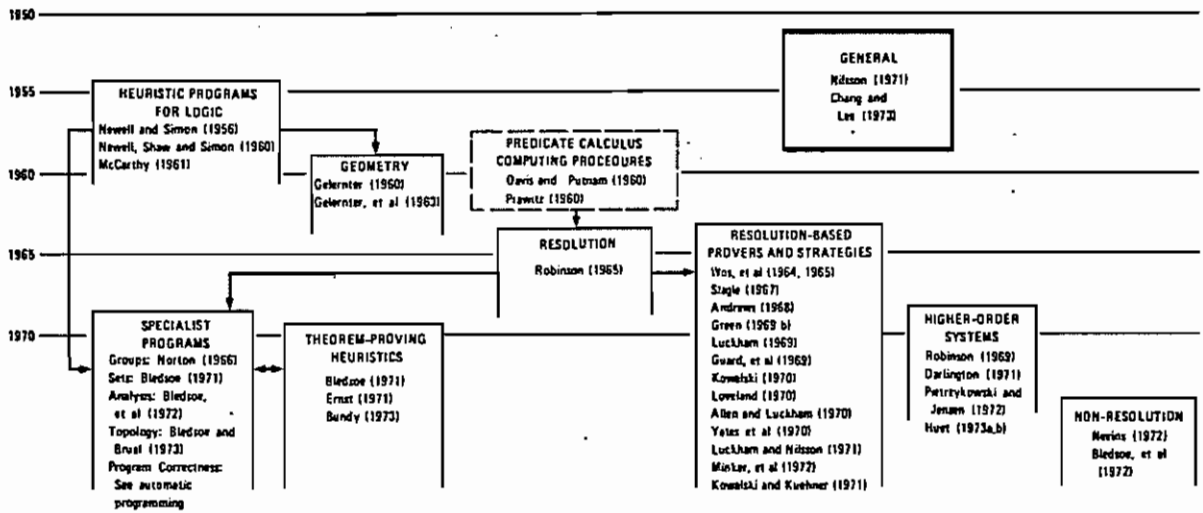


CHART 7: AUTOMATIC THEOREM PROVING

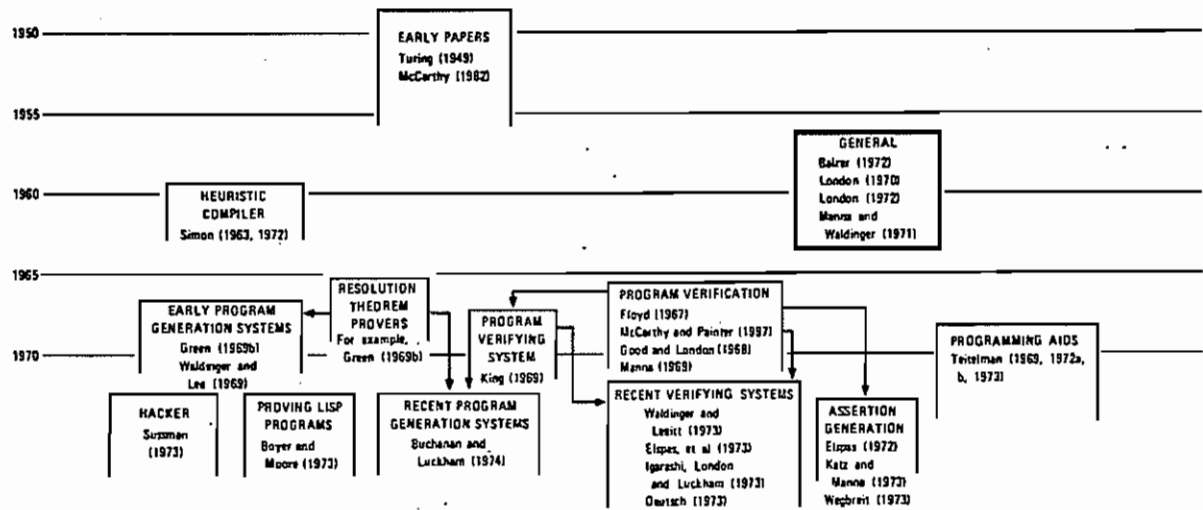


CHART 8: AUTOMATIC PROGRAMMING (Including Verification)

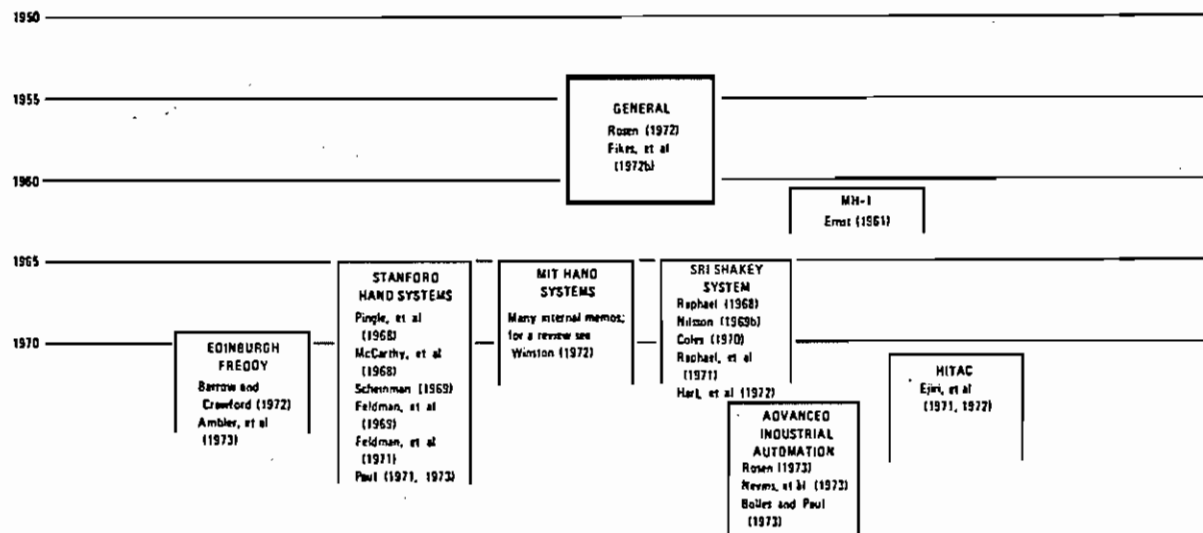


CHART 9: ROBOTS

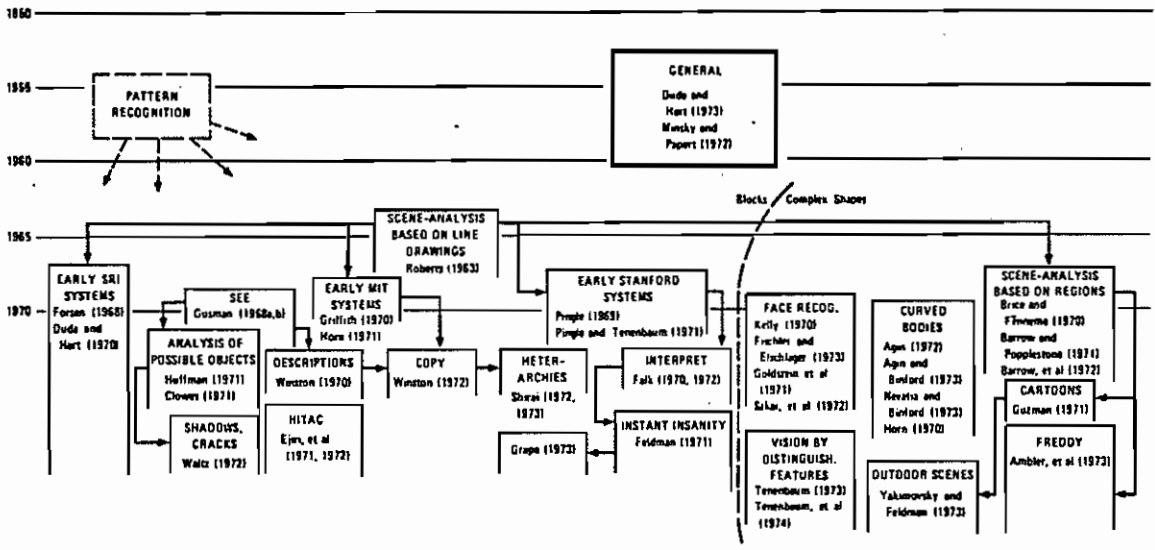


CHART 10: MACHINE VISION

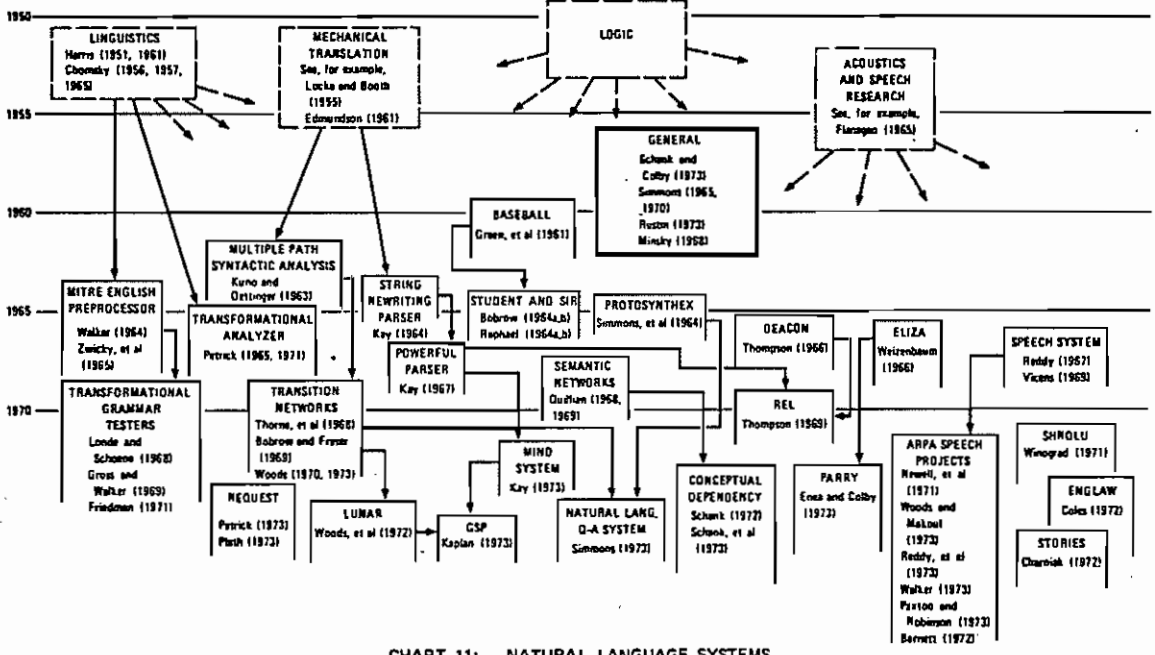


CHART 11: NATURAL LANGUAGE SYSTEMS

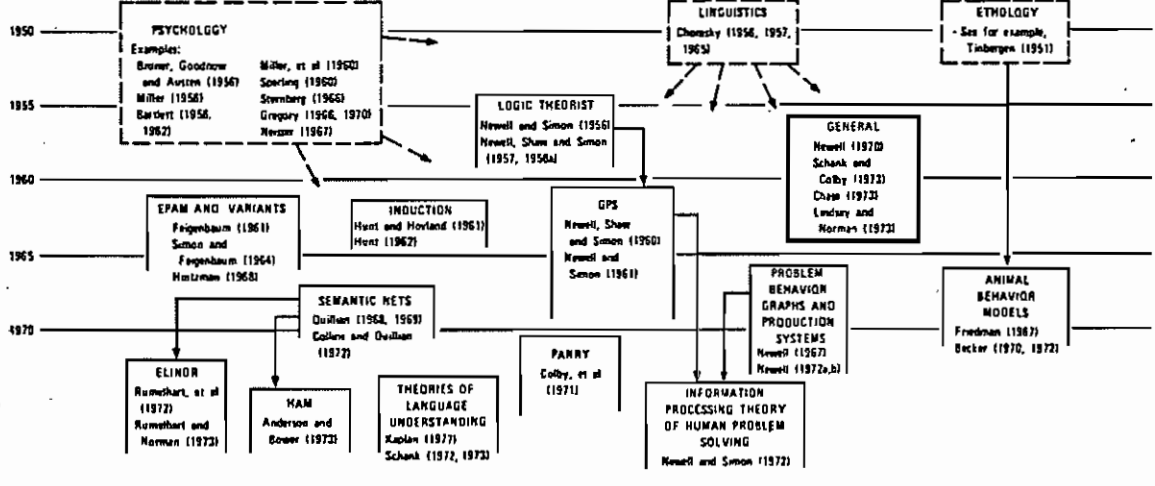


CHART 12: INFORMATION PROCESSING PSYCHOLOGY

that these useful ideas seem so familiar in AI research today testifies to the success of this early work. But the very cleanness of puzzles allowed researchers to avoid facing what has turned out to be the key problem, namely dealing with knowledge, huge amounts of knowledge, diverse, cluttered and inter-related.

Question-answering. As one step toward facing the problem of dealing with knowledge, several researchers concentrated on building inferential question-answering systems. (See, in particular, the references listed under SIR, QA2, and QA3 in Chart 1.) Such systems should be able to store a large number of facts and should be able to respond to reasonable questions whose answers could be deduced from these facts. Those systems required mechanisms for logical inference and led AI researchers into a romance with logic in general and with Robinson's resolution principle in particular. (See Chart 7.) This line of research clarified our concepts of applying inference techniques to common-sense knowledge and led to various useful schemes for associative retrieval of stored data. We also learned that for large question-answering systems the question of when to use inference methods was more important than the nature of the inference mechanism itself. Thus, we learned that we would need large amounts of secondary knowledge about how and when to use the primary knowledge of the domain.

Common-sense reasoning. In 1958, McCarthy proposed an ADVICE-TAKER that would be able to accept knowledge and use it to deduce answers to questions and to figure out simple plans for courses of action. One might ask such a system, for example, how to get to Timbuktu (a favorite example of McCarthy's). If the system knew about airline schedules, airports, how to get to airports, and other common (but immensely diverse) knowledge, it might answer thus: (1) go to your travel agent and find out about flights to Timbuktu, (2) using this information, select a flight and make a reservation, (3) drive to the airport at the appropriate time, (4) park your car, and (5) get on the appropriate airplane. Each of these steps, of course, could be expanded in detail.

Problems of this sort are clearly not so clean as puzzles; they demand the use of large amounts of knowledge; yet they have in common with puzzles the feature of planning a course of action to accomplish a goal.

Robotics research (see Chart 9) has probably contributed the most to our knowledge of how to generate plans based on large amounts of common-sense knowledge. Researchers at MIT, using an arm in a domain of simple blocks (called the BLOCKS world) and at SRI using a mobile robot in a domain of corridors and rooms, have developed various reasoning systems that can generate plans of action for a robot. Of these, we might mention in particular STRIPS, SHRDLU, and HACKER (see Chart 1).

There has been a lot of useful internal controversy about how to build reasoning systems and about the best directions for research. For a while, there was hope in some quarters that some universal system (based, for example, like QA3 on Robinson's resolution principle) could be used for all of the tasks we have mentioned so far: puzzle-solving,

question-answering, and common-sense reasoning. First attempts to build such universal systems were unsuccessful in the incorporation of the necessary domain-specific knowledge and techniques and, as far as I knew, there are at present no serious advocates of a simple universal system.

At the opposite extreme of this controversy, however, are the proponents of what I would call ad hocism. To them, following any systematic approach is anathema. Each task should simply be programmed on its own using whatever tricks might be needed. There is no doubt that this kind of opportunism is healthy for a growing field still in search of its general principles. Still, the following point must be made against rampant ad hocism: One part of developing a science is to discover those concepts that are important. We must try to produce intelligent behavior out of systems limited to various combinations of trial concepts. Our failures tell us where our present concepts are weak and give us hints about new ones that might be needed. If our trial concepts are always allowed the crutch of ad hocism, we do not learn enough about where the concepts are weak.

Another controversy concerns how much knowledge we ought to give our reasoning programs. At one extreme are researchers who insist that the program should be given only some basic premises from which it must derive any intermediate knowledge it needs to arrive at an answer. At the other (and impossible) extreme, programs would be provided explicitly with answers to all problems. There are some who feel that derivation of answers ultimately will play such a large role in intelligent systems that we may as well concentrate now on derivation techniques. To force derivation, they tend to work with knowledge-impooverished systems.

The consensus just now emerging from this controversy is that, because of combinatoric problems, an intelligent system probably will be able to make only reasonably direct derivations at any stage. Thus, to deal with a large domain, such a system must begin with a large skeletal network of basic knowledge about the domain and knowledge about how to use its knowledge. Any excursion from the known (explicitly represented) knowledge into the unknown (derived) can thus be well-guided (i.e., practical) even though the "volume" of the unknown part itself can be extremely large. It is senseless to insist that, to answer a single question, an intelligent system must repeat the tedious trial and error evolution of a large part of our cultural and scientific knowledge to say nothing of possibly having to repeat much of biological evolution itself. Even the "let's derive all" school would agree. What members of this school and some others did not realize was just how much knowledge would finally be needed by intelligent systems. Given this realization, the only possible course is to build "knowledge-based" programs.*

2.2.2 Modeling and representation of knowledge (Chart 2)

Our ideas about how to represent knowledge have come from several of the applications areas. (Quite

* Minsky (1974) guesses that a knowledge-based system reasoning about visual images (a system such as might be possessed by a typical human) "might need a few millions, but not billions, of structural units, interconnections, pointers."

obviously, every AI program uses some representational scheme. We cite in Chart 2 just a few of the important contributions.) Researchers in machine vision and perception and in natural language understanding were perhaps the first to realize how much knowledge would be needed by high performance programs. These two applications areas have thus probably contributed the most to our repertoire of representational techniques.

The systems mentioned in Chart 2 cover some of the major suggestions. For example:

Green (1969a,b,c): Statements in the first order predicate calculus.
Quillian (1968): Concept nodes in a graph structure linked by various relationships.
Schank et al. (1972): Canonical concept structures having "slots" for case information.
Hewitt (1969,71) and Winograd (1971): Pattern-invoked procedures plus assertions.
Rulifson et al. (1968): Pattern-invoked procedures plus special list structures such as n-tuples, bags and sets with property lists all organized in a discrimination net.
Newell (1967): Sets of productions organized as Markov tables.
Minsky (1974): Hierarchically organized structures called "frame systems." These have "free variables" (analogous to Schank's slots) that can be matched against constants occurring in the data to be analyzed.

For a period there was some controversy over whether knowledge should be represented assertationally or procedurally. (As an extreme case, a spiral, say, can be represented assertationally by a list of the points in the plane through which it passes, or it can be represented procedurally by a program that draws it.) Something of a cult was made of the "procedural embedding" of knowledge, but this controversy seems to be settling down now to an acceptance of the value of a combination of assertional and procedural knowledge.

Another concern, having antecedents in logic, is how to represent certain "modal" concepts involving time, necessity, possibility, and so forth. McCarthy & Hayes (1969) have analyzed some of the difficulties in formalizing these concepts; meanwhile, Hendrix (1973) and Bruce (1972) have developed systems that begin to deal with some of them.

McCarthy and Hayes (1969) also discuss two fundamental problems concerning representation and reasoning. One is called the frame problem, and it concerns certain difficulties of model maintenance. If we have a representation of the world at a certain instant (based on observations and a priori knowledge), how should we represent and use "laws of physics" to update the model so that it represents the world (reasonably accurately) at some future instant? If a robot removes a book from a shelf, can we assume that a door across the room remains open without having to derive this fact or observe it again? There are several ways of dealing with this problem, e.g., Green (1969), Fikes and Nilsson (1971), Sandowall (1972), and Hewitt (1969). These are nicely discussed by Hayes (1973).

Another problem is the qualification problem. If a system uses its representation to "prove," say, that a certain plan will achieve a desired goal

(the goal of being at the airport), how are we to deal with certain difficulties arising when new information is received prior to executing the plan. Suppose, for example, someone tells us that our automobile is out of gasoline so that now our plan (that called for driving to the airport) will not work. We had proved that it would, and now new information has rendered the proof invalid even though all of the information on which the original proof was based is still present. Hayes (1973) discusses this violation of the "extension property" and shows the close connection between the qualification problem and the frame problem. System builders [e.g., Hewitt (1969) and Rulifson et al. (1972)] have invented certain constructs that apparently get around these difficulties, although in a way that is somewhat unsatisfactory to logicians.

We are still quite a way, it seems, from having a sound theoretical basis for knowledge representation. It is my view that the necessity of developing large and complex reasoning systems will produce the new concepts out of which the needed theories will be constructed.

2.2.3 Heuristic search (Chart 3)

One of the first results of early AI research was the development of a point of view toward problem-solving sometimes called "the heuristic search paradigm." There are two closely related versions of this paradigm. In one, a "problem" is transformed into the canonical problem of finding a path through a "space" of problem states from the initial state to a goal (i.e., solution) state. In the other, a problem is "reduced" to various subproblems that are also reduced in turn (and so on) until the ultimately resulting subproblems have trivial or known solutions. Each version is merely a slightly different way of thinking about basically the same problem-solving process. In each, the process involves generating alternative paths toward solutions, setting up certain key milestone states (or subproblems), and managing search resources wisely to find acceptable solutions.

The word "heuristic" is used because these techniques emphasize the use of special knowledge from the problem domain that "aids in discovering a solution" by drastically reducing the amount of search that would otherwise have to be employed. Often this knowledge takes the form of "rules-of-thumb" that help to limit or direct the search. Sometimes they are constraining relations that can be employed to limit the search needed. [A good example of the use of constraints is the work of Waltz (1972).]

I have already referred to some of the heuristic search paradigm ideas (subgoals, reasoning backwards, and so on) as being basic to common-sense reasoning, deduction, and problem solving (Chart 1). Here (in Chart 3), we want to cite mainly those aspects of heuristic search dealing with the search process itself. Once a problem is represented as a search problem, how can a solution be found efficiently?

The searching occurs in one of two graph structures, ordinary graphs (or trees), and AND-OR graphs (or trees), depending on whether the problem is viewed as one of finding a path to a goal state or one of reducing problems to subproblems, respectively. The search techniques that have been developed (by workers in AI, control theory, and operations

research) are now commonly used in many AI programs and in many of their applications. Most of these techniques make use of heuristically-based evaluation functions that rank-order the unexplored nodes in the graph and thus indicate where search can most efficiently proceed. Furthermore, there are some theorems [Hart et al. (1968)] stating conditions under which these search methods are guaranteed to find optimal paths. The problem of efficiently searching a graph has essentially been solved and thus no longer occupies AI researchers. This one core area, at least, seems to be well under control.

2.2.4 AI systems and languages (Chart 4)

The programming languages developed and used by AI researchers are included among the core topics because they embody the most useful of the core ideas already discussed. Early AI researchers saw the need for programs that could store, access, and manipulate lists of symbolic information. The means for achieving these and other operations were built into various list processing languages, primarily IPL-V and LISP.

After some years of research using these languages, it became apparent that AI systems had a common, recurring need for operations such as search, expression-retrieval, and pattern-matching. The next step was to build these operations into the languages themselves. Thus, in the late 1960s, another generation of AI languages emerged, languages such as QA4 and PLANNER.

Edward Feigenbaum once characterized progress in AI research as progress along the "what-to-how" spectrum of computer languages. At the "how" end of this spectrum are the machine languages used by programmers who must give the most detailed instructions to the computer. As one progresses toward the "what" end, the programmer leaves more and more of the details of how operations are to be carried out to the language and can be more and more concerned only with what is to be done. AI languages are now moderately far along toward the "what" end, and the proper goal of AI research (according to this view) is to create languages even closer to the "what" end. It may well be that, ultimately, the field of AI will in large part be concerned with the development of superpowerful computing languages. In this light, the best way to measure AI progress is to look at the AI languages.

We do not have space here to trace the development of AI languages nor to describe the special features that they make available to AI researchers. Fortunately, there is an excellent tutorial paper by Bobrow and Raphael (1973) that gives a very clear account of the new languages.

Currently, a large part of AI research is being conducted by experimenting with systems written in the new languages. The languages provide especially powerful mechanisms for representing the extensive knowledge needed by present programs. Furthermore, this knowledge can now be easily added incrementally as the program evolves under the tutelage of human experts in the domain. Winograd's (1971) natural language understanding system and Waldinger and Levitt's (1974) system for proving assertions about programs are good examples of how the power of these languages is being used.

It would not be unreasonable to expect that current and future experimentation will lead to the crystallization of additional concepts [such as, perhaps, Minsky's (1974) Frame Systems] that will be incorporated in a new round of AI languages, possibly in the late 1970s.

2.3 First-level applications topics

2.3.1 Game playing (Chart 5)

Programs have been written that can play several games that humans find difficult. As the most famous example, we might mention the chess playing program, MAG-HACK, of Greenblatt et al. (1967). A version of this program achieved a United States Chess Federation rating of 1720 in one tournament. Samuel's programs for checkers have beaten experts in the game. Several other programs are mentioned in the chart.

Levy (1970) described a program written by Atkins, Slate, and Gorland at Northwestern University and said that he thought it was stronger than Greenblatt's. He estimated its rating at about 1750, which would make it, he claims, the 500th best player in Britain.

Computer chess tournaments are now held routinely. Results of these and other news about computer chess have been rather extensively reported in the SIGART Newsletter since 1972.

Most game playing programs still use rather straightforward tree-searching ideas and are weak in their use of high-level strategic concepts. It is generally agreed that advances in the use of strategy and in end-game play are necessary before chess programs can become substantially better, and they must become substantially better before they can beat human champions. (World Champion Bobby Fischer is rated at about 2810.) Levy (1970) is rather pessimistic about the rate of future progress in chess and has made a £750 bet with Professors McCarthy, Papert, and Michie that a program cannot beat him in a match by August 1978. (Levy's rating in 1970 was 2380.)

2.3.2 Math, science, and engineering aids (Chart 6)

The chart lists just a few examples of AI techniques that have been applied in systems that help human professionals. The early AI work on symbolic integration, together with the work on algebraic simplification, contributed to a number of systems for symbolic mathematical computations. Moses (1971b) presents a good review. Systems presently exist that can solve symbolically an equation like $y^{2x} - 3y^x + 2 = 0$ (for x), and that can integrate symbolically an expression like $\int (x + e^x)^2 dx$. Such systems are quite usefully employed in physics research, for example, in which expressions arise having hundreds of terms.

Another quite successful application is the DENDRAL program that hypothesizes chemical structures from a combination of mass spectrogram and nuclear magnetic resonance data. The system is presented with this data from a sample of a known chemical compound (that is, its chemical formula is known). It uses several levels of knowledge about chemical structures and how they break up in mass spectroscopy to infer the structure of the compound. It can deal with

a large number of organic compounds including complex amines and estrogenic steroids. Its performance on the steroids often exceeds the best human performance.

The DENDRAL project typifies a style of AI system building that has been quite successfully applied to chemistry and some other domains. This design style involves intensive interaction between AI scientists and applications area scientists. The latter are queried in the minutest detail to extract from them rules and other knowledge that are operationally useful in the domain. These are then coded into the system by the AI scientists and tests are run to judge their effectiveness. The process is long and involves several iterations. The applications scientists are often confronted with apparent contradictions between how they say they make decisions and how they actually make decisions. Few of them have any really global or completely accurate theory of how they apply their knowledge. Furthermore, this knowledge is often informal and heuristic. As a result, the emerging system is a collection of "mini-theories" and special rules of only local effectiveness. To use this design strategy, the system must be one that can deal with many, and sometimes conflicting, mini-theories. It must also be a system to which new knowledge can gradually be added and old knowledge modified.

After several months or years of this sort of gradual shaping of the system, it comes to simulate the performance of the human experts whose knowledge it has gained. This general strategy is beginning to be employed extensively in AI applications. [For example, see also Shortliffe et al. (1973).]

2.3.3 Automatic theorem proving (Chart 7)

There are three major themes evident in attempts to get computer programs to prove theorems in mathematics and logic. First, early work by AI researchers produced heuristic programs that could prove simple theorems in propositional logic and high-school level theorems in plane geometry. These programs used (but mainly helped to refine) concepts like reasoning backwards, means-ends analysis, use of subgoals, and the use of a model to eliminate futile search paths. The fact that logicians had already developed powerful procedures that effectively eliminated propositional logic as a domain requiring heuristic problem-solving techniques does not detract from the value of this early work.

Logicians were also developing techniques for proving theorems in the first order predicate calculus. J. A. Robinson (1965) synthesized some of this work into a procedure for using a single rule of inference, resolution, that could easily be mechanized in computer programs. Building resolution-based provers quickly became a second theme in automatic theorem proving, while other approaches languished. Resolution had a great influence on other application areas as well (Charts 1 and 8). Performance of the resolution systems reached impressive, if not superhuman, levels. Programs were written that could prove reasonably complex, sometimes novel, theorems in certain domains of mathematics. The best performance, however, was achieved by man-machine systems in which a skilled human provided strategic guidance leaving the system to verify lemmas and to fill in short chains of deduction. [See especially Guard et al. (1969) and Allen and Luckham (1970)]. The latter system has been used to obtain proofs of new mathematical

results announced without proof in the Notices of the American Mathematical Society.]

Various strategies were developed to improve the efficiency of the resolution provers. These strategies were mainly based on the form or syntax of the expressions to be proved and not on any special knowledge or semantics of the domain. In automatic theorem proving, just as in other applications areas, semantic knowledge was needed to improve performance beyond the plateau reached by the late 1960s.

The work of Bledsoe and his students is typical of the third and latest theme in automatic theorem proving. Although they emphasize the importance of man-machine systems, their programs themselves have become knowledge-based specialists in certain mathematical domains. The use of semantic knowledge in theorem-proving systems has also renewed interest in heuristics for subgoaling, and so forth. The programs of this group are capable of proving some rather impressive theorems, and it can be expected that the present man-machine systems will produce ever more competent and more completely automatic offspring.

2.3.4 Automatic programming (Chart 8)

Work in automatic programming has two closely inter-related goals. One is to be able to prove that a given program acts in a given way; the other is to synthesize a program that (provably) will act in a given way. The first might be called program verification and the second program generation. Work on one goal usually contributes to progress toward the other; hence, we combine them in our discussion.

Most of the work on program verification is based on a technique proposed by Floyd (1967). This technique, inspired by Turing (1949), involves associating assertions with various points in the flow chart of a program and then proving these assertions. Originally, the assertions had to be provided by a human, but some recent work has been devoted to generating the assertions automatically. Once proposed, one can attempt to have the assertions proved either by a human or by a machine. The latter course involves a close link between this field and that of automatic theorem proving.

A recent system developed at the Stanford Research Institute [Elspas et al. (1973)] is typical of one in which the assertions are both produced [Elspas (1972)] and proved [Waldinger and Levitt (1973)] automatically. This system has been used to verify several programs including a real-number division algorithm and some sort programs. It has also proved theorems about a pattern matcher and a version of Robinson's (1965) unification algorithm. It is a good example of a modern AI program in that it makes effective use of a large amount of domain-specific knowledge.

The closely related work on program generation has succeeded in producing some simple programs. Typical of this work is the system of Buchanan and Luckham (1974). Broadly viewed, the problem of constructing a computer program includes the problem of constructing a plan, say, for a robot, and thus there are close links between work in automatic programming, robotics, and common-sense reasoning and deduction.

Sussman's (1963) HACKER is another system that writes simple programs for a limited domain (the BLOCKS world). Sussman's goal for HACKER is for it to simulate his own programming style. An important feature of HACKER is its strategy of attempting first to

write a simple "let's-hope-that-this-will-do" program, and then debugging it until it does succeed at its task. To employ this strategy, HACKER uses a great deal of knowledge about likely classes of program bugs and how to fix them.

Again, some of the most successful work has been in connection with man-machine systems. We include in this category certain aids to human programmers such as those found in the INTERLISP system [Teitelman (1972a, b, 1973)]. In fact, any techniques that help make the production of programs more efficient might be called part of automatic programming. Balzer (1972) provides a good summary of this broad view of the field.

2.3.5 Robots (Chart 9)

Every now and then, man gathers up whatever technology happens to be around and attempts to build robots. During the late 1960s, research on robots provided a central focus for integrating much of the AI technology. To build an intelligent robot is to build a model of man. Such a robot should have general reasoning ability, locomotive and manipulative skills, perceptual (especially visual) abilities, and facility with natural language. Thus, robot research is closely linked with several other applications areas. In fact, most of the research on machine vision (Chart 10) was, and is, being performed in connection with robot projects.

Our problem-solving and representational techniques are probably already adequate to allow useful general purpose robot applications; however, such robots would be perceptually impoverished until we develop much more powerful visual abilities. Robotics is a particularly good domain in which to pursue the necessary vision research.

The robot research of the late 1960s produced systems capable of forming and then intelligently executing plans of action based on an internal model of the world. The Edinburgh, Stanford, HITAC, and MIT systems consisted of manipulator arms and TV cameras or other visual input devices. These became capable of building structures out of simple blocks. In one case (Stanford), the system could assemble an automobile water pump. The Stanford Research Institute system consisted of a mobile cart and TV camera (but no arm). It could form and execute plans for navigating through a simple environment of rooms, doorways, and large blocks, and its visual system could recognize and locate doorways, floor-wall boundaries, and the large blocks. The system had sophisticated techniques to allow it to recover from errors and unforeseen circumstances, and it could store (learn) generalized versions of the plans it produced for future use.

Since practical applications of general purpose robot systems seem more remote than they do in other applications areas, the increasingly pragmatic research climate of the early 1970s has seen a lessening of activity in general robotics research. In the meantime, various projects with the practical goal of advancing industrial automation have begun to apply some of the already-developed manipulative and visual skills to factory assembly and inspection problems. It seems reasonable to predict that man's historic fascination with robots, coupled with a new round of advances in vision and reasoning abilities, will

lead to a resurgence of interest in general robot systems, perhaps during the late 1970s.

2.3.6 Machine vision (Chart 10)

The ability to interpret visual images of the world is adequate enough even in some insects to guide many complex behavior patterns. Yet the analysis of everyday visual scenes by machine still remains a largely unconquered challenge to AI researchers. Early work concentrated almost exclusively on designing systems that could classify two-dimensional images into a small number of categories--alpha-numeric character recognition, for example. In fact, much of the AI work during the 1950s was concerned with pattern recognition. Researchers, such as Frank Rosenblatt and Oliver Selfridge, were influential in shaping this early period. Pattern classification (or recognition) continues as a separate active research interest, but since about 1965, AI interest in vision has centered on the more difficult problem of interpreting and describing complex three-dimensional scenes. Both aspects, classification and description, are thoroughly and clearly treated in an excellent textbook by Duda and Hart (1973).

Much of the scene analysis work can be traced to Robert's (1963) influential thesis. It established a trend of analyzing scenes composed of prismatic solids (the so-called "blocks world"). Working with these (sometimes complex) scenes composed of simple objects helped to establish a wide range of techniques for converting raw video images into symbolic descriptions based on concepts such as lines, regions, and simple shapes. The MIT "COPY" system, for example, can use a visual input device to look at a scene consisting of a structure of blocks. The system can analyze the scene to form a representation of how the blocks are arranged. This representation can then later be used (with the robot arm system) to reproduce this exact block structure from disarranged blocks.

Some successful excursions outside the blocks world have been made. (See the entries to the right of the dashed line in Chart 10). Indeed, many researchers contend that continuing to work with blocks has actually hindered research progress in machine vision because it has allowed workers to avoid facing certain key problems associated with domain semantics, distinguishing features of complex objects, and new representational schemes. In any case, working with more complex scenes is now well established. The spirit of recent work is well described in a note by Tenenbaum (1973). Again, knowledge about the domain is crucial to scene interpretation!

2.3.7 Natural language systems (Chart 11)

Processing written and spoken forms of natural languages has been the subject of so much research that the task of surveying progress in this field alone should humble anyone. How, then, are we to say anything very helpful in just a few sentences. The interested reader should see the GENERAL volumes cited in Chart 11.

Perhaps a description of the power of some representative systems will suffice. Winograd (1971)

implemented a system that carried on "blocks-world" written dialogs like the following:

Man (typing into a terminal): Which cube is sitting on the table?
System (generating text on a display): The large green one which supports the red pyramid.
M: Is there a large block behind a pyramid?
S: Yes, three of them: A large red one, a large green cube, and the blue one.
M: Put a small one onto the green cube which supports a pyramid.
S: OK
M: Put the littlest pyramid on top of it.
S: OK

The system demonstrates its understanding of the last two commands by having a simulated robot arm carry out appropriate actions in a simulated blocks world.

The work of Schank (1972) typifies a rather successful trend in natural language understanding. Many of the recent systems, in one way or another, attempt to match a section of input text or utterance against semantically likely stored structures (that are more or less complex.) These structures are themselves schemas or scenario families having variables that are bound to constants in the input during matching. The instantiated scenarios serve as a sort of deep structure that represent the meaning of the utterance. [See also Minsky (1974).]

The goals of a coordinated scientific effort to produce systems to understand limited utterances of continuous speech are clearly outlined in a plan by Newell et al. (1973). If the goals are met, by 1976 a prototype system should be able (in the context of a limited domain of discourse) to understand (in a few times 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 1000-word vocabulary. These projects bring together workers in acoustics and speech research as well as in AI. The projects seem to be more or less on schedule and will probably achieve creditable performance by 1976. (In the spirit of the vagueness of the phrase "a few times real time," the projects ought to achieve the 1976 goals at least sometime in the late 1970s.)

In my opinion, the work in natural language understanding is extremely important both for its obvious applications and for its future potential contributions to the core topics of AI. It is the prime example of a field in which reasonable performance could not be achieved by knowledge-impooverished systems. We now know that understanders need large amounts of knowledge; the challenge is to attempt to build some really large systems that have the adequate knowledge and to learn, by our mistakes, the organizational principles needed to keep these large systems from becoming unwieldy.

2.3.8 Information processing psychology (Chart 12)

Computer science in general and AI in particular have had a tremendous impact on psychology. They have and will continue to provide the concepts and the very vocabulary out of which to construct the most useful theories of human behavior. In my opinion the reason that, say, prior to 1955, there were, in fact, no adequate theories of human behavior, perception, and cognition is because the concepts out of which to

construct these theories had not yet been formulated. Before we have the concepts (and they are now gradually accumulating) it is as impossible to understand human thought as it was impossible to understand bat navigation, say, before we had the concept of sonar. Man understands the world by constructing models, and his models are often based on concepts drawn from his technological inventions. We may not understand man immediately after building the first robot, but we certainly won't understand him before! (We note in passing that knowledge about the structure and function of the neuron—or any other basic component of the brain—is irrelevant to the kind of understanding of intelligence that we are seeking. So long as those components can perform some very simple logical operations, then it doesn't really matter whether they are neurons, relays, vacuum-tubes, transistors, or whatever.)

An excellent short account of the relationship between AI and psychology has been written by Newell (1970). While he, perhaps prudently, adopts a somewhat less extreme position than mine about the dependence of psychology on AI, he nevertheless shows how thoroughly information processing ideas have penetrated psychological theory.

Most of the information-processing-based psychology to date has been devoted to explaining either memory (e.g., EPAM and HAM in Chart 12), perception [e.g., Sternberg (1966)], or problem solving [e.g., Newell and Simon (1972)]. Probably the most complete attempt at understanding human problem-solving ability is the last-mentioned work of Newell and Simon. This volume proposes an information processing theory of problem-solving based on the results of many years of research in psychology and AI.

Animal behavior, while long the special interest of experimental psychologists, has had little information-processing-based theoretical attention. Some models inspired by ethologists have been proposed by Friedman (1967). I think that the production system model advanced to explain certain human problem solving behavior by Newell (1967) and colleagues might be a starting point for an extensive theory of animal behavior. Newell, himself, notes that these production systems can be viewed as generalizations of stimulus-response systems. [Incidentally, the entire repertoire of what was called "intermediate-level actions" of the Stanford Research Institute robot system (Raphael et al. 1971) was independently programmed in almost exactly this production formalism. Production systems have been used in other AI programs as well.] Newell and Simon (1972, p. 803) have also stated that they "have a strong premonition that the actual organization of human problem solving programs closely resembles the production system organization . . ." It would seem profitable then to attempt to trace the evolutionary development of this hypothesized production system organization down through some of the higher animals at least.

3. CONCLUSIONS

In summary, we see that the AI campaign is being waged on several different fronts, and that the victories, as well as the setbacks, contribute to a growing common core of ideas that aspires to be a science of intelligence. Against this background,

it is worth mentioning some of the popular criticisms of AI:

(1) AI hasn't really done anything yet. There are a few "toy" programs that play middling chess and solve simple puzzles like "missionaries and cannibals," but the actual accomplishments of AI measured against its promises are disappointing. [See, for example, Dreyfus (1965, 1972).] [My comment about this kind of criticism is that its authors haven't really looked at AI research past about 1960.]

(2) Not only has AI not achieved anything, but its goals are actually impossible. Thus, AI is something like alchemy. It is impossible in principle to program into computers such necessities of intelligence as "fringe consciousness" and "perspicuous grouping." [Again, see Dreyfus (1965, 1972).] [This kind of criticism is actually rather brave in view of the fate of many previous impossibility predictions. This attack simply looks like a poor bet to me.]

(3) The subject matter of AI, namely intelligence, is too broad. It's like claiming science is a field. [This criticism may have some merit.]

(4) Everything happening in AI could just as well happen in other parts of computer science, control engineering, and psychology. There is really no need for this AI "bridge" between already established disciplines. [See Lighthill (1973).] [This kind of criticism caused quite a stir in Great Britain recently. I think I have shown that the so-called bridge has quite a bit of internal structure and is contributing a heavy traffic of ideas into its terminal.]

(5) AI is impossible because it is attempting to reduce (to understanding) something fundamentally "irreducible." Furthermore, this very attempt is profane; there are certain awesome mysteries in life that best remain mysterious. [See Roszak (1972).] [My prejudice about this view is that, at best, it is, of course, nonsense. A blind refusal even to attempt to understand is patently dangerous. By all means, let us not foreclose a "rhapsodic understanding" of these mysteries, but let us also really understand them.]

(6) AI is too dangerous, so it probably ought to be abandoned—or at least severely limited. [See Weizenbaum (1972).] [My view is that the potential danger of AI, along with all other dangers that man presents to himself, will survive at least until we have a science that really understands human emotions. Understanding these emotions, no less than understanding intelligence and perception, will be an ultimate consequence of AI research. Not to understand them is to be at their mercy forever, anyway.]

The one criticism having any weight at all, I think, is that AI may be too broad and diverse to remain a cohesive field. So far, it has stayed together reasonably well. Whether it begins to fractionate into separate exotic applications areas of computer science depends largely, I think, on whether these applications continue to contribute core ideas of great generality.

What is the status of these core ideas today? There are two extreme views. I have heard John McCarthy say (perhaps only provocatively to students) that really intelligent programs are a long way off and that when we finally achieve them they will be based on ideas that aren't around yet. Their builders will look back at AI in 1974 as being a period of pre-history of the field.

On the other hand, what if we already have most of the ideas that we are going to get, ideas like millions of coordinated mini-theories, procedural embedding of knowledge, associative retrieval, and scenario frames. Suppose that we have now only to devote the large effort required to build really huge intelligent systems based on these ideas. To my knowledge, no one advocates this alternative view, but consider this: Whatever the nature of an intelligent system, it will be exceedingly complex. Its performance will derive in large part from its complexity. We will not be sure that AI is ready to build a large, intelligent system until after we have done so. The elegance of the basic ideas and the new and powerful languages alone will not be sufficient indication of our maturity. At some time, we will have to put together exceedingly complex systems. The time at which it is appropriate to try will always be a guess.

My guess is that we still have a good deal of work to do on the problem of how to obtain, represent, coordinate, and use the extensive knowledge we now know is required. But these ideas will not come to those who merely think about the problem. They will come to those who both think and experiment with much larger systems than we have built so far.

Another problem, of a more practical type, concerns knowledge acquisition. Today, the knowledge in a program must be put in "by hand" by the programmer although there are beginning attempts at getting programs to acquire knowledge through on-line interaction with skilled humans. To build really large, knowledgeable systems, we will have to "educate" existing programs rather than attempt the almost impossible feat of giving birth to already competent ones. [Some researchers (e.g., Papert, 1972) expect that at least some of the principles we discover for educating programs will have an impact, perhaps revolutionary, on how we educate people.]

In this connection, we have already mentioned that several successful AI systems use a combination of man and machine to achieve high performance levels. I expect this research strategy to continue and to provide the setting in which the human expert(s) can gradually transfer skills to the machine. [Woods and Makhoul (1973) consciously apply a strategy such as this and call it "incremental simulation."]

I have not yet mentioned in this paper the subject of learning. It is because I have come to agree with John McCarthy that we cannot have a program learn a fact before we know how to tell it that fact and before the program knows how to use that fact. We have been busy with telling and using facts. Learning then is still in the future, although some isolated successes have, in fact, occurred. [See especially, Samuel (1959, 1967), Winston (1970), Fikes et al. (1972a), and Sussman (1973).]

Continuing our discussion of the likely future of AI, we note that the increasingly pragmatic attitude of those who have been sponsoring AI research will have a great effect on the course of this research. There may even be a temporary reduction of effort by AI researchers in the core topics and the first-level applications areas in favor of increased support of engineers and scientists building second-level applications. The results of these second-level efforts may, in fact, be rather spectacular. I have in mind such things as automated factories, automatic robots

for factories and warehouses, medical diagnosis systems, systems that will automate a large amount of office work, legal aids, teaching aids, interactive software production systems, and so on. [Firschein et al. (1973), make some predictions about when these and other intelligent systems may come.]

The short range result of this increased pragmatism may tend to fractionate the field. In the long run, though, if there really are many more core ideas to be discovered, these technological efforts will stimulate their discovery, provided that a sufficient level of basic investigation continues.

In closing, I have one final prediction. As AI successes grow, so will the criticisms of AI, especially from those who are certain that intelligence cannot be mechanized. These critics, having been forced out of various mystical tranches in the past, will be especially vigorous in their defense of what little ground remains to them. The ensuing debates will have the crucially important side effect of getting us all to consider how we want to use and control our new intellectual powers. I hope that society assesses these powers accurately and is not lulled by certain otherwise well-meaning humanists into believing that Artificial Intelligence is not real.

ACKNOWLEDGMENTS

I am grateful for the comments and criticisms of the following people: Woodrow Bledsoe, Stephen Coles, Edward Feigenbaum, Jerome Feldman, Richard Fikes, Cordell Green, Peter Hart, Michael Kassler, John McCarthy, Alton Nowell, Charles Rosen, Earl Sacerdoti, Jay Tenenbaum, Richard Waldinger, and Donald Walker.

BIBLIOGRAPHY

Each entry has a code symbol or symbols associated with one or more of the twelve subheadings of AI that we have discussed in the paper. These symbols are: DED = Common-Sense Reasoning, Deduction, and Problem Solving; REP = Modeling and Representation of Knowledge; SEARCH = Heuristic Search; SYS = AI Systems and Languages; GAME = Game Playing; AIDS = Math, Science, and Engineering Aids; TP = Automatic Theorem Proving; PROG = Automatic Programming; ROB = Robots; VIS = Machine Vision; LANG = Natural Language Systems; PSYC = Information Processing Psychology. A prefix "-G" after a symbol means that the reference contains a general discussion or survey.

The code symbol "GEN" identifies the reference as being general to the whole field of AI. These general references are:

Collins and Michie (1968)	Minsky (1961, 1965, 1968)
Dale and Michie (1968)	Newell (1973)
Dreyfus (1965, 1972)	Papert (1968)
Feigenbaum (1963, 1969)	Papert (1972)
Feigenbaum and Feldman (1963)	Roszak (1972)
Firschein, et al. (1973)	Simon (1969)
Hunt (1974)	Slagle (1971)
Jackson (1974)	Solomonoff (1968)
Lighthill (1973)	Turing (1950)
Meltzer and Michie (1969, 1970, 1971, 1972)	Weizenbaum (1972)
Michie (1968)	

REFERENCES

- Agin, G. (1972), "Representation and Description of Curved Objects," Ph.D. thesis, Elec. Eng. Dept., Stanford Univ., Stanford, CA, July 1972. Available as Stanford Artificial Intelligence Laboratory Memo AIM-173, Oct. 1972. (VIS)
- Agin, Gerald J. and Binford, Thomas O. (1973), "Computer Description of Curved Objects," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (VIS)
- Allen, John and Luckham, David (1970), "An Interactive Theorem-Proving Program," Machine Intelligence, Vol. 5, 321-336, 1970. (TP)
- Amarel, S. (1968), "On Representations of Problems of Reasoning About Actions," Machine Intelligence, Vol. 3, D. Michie, ed., 131-170, Edinburgh Univ. Press, Edinburgh, 1968. (REP)
- Amarel, S. (1969), "On the Representation of Problems and Goal Directed Procedures for Computers," Comm. Am. Soc. Cybernetics, Vol. 1, No. 2, 1969. (SEARCH)
- Ambler, A. P. et al. (1973), "A Versatile Computer-Controlled Assembly System," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (ROB, VIS)
- Anderson, J. R. and Bower, G. H. (1973), Human Associative Memory, V. H. Winston and Sons, Washington, D.C., 1973. (PSYC)
- Andrews, P. B. (1968), "Resolution with Merging," J. ACM, Vol. 15, 367-381, 1968. (TP)
- Balzer, R. M. (1972), "Automatic Programming," Institute Technical Memo, Univ. of So. Calif./Information Sciences Institute, Sept. 1972. (PROG-G)
- Banerji, R. B. (1969), Theory of Problem Solving, American Elsevier Publishing Company, New York, 1969. (GAME)
- Banerji, R. B. and Ernst, G. W. (1972), "Strategy Construction Using Homomorphisms Between Games," Artificial Intelligence, Vol. 3, No. 4, 223-249, Winter 1972. (GAME)
- Barrett, Jeffrey (1972), "A Vocal Data Management System," 1972 Conf. on Speech Communication and Processing, Newton, MA, 24-26 April 1972, U.S. Air Force Cambridge Research Laboratories, Bedford, MA, 22 Feb. 1972, 340-343 (AFCRL-72-0120, Special Reports Number 131). (LANG)
- Barrow, H. and Popplestone, R. (1971), "Relational Descriptions in Picture Processing," Machine Intelligence, Vol. 6, B. Meltzer and D. Michie, eds., 377-396, Edinburgh Univ. Press, 1971. (VIS)
- Barrow, H. G. and Crawford, G. F. (1972), "The Mark 1.5 Edinburgh Robot Facility," Machine Intelligence, Vol. 7, B. Meltzer and D. Michie, eds., 465-480, American Elsevier, 1972. (ROB)
- Barrow, H. G., Ambler, A. P., and Burstall, R. M. (1972), "Some Techniques for Recognizing Structures in Pictures," Frontiers of Pattern Recognition, S. Watanabe, ed., 1-29, Academic Press, New York, 1972. (VIS)
- Bartlett, F. C. (1958), Thinking, Basic Books, New York. (PSYC)
- Bartlett, F. C. (1962), Remembering, Cambridge Univ. Press, Cambridge. (PSYC)
- Becker, J. D. (1970), "An Information-Processing Model of Intermediate-Level Cognition," Memo No. 119, Stanford Artificial Intelligence Project, Computer Science Dept., Stanford Univ., Stanford, CA. Also Report No. 2335, Bolt, Beranek, and Newman, Inc., Cambridge, MA. (PSYC)

- Becker, J. D. (1973), "A Model for the Encoding of Experiential Information," Computer Models of Thought and Language, Schank and Colby, eds., W. H. Freeman and Co., San Francisco, 1973. (PSYC)
- Borlekamp, E. (1963), "Program for Double-Dummy Bridge Problems, A Strategy for Mechanical Game Playing," J. ACM, Vol. 10, No. 3, 357-364, July 1963. (GAME)
- Berliner, Hans (1970), "Experiences Gained in Constructing and Testing a Chess Program," Proc. IEEE Systems Science and Cybernetics Conf., 216-223, Pittsburgh, PA, Oct. 1970. (GAME)
- Berliner, Hans J. (1973), "Some Necessary Conditions for a Master Chess Program," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (GAME)
- Bernstein, A. et al. (1959), "A Chess Playing Program for the IBM 704," Proc. Western Joint Comp. Conf. AIEE, 157-159, Mar. 1959. (GAME)
- Blair, F. W., Griesmer, J. H., and Jenks, R. D. (1970), "An Interactive Facility for Symbolic Mathematics," Proc. Intl. Comp. Symposium, 394-419, Bonn, Germany, 1970. (AIDS)
- Bledsoe, W. W. (1971), "Splitting and Reduction Heuristics in Automatic Theorem Proving," Artificial Intelligence, Vol. 2, No. 1, 55-77, Spring 1971. (TP)
- Bledsoe, W. W., Boyer, R. S., and Henneman, W. H. (1972), "Computer Proofs of Limit Theorems," Artificial Intelligence, Vol. 3, 27-60, 1972. (TP)
- Bledsoe, W. W. and Bruel, P. (1973), "A Man-Machine Theorem Proving System," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, 1973. (TP)
- Bobrow, D. G. (1964a), "Natural Language Input for a Computer Problem-Solving System," Doctoral dissertation, MIT, Sept. 1964. Reprinted in Semantic Information Processing, N. Minsky, ed., MIT Press, Cambridge, MA, 1968. (LANG)
- Bobrow, D. G. (1964b), "A Question-Answering System for High-School Algebra Word Problems," Proc. AFIPS Fall Joint Comp. Conf., 591-614, 1964. (LANG)
- Bobrow, D. G. and Fraser, J. B. (1969), "An Augmented State Transition Network Analysis Procedure," Proc. Intl. Joint Conf. on Artificial Intelligence, 557-567, Washington, D.C., 1969. (LANG)
- Bobrow, D. G. and Raphael, B. (1973), "New Programming Languages for AI Research," SRI Artificial Intelligence Center Tech. Note 82, Stanford Research Institute, Menlo Park, CA, Aug. 1973. To appear in Computer Surveys, 1974. (SYS-G)
- Bobrow, D. G. and Wegbreit, B. (1973a), "A Model for Control Structures for Artificial Intelligence Programming Languages," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (SYS)
- Bobrow, D. G. and Wegbreit, B. (1973b), "A Model and Stack Implementation of Multiple Environments," CACM, Vol. 16, No. 10, Oct. 1973. (SYS)
- Bolles, R. and Paul, R. (1973), "The Use of Sensory Feedback in a Programmable Assembly System," Memo CS-396, Comp. Sci. Dept., Stanford Univ., Artificial Intelligence Lab. Memo AIM-220, Oct. 1973. (ROB)
- Boyer, Robert S. and Moore, J. Strother (1973), "Proving Theorems About LISP Functions," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (PROG)
- Brice, C. R. and Fennema, C. L. (1970), "Scene Analysis Using Regions," Artificial Intelligence, Vol. 1, No. 3, 205-226. (VIS)
- Bruce, Bertram (1972), "A Model for Temporal Reference and Its Application in a Question Answering Program," Artificial Intelligence, Vol. 3, 1-26, 1972. (REP)
- Bruner, J. S., Goodnow, J. J., and Austin, G. A. (1956), A Study of Thinking, Wiley, New York. (PSYC)
- Buchanan, B. G., Sutherland, G., and Foigenbaum, E. (1969), "Heuristic DENDRAL: A Program for Generating Explanatory Hypotheses in Organic Chemistry," Machine Intelligence, Vol. 4, B. Meltzer and D. Michie, eds., 209-254, American Elsevier Publishing Company, New York, 1969. (AIDS)
- Buchanan, B. G. and Lederberg, J. (1971), "The Heuristic DENDRAL Program for Explaining Empirical Data," Proc. IFIP Congress, Vol. 71, Ljubljana, Yugoslavia, 1971. Also Stanford Univ. AIM 141. (AIDS)
- Buchanan, J. R. and Luckham, D. C. (1974), "On Automating the Construction of Programs," Stanford Artificial Intelligence Lab. Memo, forthcoming 1974. (DED, PROG)
- Bundy, Alan (1973), "Doing Arithmetic with Diagrams," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford CA, Aug. 1973. (TP)
- Burstein, R. M., Collins, J. S., and Popplestone, R. J. (1971), Programming in POP2, 279-282, Edinburgh Univ. Press, Edinburgh, 1971. (SYS)
- Carbonell, J. R. (1971), "AI in CAI: An Artificial Intelligence Approach to Computer-Assisted Instruction," IEEE Trans. on Man-Machine Systems, Vol. MMS-11, No. 4, 190-202, Dec. 1970. (AIDS)
- Chang, C. L. and Slagle, J. R. (1971), "An Admissible and Optimal Algorithm for Searching and/or Graphs," Artificial Intelligence, Vol. 2, 117-128, 1971. (SEARCH)
- Chang, C. L. and Lee R. C. (1973), Symbolic Logic and Mechanical Theorem Proving, Academic Press, 1973. (TP-G)
- Charniak, E. C. (1972), "Toward a Model of Children's Story Comprehension," AI TR-266, MIT, Cambridge, MA. (LANG, DED)
- Chase, W. G. (1973), Visual Information Processing, Academic Press, 1973. (PSYC-G)
- Chomsky, N. (1956), "Three Models for the Description of Language," IRE Trans. on Info. Theory, Vol. IT-2(3), 113-124, 1956. (PSYC, LANG)
- Chomsky, N. (1957), Syntactic Structures, Mouton, The Hague, 1957. (PSYC, LANG)
- Chomsky, N. (1965), Aspects of the Theory of Syntax, MIT Press, Cambridge, MA, 1965. (LANG)
- Cioux, M. G. (1971), "On Seeing Things," Artificial Intelligence, Vol. 2, No. 1, 79-118, 1971. (VIS)
- Colby, K. M. and Enea, H. (1987), "Heuristic Methods for Computer Understanding of Natural Language in Context Restricted On-Line Dialogues," Mathematical Biosciences, Vol. 1, 1-25, 1967. (LANG)
- Colby, K. M., Weber, S., and Hilf, F. D. (1971), "Artificial Paranoia," Artificial Intelligence, Vol. 2, No. 1, 1-25, Spring 1971. (PSYC)
- Colles, L. S. (1970), "An Experiment in Robot Tool Using," Stanford Research Institute Tech. Note No. 41, Stanford Research Institute, Menlo Park, CA, Oct. 1970. (ROB)
- Colles, L. S. (1972), "Techniques for Information Retrieval Using an Inferential Question-Answering System with Natural-Language Input," SRI Artificial Intelligence Center Tech. Note 74, Stanford Research Institute, Menlo Park, CA, Nov. 1972. (LANG)

- Collins, A. M. and Quillian, M. R. (1972), "Retrieval Time from Semantic Memory," J. Verbal Learning and Verbal Behavior, Vol. 8, 240-247, 1969. (PSYC)
- Collins, N. and Michie, D., eds. (1968), Machine Intelligence, Vol. 1, American Elsevier Publishing Company, New York, 1967. (GEN)
- Corcy, E. J. (1969), "Computer-Assisted Design of Complex Organic Synthesis," Science, 10 Oct. 1969. (AIDS)
- Dale, E. and Michie, D., eds. (1968), Machine Intelligence, Vol. 2, American Elsevier Publishing Company, New York, 1968. (GEN)
- Darlington, J. L. (1971), "A Partial Mechanization of Second-Order Logic," Machine Intelligence, Vol. 6, 91-100, B. Meltzer and D. Michie, eds., Edinburgh Univ. Press, Edinburgh. (TP)
- Davies, D. J. M. (1971), "POPLER: A POP-2 Planner," Memo MIP-R-89, School of Artificial Intelligence, Univ. of Edinburgh. (SYS)
- Davis, M. and Putnam, H. (1960), "A Computing Procedure for Quantification Theory," J. ACM, Vol. 7, 201-215, 1960. (TP)
- Derksen, J., Rulifson, J. F., and Waldinger, R. J. (1972), "The QA4 Language Applied to Robot Planning," AFIPS Conf. Proc., Vol. 41, Part II, 1181-1187, Fall Joint Comp. Conf., 1972. (DED)
- Deutsch, L. P. (1973), "An Interactive Program Verifier," Ph.D. thesis, Dept. of Computer Science, Univ. of California, Berkeley, 1973. (PROG)
- Doran, J. and Michie, D. (1966), "Experiments with the Graph Traverser Program," Proc. Roy. Soc. A, Vol. 294, 235-259, 1966. (SEARCH)
- Dreyfus, H. (1965), "Alchemy and Artificial Intelligence," RAND Corporation Paper P3244 (AD 625 719), Dec. 1965. (GEN)
- Dreyfus, H. L. (1972), What Computers Can't Do, Harpor and Row, 1972. (GEN)
- Duda, R. and Hart, P. (1970), "Experiments in Scene Analysis," Proc. 1st Natl. Symposium on Industrial Robots, IIT Research Institute, Chicago, IL, Apr. 1970. (VIS)
- Duda, R. and Hart, P. (1973), Pattern Classification and Scene Analysis, John Wiley & Sons, New York, 1973. (VIS-C)
- Eastman, C. M. (1971a), "GSP: A System for Computer Assisted Space Planning," Proc. 6th Annual Design Automation Workshop, Atlantic City, NJ. (AIDS)
- Eastman, C. M. (1971b), "Heuristic Algorithms for Automated Space Planning," Proc. 2d Intl. Joint Conf. on Artificial Intelligence, Imperial College, London, 1971. (AIDS)
- Edmundson, H. P., ed. (1961), Proc. of the Natl. Symposium on Machine Translation, Prentice-Hall, Englewood Cliffs, NJ, 1961. (LANG)
- Edwards, D. and Hart, T. (1961), "The Alpha-Beta Heuristic," MIT Artificial Intelligence Memo No. 30 (revised), 28 Oct. 1963. Originally printed as "The Tree Prune (TP) Algorithm," 4 Dec. 1961. (GAME)
- Ejiri, M. et al. (1971), "An Intelligent Robot with Cognition and Decision-Making Ability," Proc. of 2d Intl. Joint Conf. on Artificial Intelligence, Imperial College, London, 350-358, 1971. (ROB, VIS)
- Ejiri, E. et al. (1972), "A Prototype Intelligent Robot that Assembles Objects from Plan Drawings," IEEE Trans. Comp., 161-170, Feb. 1972. (ROB, VIS)
- Elcock, E. W. et al. (1971), "ANSSET, A Programming Language Based on Sets: Motivation and Examples," Machine Intelligence, Vol. 6, Edinburgh Univ. Press, Edinburgh, 1971. (SYS)
- Elspas, B. (1972), "The Semiautomatic Generation of Inductive Assertions for Program Correctness Proofs," Report No. 55, Seminar, Des Instituts fur Theorie der Automaten und Schaltnetzwerke, Gesellschaft fur Mathematik und Datenverarbeitung, Bonn, 21 Aug. 1972. (PROG)
- Elspas, D. et al. (1973), "Design of an Interactive System for Verification of Computer Programs," SRI Report, Project 1891, Stanford Research Institute, Menlo Park, CA, July 1973. (PROG)
- Enea, H. and Colby, K. M. (1973), "Idioelectric Language-Analysis for Understanding Doctor-Patient Dialogues," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, 278-284, Stanford Univ., Stanford, CA, 1973. (LANG)
- Engelman, C. (1969), "MATLAB 68," Information Processing 68, A. J. H. Morrell, ed., 462-467, North-Holland Publishing Company, Amsterdam, 1969. (AIDS)
- Ernst, G. W. and Newell, Allen (1969), GPS: A Case Study in Generality and Problem Solving, Academic Press, New York, 1969. (DED)
- Ernst, G. W. (1971), "The Utility of Independent Subgoals in Theorem Proving," Information and Control, Apr. 1971. (TP)
- Ernst, H. (1961), "MH-1, A Computer-Operated Mechanical Hand," O.Sc. dissertation, Dept. of Elec. Eng., MIT, Cambridge, MA. (ROB)
- Fahlman, S. (1973), "A Planning System for Robot Construction Tasks," Report AI TR-283, Artificial Intelligence Lab., MIT, Cambridge, MA, May 1973. (DED)
- Falk, G. (1970), "Computer Interpretation of Imperfect Line Data as a Three-Dimensional Scene," Ph.D. thesis, Stanford Univ., Comp. Sci. Dept., 1970. Available as CS-180 and AIM-132. (VIS)
- Falk, G. (1972), "Interpretation of Imperfect Line Data as a Three Dimensional Scene," Artificial Intelligence, Vol. 3, No. 2, 101-144, 1972. (VIS)
- Feigenbaum, E. A. (1961), "The Simulation of Verbal Learning Behavior," Proc. Western Joint Comp. Conf., 121-132, 1961. Also in Computers and Thought, E. A. Feigenbaum and J. Feldman, eds., 297-309, McGraw-Hill Book Company, New York, 1963. (PSYC)
- Feigenbaum, E. (1963), "Artificial Intelligence Research," IEEE Trans. Info. Theory, Vol. IT-9, No. 4, 248-261, Oct. 1963. (GEN)
- Feigenbaum, E. (1969), "Artificial Intelligence: Themes in the Second Decade," Information Processing 68, Vol. 2, A.J.H. Morrell, ed., 1008-1022, North-Holland Publishing Company, Amsterdam, 1969. Also printed as Stanford Univ. Artificial Intelligence Project Memo No. 67, 15 Aug. 1968. (GEN)
- Feigenbaum, E. and Feldman, J., eds. (1963), Computers and Thought, McGraw-Hill Book Company, New York, 1963. (GEN)
- Feldman, J. A. et al. (1969), "The Stanford Hand-Eye Project," Proc. 1st Intl. Joint Conf. on Artificial Intelligence, 521-526, Washington, D.C., 1969. (ROB)
- Feldman, J. A. et al. (1971), "The Use of Vision and Manipulation to Solve the Instant Insanity Puzzle," Proc. 2d Intl. Joint Conf. on Artificial Intelligence, London, 1971. (ROB, VIS)
- Feldman, J. A. et al. (1972), "Recent Developments in SAIF--An ALGOL-Based Language for Artificial Intelligence," 1972 FJCC Proc., 5-7 Dec. 1972, Anaheim, CA. (SYS)
- Feldman, J. A. and Rovner, P. D. (1969), "An ALGOL-Based Associative Language," Comm. ACM, 434-449, Aug. 1969. (SYS)

- Fikes, R. E. (1968), "A Heuristic Program for Solving Problems Stated as Nondeterministic Procedures," Ph.D. thesis, Carnegie-Mellon Univ., 1968. (DED, SYS)
- Fikes, R. E. (1970), "REF-ANF: A System for Solving Problems Stated as Procedures," Artificial Intelligence, Vol. 1(1), 1970. (DED, SYS)
- Fikes, R. E. and Nilsson, N. J. (1971), "STRIPS: A New Approach to the Application of Theorem Proving in Problem Solving," Artificial Intelligence, Vol. 2, 189-208, 1971. (REP, DED)
- Fikes, R. E., Hart, P. E., and Nilsson, N. J. (1972a), "Learning and Executing Generalized Robot Plans," Artificial Intelligence, Vol. 3, 251-288, 1972. (DED)
- Fikes, R. E., Hart, P. E., and Nilsson, N. J. (1972b), "Some New Directions in Robot Problem Solving," B. Meltzer and D. Michie, eds., Machine Intelligence, Vol. 7, Edinburgh Univ. Press, Edinburgh, 1972. (ROB-G)
- Firschein, Oscar et al. (1973), "Forecasting and Assessing the Impact of Artificial Intelligence on Society," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (GEN)
- Fischler, Martin A. and Elschlager, Robert A. (1973), "The Representation and Matching of Pictorial Structures," IEEE Trans. on Computers, Vol. C-22, No. 1, 67-92, Jan. 1973. (VIS)
- Flanagan, J. L. (1965), Speech Analysis, Synthesis and Perception, Academic Press, New York. (LANG)
- Floyd, R. W. (1967), "Assigning Meanings to Programs," Proc. of a Symposium in Applied Mathematics, Vol. 19, J. T. Schwartz, ed., Am. Math. Soc., 19-32, 1967. (PROG)
- Forsen, G. (1968), "Processing Visual Data with an Automaton Eye," Pictorial Pattern Recognition, 471-502, Thompson Book Co., Washington, D.C., 1968. (VIS)
- Friedman, Joyce (1971), A Computer Model of Transformational Grammar, 166, American Elsevier Publishing Company, New York, 1971. (LANG)
- Friedman, L. (1967), "Instinctive Behavior and Its Computer Synthesis," Behavioral Science, Vol. 12, No. 2, Mar. 1967. (PSYC)
- Fuller, S., Gaschnig, J., and Gillogly, J. (1973), "Analysis of the Alpha-Beta Pruning Algorithm," Carnegie-Mellon Univ., Dept. of Comp. Sci. Report, Pittsburgh, PA, July 1973. (GAME)
- Gelernter, H. (1960), "Realization of a Geometry Theorem-Proving Machine," Proc. 1959 Intl. Conf. on Info. Proc., 273-282, UNESCO, Paris. Also in Computers and Thought, E. A. Feigenbaum and J. Feldman, eds., 134-152, McGraw-Hill Book Company, New York, 1963. (TP, SEARCH)
- Gelernter, H., Hansen, J. R., and Loveland, D. W. (1963), "Empirical Explorations of the Geometry Theorem Machine," Computers and Thought, E. A. Feigenbaum and J. Feldman, eds., 153-163, McGraw-Hill Book Company, New York, 1963. (TP)
- Gere, W. S. (1966), "Heuristics in Job Shop Scheduling," Management Science, Vol. 13, 167-190. (AIDS)
- Gillogly, J. (1972), "The Technology Chess Program," Artificial Intelligence, Vol. 3, No. 3, 145-163, Fall 1972. (GAME)
- Goldstein, A. Jay, Harmon, Leon D., and Lask, Ann D. (1971), "Identification of Human Faces," Proc. IEEE, Vol. 59, No. 5, 748-760, May 1971. (VIS)
- Good, I. J. (1967), "A Five-Year Plan for Automatic Chess," Machine Intelligence, Vol. 2, E. Dale and D. Michie, eds., 89-118, Edinburgh Univ. Press, Edinburgh, 1967. (GAME)
- Good, D. I. and London, R. L. (1968), "Interval Arithmetic for the Burroughs B5500: Four ALGOL Procedures and Proofs of Their Correctness," Comp. Sci. Tech. Report No. 26, Univ. of Wisconsin, 1968. (PROG)
- Gordon, G. (1969), System Simulation, Prentice-Hall, Englewood Cliffs, NJ, 1969. (SYS)
- Grape, G. (1973), "Model Based (Intermediate Level) Computer Vision," Ph.D. thesis, Comp. Sci. Dept., Stanford Univ., Stanford, CA, 1973. (VIS)
- Greenblatt, R. et al. (1967), "The Greenblatt Chess Program," Proc. AFIPS Fall Joint Comp. Conf., 801-810, 1967. (GAME)
- Green, B. F. et al. (1961), "Baseball: An Automatic Question Answerer," Proc. Western Joint Comp. Conf., 219-224. (LANG)
- Green, C. (1969a), "Theorem-Proving by Resolution as a Basis for Question-Answering Systems," Machine Intelligence, Vol. 4, B. Meltzer and D. Michie, eds., 183-205, American Elsevier Publishing Company, New York, 1969. (REP, DED)
- Green, C. (1969b), "The Application of Theorem-Proving to Question-Answering Systems," Doctoral dissertation, Elec. Eng. Dept., Stanford Univ., Stanford, CA, June 1969. Also printed as Stanford Artificial Intelligence Project Memo AI-96, June 1969. (REP, DED, PROG, TP)
- Green, C. (1969c), "Application of Theorem-Proving to Problem Solving," Proc. Intl. Joint Conf. Artificial Intelligence, Donald E. Walker and Lewis M. Norton, eds., Washington, D.C., May 1969. (REP, DED)
- Gregory, R. L. (1966), Eye and Brain, McGraw-Hill Book Company, New York. (PSYC)
- Gregory, R. L. (1970), The Intelligent Eye, Weidenfeld and Nicolson, London, 1970. (PSYC)
- Griesmer, J. H. and Jenks, R. D. (1971), "SCRATCHPAD/1--An Interactive Facility for Symbolic Mathematics," Proc. ACM 2d Symposium on Symbolic and Algebraic Manipulation, S. R. Petrick, ed., Los Angeles, CA, 23-25 Mar. 1971. (AIDS)
- Griffith, A. K. (1970), "Computer Recognition of Prismatic Solids," MAC Tech. Report 73, Project MAC, MIT, Cambridge, MA. (VIS)
- Griffith, A. K. (1973), "Mathematical Models for Automatic Line Detection," J. ACM, 62-80, 1973. (VIS)
- Gross, Louis N. and Walker, Donald E. (1969), "On-Line Computer Aids for Research in Linguistics," Information Processing, Vol. 68, A. J. H. Morrell, ed., North-Holland Publishing Company, Amsterdam, 1969. (LANG)
- Guard, J. R. et al. (1969), "Semi-Automated Mathematics," J. ACM, Vol. 16, 49-62, 1969. (TP)
- Guzman, A. (1968a), "Decomposition of a Visual Scene Into Three-Dimensional Bodies," Proc. AFIPS 1968 Fall Joint Comp. Conf., Vol. 33, 291-304, Thompson Book Co., Washington, D.C. (VIS)
- Guzman, A. (1968b), "Computer Recognition of Three-Dimensional Objects in a Visual Scene," MAC Tech. Report 59, thesis, Project MAC, MIT, Cambridge, MA, 1968. (VIS)
- Guzman, A. (1971), "Analysis of Curved Line Drawings Using Context and Global Information," Machine Intelligence, Vol. 6, D. Meltzer and D. Michie, eds., 325-376, Edinburgh Univ. Press, Edinburgh. (VIS)
- Hacssler, R. W. (1971), "A Heuristic Programming Solution to a Nonlinear Cutting Stock Problem," Management Science, Vol. 17B, 793-802. (AIDS)
- Harris, Z. (1951), Structural Linguistics, Univ. of Chicago Press, Chicago, 1951. (LANG)

- Harris, Z. (1961), String Analysis of Sentence Structure, Mouton, The Hague, 1961. (LANG)
- Hart, P. E. et al. (1972), "Artificial Intelligence-- Research and Applications," Annual Tech. Report to ARPA, Contract DAHCO4-72-C-0008, Stanford Research Institute, Menlo Park, CA, Dec. 1972. (ROB)
- Hart, P., Nilsson, N., and Raphael, B. (1968), "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," IEEE Trans. Sys. Sci. Cybernetics, Vol. 4, No. 2, 100-107, 1968. (SEARCH)
- Hart, T. (1961), "SIMPLIFY," Memo 27, Artificial Intelligence Group, Project MAC, MIT, Cambridge, MA, 1961. (AIDS)
- Hayes, P. (1971), "A Logic of Action," Machine Intelligence, Vol. 6, B. Meltzer and D. Michie, eds., 495-520, American Elsevier Publishing Company, 1971. (REP)
- Hayes, P. F. (1973), "The Frame Problem and Related Problems in Artificial Intelligence," Artificial and Human Thinking, A. Elithorn and D. Jones, eds., Elsevier Scientific Publishing Co., New York, 1973. (REP)
- Hearn, A. C. (1968), "REDUCE, A User-Oriented Interactive System for Algebraic Simplification," Interactive Systems for Experimental Applied Mathematics, 79-90, M. Klerer and J. Reinfelds, eds., Academic Press, New York and London, 1968. (AIDS)
- Hearn, Anthony C. (1971), "REDUCE 2: A System and Language for Algebraic Manipulation," Proc. ACM 2d Symposium on Symbolic and Algebraic Manipulation, S. R. Petrick, ed., 23-25 Mar. 1971, Los Angeles, CA. (AIDS)
- Hendrix, G. (1973), "Modeling Simultaneous Actions and Continuous Processes," Artificial Intelligence, Vol. 4, 145-180, 1973. (REP)
- Hendrix, Gary G., Thompson, Craig W., and Slocum, Jonathan (1973), "Language Processing via Canonical Verbs and Semantic Models," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (LANG)
- Hewitt, C. (1969), "PLANNER: A Language for Proving Theorems in Robots," 1st Intl. Joint Conf. on Artificial Intelligence, Washington, D.C., 1969. (REP, SYS, DED)
- Hewitt, C. (1971), "Procedural Embedding of Knowledge in PLANNER," Proc. 2d Intl. Joint Conf. on Artificial Intelligence, British Computer Society, London, England, 167-182, 1971. (REP, SYS, DED)
- Hewitt, C. (1972), "Description and Theoretical Analysis (Using Schemata) of PLANNER: A Language for Proving Theorems and Manipulating Models in a Robot," Ph.D. thesis, Dept. of Math., MIT, Cambridge, MA, 1972. (SYS, REP, DED)
- Hewitt, C., Bishop, P., and Steiger, R. (1973), "A Universal Modular Actor Formalism for Artificial Intelligence," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (SYS)
- Hintzman, D. L. (1968), "Explorations with a Discrimination Net Model for Paired-Associate Learning," J. Mathematical Psychology, Vol. 5, 123-162, 1968. (PSYC)
- Horn, B. K. P. (1970), "Shape from Shading: A Method for Obtaining the Shape of a Smooth Opaque Object from One View," MAC Tech. Report 79, Project MAC, MIT, Cambridge, MA, Nov. 1970. (VIS)
- Horn, B. K. P. (1971), "The Binford-Horn Line Finder," Vision Flash, No. 16, Artificial Intelligence Laboratory, MIT, Cambridge, MA. Later issued as MIT Artificial Intelligence Lab Memo 285, Mar. 1973. (VIS)
- Huet, G. P. (1973a), "A Unification Algorithm for Type Theory," IRIA Laborin, 1973. (TP)
- Huet, G. P. (1973b), "A Mechanization of Type Theory," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (TP)
- Huffman, D. A. (1971), "Impossible Objects as Nonsense Sentences," Machine Intelligence, Vol. 6, B. Meltzer and D. Michie, eds. 295-323, Edinburgh Univ. Press, Edinburgh, 1971. (VIS)
- Hunt, E. B. (1962), Concept Formation, John Wiley & Sons, New York. (PSYC)
- Hunt, E. B. (1974), Artificial Intelligence, Academic Press, 1974. (GEN)
- Hunt, E. B. and Howland, C. I. (1961), "Programming a Model of Human Concept Formation," Proc. Western Joint Comp. Conf., Vol. 19, 145-155. (PSYC)
- Igarashi, W., London, R., and Luckham, D. (1973), "Automatic Verification of Programs I: A Logical Basis and Implementation," Stanford Univ. Artificial Intelligence Lab. Memo, No. 200, Stanford Univ., Stanford, CA, May 1973. (PROG)
- Jackson, P. C. (1974), Introduction to Artificial Intelligence, Mason and Lipscomb, New York, 1974. (GEN)
- Kaplan, R. M. (1972), "Augmented Transition Networks as Psychological Models of Sentence Comprehension," Artificial Intelligence, Vol. 3, 77-100, 1972. (PSYC)
- Kaplan, R. M. (1973), "A General Syntactic Processor," Natural Language Processing, R. Rustin, ed., 293-241, Algorithmic Press, New York, 1973. (LANG)
- Katz, S. M. and Nanna, Z. (1973), "Heuristic Approach to Program Verification," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, 500-512, Stanford Univ., Stanford, CA, Aug. 1973. (PROG)
- Kay, Martin (1964), "A General Procedure for Rewriting Strings," presented at the 1964 Annual Meeting, Association for Machine Translation and Computational Linguistics, Indiana Univ., Bloomington. (LANG)
- Kay, Martin (1967), "Experiments with a Powerful Parser," RM-5452-PR, RAND Corporation, Santa Monica, CA, 1967. (LANG)
- Kay, Martin (1973), "The Mind System," Natural Language Processing, R. Rustin, ed., 155-167, Courant Computer Science Symposium 8, 20-21 Dec., Algorithmics Press, Inc., New York, 1973. (LANG)
- Kelly, M. (1970), "Visual Identification of People by Computer," Memo AI-130, Comp. Sci. Dept., Stanford Univ., Stanford, CA, July 1970. (VIS)
- King, J. (1969), "A Program Verifier," Doctoral dissertation, Comp. Sci. Dept., Carnegie-Mellon Univ., Pittsburgh, PA, 1969. (PROG)
- Kister, J. et al. (1957), "Experiments in Chess," J. ACM, Vol. 4, 174-177, Apr. 1957. (GAME)
- Kling, R. E. (1971), "A Paradigm for Reasoning by Analogy," Artificial Intelligence, Vol. 2, No. 2, 147-178, Fall 1971. (DED)
- Koffman, Elliot B. and Blount, Sumner E. (1973), "Artificial Intelligence and Automatic Programming in CAI," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (AIDS)
- Korsvold, K. (1965), "An On-Line Algebraic Simplification Program," Artificial Intelligence Project Memo No. 37, Stanford Univ., Stanford, CA, Nov. 1965. (AIDS)
- Kotok, A. (1962), "A Chess Playing Program for the IBM 7090," Bachelor's thesis, MIT, 1962. (GAME)

- Kowalski, R. (1970), "Search Strategies for Theorem Proving," Machine Intelligence, Vol. 5, B. Meltzer and D. Michie, eds., 181-200, American Elsevier Publishing Company, New York, 1970. (TP)
- Kowalski, R. and Kuehner, D. (1971), "Linear Resolution with Selection Function," Artificial Intelligence, Vol. 2, 227-260, 1971. (TP)
- Kulikowski, C. A. and Weiss, S. (1972), "The Medical Consultant Program--Glaucoma," Tech. Report TR-5, Computers in Biomedicine, Dept. of Comp. Sci., Rutgers Univ., New Brunswick, NJ, July 1972. (AIDS)
- Kuno, Susumu and Oettinger, Anthony G. (1963), "Multiple-Path Syntactic Analyzer," Information Processing 1962, 306-312, North-Holland Publishing Company, Amsterdam, 1963. (LANG)
- Levy, B. N. L. (1970), "Computer Chess--A Case Study," Machine Intelligence, Vol. 6, B. Meltzer and D. Michie, eds., 151-163, Edinburgh Univ. Press, 1970. (GAME)
- Lighthill, J. (1973), "Artificial Intelligence: A General Survey," Artificial Intelligence: A Paper Symposium, Science Research Council Pamphlet, Science Research Council, State House, High Holburn, London, Apr. 1973. (GEN)
- Lin, Shen (1970), "Heuristic Techniques for Solving Large Combinatorial Problems on a Computer," Theoretical Approaches to Non-Numerical Problem-Solving, R. Banerji and M. Mesarovic, eds., 410-418, Springer-Verlag, New York, 1970. (AIDS)
- Lindsay, P. H. and Norman, D. A. (1972), Human Information Processing: An Introduction to Psychology, Academic Press, 1972. (PSYC-G)
- Locke, William N. and Booth, A. Donald, eds. (1955), Machine Translation of Languages: Fourteen Essays, John Wiley & Sons, New York, 1955. (LANG)
- Londe, Davo L. and Schoene, William J. (1968), "TGT: Transformational Grammar Tester," Proc. AFIPS Conf., Vol. 32, 1968 Spring Joint Comp. Conf., 385-393, Thompson Book Co., Washington, D.C. (LANG)
- London, R. (1970), "Bibliography on Proving the Correctness of Computer Programs," Machine Intelligence, Vol. 5, B. Meltzer and D. Michie, eds., 569-580, American Elsevier Publishing Company, New York, 1970. (PROG-G)
- London, R. L. (1972), "The Current State of Proving Programs Correct," Proc. ACM 25th Annual Conf., 39-46, 1972. (PROG-G)
- Loveland, D. W. (1970), "A Linear Format for Resolution," Symposium on Automatic Demonstration, M. Laudet, ed., Springer-Verlag, New York, 1970. (TP)
- Luckham, D. (1969), "Refinement Theorems in Resolution Theory," Stanford Artificial Intelligence Project Memo AI-81, 24 Mar. 1969. Also in Proc. IRIA 1968 Symp. Autom. Demonstration, Lecture Notes on Mathematics No. 125, Springer-Verlag, New York, 1970. (TP)
- Luckham, D. and Nilsson, N. (1971), "Extracting Information from Resolution Proof Trees," Artificial Intelligence, 1971. (TP)
- Manna, Z. (1969), "The Correctness of Programs," J. Computer Syst. Sci., Vol. 3, May 1969. (PROG)
- Manna, Z. and Waldinger, R. J. (1971), "Toward Automatic Program Synthesis," Comm. ACM, Vol. 14, No. 3, 151-165, Mar. 1971. (PROG-G)
- Martin, W. A. (1967), "Symbolic Mathematical Laboratory," MAC Tech. Report 36 (thesis), Project MAC, MIT, Cambridge, MA, Jan. 1967. (AIDS)
- Martin, W. A. and Fateman, R. J. (1971), "The MACSYMA System," Proc. ACM 2d Symposium on Symbolic and Algebraic Manipulation, S. R. Petrick, ed., Los Angeles, CA, 23-25 Mar. 1971. (AIDS)
- McCarthy, J. (1958), "Programs with Common Sense," in "Mechanization of Thought Processes," Vol. 1, 77-84, Proc. Symp., Nat. Phys. Lab., London, 24-27, Nov. 1958. Reprinted in Semantic Information Processing, M. Minsky, ed., 40J-410, MIT Press, Cambridge, MA, 1968. (REP, DED)
- McCarthy, J. (1960), "Recursive Functions of Symbolic Expressions and Their Computation by Machine, Part 1," Comm. ACM, Vol. 3, No. 4, 184-195, Apr. 1960. (SYS)
- McCarthy, John (1961), "Computer Programs for Checking Mathematical Proofs," Proc. Amer. Math. Soc. on Recursive Function Theory, New York, Apr. 1961. (TP)
- McCarthy, J. (1962), "Towards a Mathematical Science of Computation," Proc. IFIP Congr., Vol. 62, North-Holland Publishing Company, Amsterdam, 1962. (PROG)
- McCarthy, J. (1963), "Situations, Actions and Causal Laws," Memo. No. 2, Stanford Univ. Artificial Intelligence Project, 1963. Reprinted in Semantic Information Processing, M. Minsky, ed., 410-418, MIT Press, Cambridge, MA, 1968. (REP)
- McCarthy, J. and Painter, J. A. (1967), "Correctness of a Compiler for Arithmetic Expressions," in Proc. Symp. Applied Mathematics, Vol. 19, Math. Aspects of Computer Science, J. T. Schwartz, ed., 33-41, American Mathematical Society, Providence, RI, 1967. (PROG)
- McCarthy, J. et al. (1968), "A Computer with Hands, Eyes, and Ears," Proc. 1968 Fall Joint Comp. Conf., Vol. 33, 329-338, Thompson Book Company, Washington, D.C. (ROB)
- McCarthy, J. and Hayes, P. (1969), "Some Philosophical Problems from the Standpoint of Artificial Intelligence," Machine Intelligence, Vol. 4, American Elsevier Publishing Company, New York, 1969. (REP)
- McDermott, D. V. and Sussman, G. J. (1972), "The CONNIVER Reference Manual," MIT, Artificial Intelligence Lab., Memo No. 259, May 1972. (SYS)
- Meltzer, B. and Michie, D., eds. (1969), Machine Intelligence, Vol. 4, American Elsevier Publishing Company, New York, 1969. (GEN)
- Meltzer, B. and Michie, D., eds. (1970), Machine Intelligence, Vol. 5, American Elsevier Publishing Company, New York, 1970. (GEN)
- Meltzer, B. and Michie, D., eds. (1971), Machine Intelligence, Vol. 6, American Elsevier Publishing Company, New York, 1971. (GEN)
- Meltzer, B. and Michie, D., eds. (1972), Machine Intelligence, Vol. 7, American Elsevier Publishing Company, New York, 1972. (GEN)
- Michie, D., ed. (1968), Machine Intelligence, Vol. 3, American Elsevier Publishing Company, Princeton, NJ, 1968. (GEN)
- Miller, G. A. (1956), "The Magical Number Seven, Plus or Minus Two," Psychological Review, Vol. 63, 81-97. (PSYC)
- Miller, G. A., Galanter, E., and Pribram, K. H. (1960), Plans and the Structure of Behavior, Holt, Rinehart & Winston, New York. (PSYC)
- Minker, J., Fishman, D. H., and McSkimin, J. R. (1972), "The Maryland Refutation Proof Procedure System (MRPPS)," TR-208, Comp. Sci. Center, Univ. of Maryland, College Park, MD, 1972. (TP)
- Minsky, M. (1961), "Steps Toward Artificial Intelligence," Proc. IRE, Vol. 49, 8-30, Jan. 1961. (SEARCH, GEN)
- Minsky, M. L. (1965), "Matter, Mind, and Models," IFIP, 1965. (GEN)
- Minsky, M., ed. (1968), Semantic Information Processing, MIT Press, Cambridge, MA. (GEN, LANG-G)
- Minsky, M. (1974), "Frame-Systems: A Framework for Representation of Knowledge," forthcoming 1974. (REP)

- Minsky, M. and Papert, S. (1972), Progress Report, AI Memo 252, Artificial Intelligence Laboratory, Cambridge, MA. (REP, VIS-G)
- Mittman, B. (1973), "Can a Computer Beat Bobby Fischer?", Dataamation, 84-87, June 1973. (GAME)
- Moore, J. and Newell, A. (1973), "How Can Merlin Understand?", Dept. of Com. Sci. Report, Carnegie-Mellon Univ., Nov. 15, 1973. (REP)
- Moses, J. (1967), "Symbolic Integration," MAC Tech. Report 47, Project MAC, MIT, Cambridge, MA, Dec. 1967. (AIDS)
- Moses, Joel (1971a), "Symbolic Integration: The Stormy Decade," Proc. ACM 2d Symposium on Symbolic and Algebraic Manipulation, S. R. Petrick, ed., Los Angeles, CA, 23-25 Mar. 1971. (AIDS-G)
- Moses, Joel (1971b), "Algebraic Simplification: A Guide for the Perplexed," Proc. ACM 2d Symposium on Symbolic and Algebraic Manipulation, S. R. Petrick, ed., Los Angeles, CA, 23-25 Mar. 1971. (AIDS-G)
- Neisser, U. (1967), Cognitive Psychology, Appleton-Century-Crofts, New York. (PSYC)
- Novatia, Ramakant and Binford, Thomas O. (1973), "Structured Descriptions of Complex Objects," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, 641-645, Stanford Univ., Stanford, CA, Aug. 1973. (VIS)
- Nevins, A. J. (1972), "A Human Oriented Logic for Automatic Theorem Proving," Tech. Report 268, MIT Artificial Intelligence Laboratory, MIT, Cambridge, MA, Oct. 1972. (TP)
- Nevins, J. L., Whitney, D. E., and Simunovic, S. N. (1973), "Report on Advanced Automation," No. R-764, prepared for National Science Foundation, Grant No. GK-34094 and GI-39432X, The Charles Stark Draper Lab., Inc., Cambridge, MA, Nov. 1973. (ROB)
- Newell, A., ed. (1961), Information Processing Language V Manual, Prentice-Hall, Englewood Cliffs, NJ. (SYS)
- Newell, A. (1967), Studies in Problem Solving: Subject 3 on the Cryptarithmic Task: Donald + Gerald = Robert, Carnegie-Mellon Univ., Pittsburgh, PA. (PSYC, REP)
- Newell, A. (1970), "Remarks on the Relationship Between Artificial Intelligence and Cognitive Psychology," Theoretical Approaches to Non-Numerical Problem Solving, R. B. Banerji and M. D. Mesarovic, eds., Springer-Verlag, Berlin. (PSYC-G)
- Newell, A. (1972a), "Production Systems: Models of Control Structures," Visual Information Processing, Wm. Chase, ed., Academic Press, New York, 1972. (REP, PSYC)
- Newell, A. (1972b), "A Theoretical Exploration of Mechanisms for Coding the Stimulus," Coding Processes in Human Memory, A. Melter and E. Martin, eds., V. H. Winston, Washington, D.C., 1972. (REP, PSYC)
- Newell, A. (1973), "Artificial Intelligence and the Concept of Mind," Computer Models of Thought and Language, Roger Schank and Kenneth Colby, eds., W. H. Freeman & Co., San Francisco, 1973. (GEN)
- Newell, A. and Shaw, J. C. (1957), "Programming the Logic Theory Machine," Proc. West. Joint Comp. Conf., 230-240, 1957. (SYS)
- Newell, A., Shaw, J. C., and Simon, Herbert (1957), "Empirical Explorations with the Logic Theory Machine: A Case Study in Heuristics," Proc. Western Joint Comp. Conf. 1957, 218-239. Also in Computers and Thought, E. A. Feigenbaum and J. Feldman, eds., 109-133, McGraw-Hill Book Company, New York, 1963. (DED, PSYC)
- Newell, A., Shaw, J. C., and Simon, H. A. (1958a), "Elements of a Theory of Human Problem Solving," Psychological Review, Vol. 65, 151-166, 1958. (PSYC)
- Newell, A., Shaw, J., and Simon, H. (1958b), "Chess Playing Programs and the Problem of Complexity," IBM J. Res. Develop., Vol. 2, 320-335, Oct. 1958. Reprinted in Computers and Thought, Feigenbaum and Feldman, eds., 39-70, McGraw-Hill Book Company, New York, 1963. (GAME)
- Newell, A., Shaw, J. C., and Simon, H. A. (1960), "Report on a General Problem-Solving Program for a Computer," Information Processing: Proc. Intl. Conf. Information Processing, 256-264, UNE500, Paris. Also printed in Computers and Automation, July 1959. (DED, PSYC, TP)
- Newell, A. and Simon, H. A. (1956), "The Logic Theory Machine: A Complex Information Processing System," IRE Trans. on Information Theory, Vol. IT-2, No. 3, 61-79. (SEARCH, DED, PSYC, TP)
- Newell, A. and Simon, H. (1961), "GPS, A Program that Simulates Human Thought," Lernende Automaten, H. Billig, ed., 109-124, E. Oldenbourg, Munich. Reprinted in Computers and Thought, Feigenbaum and Feldman, eds., McGraw-Hill Book Company, New York, 1963. (PSYC)
- Newell, A. and Simon, Herbert A. (1972), Human Problem Solving, Prentice-Hall, Englewood Cliffs, NJ, 1972. (PSYC)
- Newell, A. and Tonge, F. M. (1960), "An Introduction to Information Processing Language V," Comm. ACM, Vol. 3, 208-211. (SYS)
- Nilsson, N. J. (1969a), "Searching Problem-Solving and Game-Playing Trees for Minimal Cost Solutions," Information Processing 68, Vol. 2, A. J. H. Morrel, ed., 1556-1562, North-Holland Publishing Company, Amsterdam, 1969. (SEARCH)
- Nilsson, N. J. (1969b), "A Mobile Automaton: An Application of Artificial Intelligence Techniques," Proc. IJCAL, 509-515, May 1969. (ROB)
- Nilsson, N. J. (1971), Problem-Solving Methods in Artificial Intelligence, McGraw-Hill Book Company, New York, 1971. (SEARCH-G, TP-G)
- Norton, L. (1966), "ADEPT--A Heuristic Program for Proving Theorems of Group Theory," MAC Tech. Report 33, thesis, Project MAC, MIT, Cambridge, MA, 1966. (TP)
- Papert, S. (1968), "The Artificial Intelligence of Hubert L. Dreyfus, A Budget of Fallacies," MIT Artificial Intelligence Memo No. 54, Jan. 1968. (GEN)
- Papert, Seymour A. (1972), "Teaching Children Thinking," Programmed Learning and Educational Technology, Vol. 9, No. 5, Sept. 1972. (GEN)
- Paul, R. C. (1971), "Trajectory Control of a Computer Arm," Proc. 2d Intl. Joint Conf. on Artificial Intelligence, London, England, Sept. 1971. (ROB)
- Paul, R. C. (1973), "Modeling, Trajectory Calculation and Servoing of a Computer Controlled Arm," Ph.D. dissertation, Comp. Sci. Dept., Stanford Univ., Stanford, CA, 1973. (ROB)
- Paxton, William H. and Robinson, Ann E. (1973), "A Parser for a Speech Understanding System," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (LANG)
- Petrick, S. R. (1965), "A Recognition Procedure for Transformational Grammars," Ph.D. thesis, MIT, Cambridge, MA. (LANG)
- Petrick, S. R. (1971), "Transformational Analysis," Natural Language Processing, Courant Computer Science Symposium 8, 20-21 Dec. 1971, R. Rustin, ed., 27-41, Algorithmics Press, New York. (LANG)
- Petrick, S. R. (1973), "Semantic Interpretation in the Request System," Proc. Intl. Conf. on Computational Linguistics, Pisa, Italy, 1973. (LANG)

- Piotrzykowski, T. and Jensen, D. (1972), "A Complete Mechanization of Omega-Order Typo. Theory," Proc. ACM Natl. Conf., Vol. 1, 82-92, 1972. (TP)
- Pingle, K. K. (1969), "Visual Perception by a Computer," Automatic Interpretation and Classification of Images, A. Grasselli, ed., 277-284, Academic Press, New York, London, 1969. (VIS)
- Pingle, K. K., Singer, J., and Wichman, W. (1968), "Computer Control of a Mechanical Arm Through Visual Input," Proc. IFIP Cong. 1968, Vol. 2. (ROB)
- Pinglo, K. K. and Tenenbaum, J. M. (1971), "An Accommodating Edge Follower," IJCAI-2, 1971. (VIS)
- Plath, Warren, J. (1973), "Transformational Grammar and Transformational Parsing in the Request System," Proc. Intl. Joint Conf. on Computational Linguistics, Pisa, Italy, 1973. (LANG)
- Pohl, I. (1970), "Heuristic Search Viewed as Path Finding in a Graph," Artificial Intelligence, Vol. 1, 193-204, 1970. (SEARCH)
- Prawitz, D. (1960), "An Improved Proof Procedure," Theoria, Vol. 26, 102-139, 1960. (TP)
- Quillian, M. R. (1968), "Semantic Memory," Semantic Information Processing, MIT Press, Cambridge, MA. (REP, PSYC, LANG)
- Quillian, Ross (1969), "The Teachable Language Comprehender: A Simulation Program and Theory of Language," Comm. ACM, Vol. 12, 459-476, 1969. (REP, PSYC, LANG)
- Raphael, B. (1964a), "A Computer Program Which Understands," Proc. AFIPS Fall Joint Comp. Conf., 577-589, 1964. (DED, LANG)
- Raphael, B. (1964b), "SIR: A Computer Program for Semantic Information Retrieval," MIT, June 1964. Reprinted in Semantic Information Processing, M. Minsky, ed., MIT Press, Cambridge, MA, 1968. (DED, LANG)
- Raphael, B. (1968), "Programming A Robot," Proc. IFIP Cong. 68, H135-H140, Edinburgh, 1968. (ROB)
- Raphael, B. et al. (1971), "Research and Applications--Artificial Intelligence," Stanford Research Institute Final Report, Contract NASW-2164, National Aeronautics and Space Administration, Doc. 1971. (ROB)
- Raphael, B. and Green, C. (1968), "The Use of Theorem Proving Techniques in Question Answering Systems," J. ACM, 169, 1968. (DED, REP)
- Reboh, R. and Sacerdoti, E. (1973), "A Preliminary QLISP Manual," SRI Artificial Intelligence Center Tech. Note 81, Stanford Research Institute, Menlo Park, CA, Aug. 1973. (SYS)
- Reddy, D. R. (1967), "Computer Recognition of Connected Speech," J. ASA, Vol. 42, 329-347. (LANG)
- Reddy, D. R. et al. (1973), "The Hearsay Speech Understanding System: An Example of the Recognition Process," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (LANG)
- Roberts, L. G. (1963), "Machine Perception of Three-Dimensional Solids," MIT Lincoln Lab., Lexington, MA, Tech. Report No. 315, May 1963. Also in Optical and Electro-Optical Information Processing, J. T. Tippett et al., eds., MIT Press, Cambridge, MA, 1965. (VIS)
- Robinson, J. A. (1965), "A Machine-Oriented Logic Based on the Resolution Principle," J. ACM, Vol. 12, No. 1, 23-41, Jan. 1965. (DED, TP)
- Robinson, J. A. (1969), "A Note on Mechanizing Higher-Order Logic," Machine Intelligence, Vol. 5, B. Meltzer and D. Michie, eds., 123-133, Edinburgh Univ. Press, Edinburgh, 1969. (TP)
- Roan, C. A. (1972), "Robots, Productivity, and Quality," Proc. ACM Natl. Conf. 1972, Boston, MA, Aug. 14-16, 1972. (ROB-G)
- Rosen, C. (1973), "Exploratory Research in Advanced Automation," SRI Report to National Science Foundation, Grant GI-38200X, Stanford Research Institute, Menlo Park, CA, Dec. 1973. (ROB)
- Rozsak, T. (1972), Where the Wasteland Ends: Politics and Transcendence in Post-Industrial Society, Doubleday, 1972. (GEN)
- Rulifson, J. F., Waldinger, R. J., and Derksen, J. A. (1968), "QA4, A Language for Writing Problem-Solving Programs," Proc. IFIP, TA-2, 111-115, 1968. (REP, SYS)
- Rulifson, J. F., Derksen, J. A., and Waldinger, R. J. (1972), "QA4: A Procedural Calculus for Intuitive Reasoning," Tech. Note 73, Artificial Intelligence Center, Stanford Research Institute, Menlo Park, CA, 1972. (REP, SYS, DED)
- Rumelhart, D. E., Lindsay, P. H., and Norman, D. A. (1972), "A Process Model for Long-Term Memory," Organization and Memory, E. Tulving and W. Donaldson, eds., Academic Press, New York, 1972. (PSYC)
- Rumelhart, David E. and Norman, Donald A. (1973), "Active Semantic Networks as a Model of Human Memory," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (PSYC)
- Russell, R. (1964), "KALAH--The Game and the Program," Stanford Univ. Artificial Intelligence Project Memo No. 22, 3 Sept. 1964. (GAME)
- Rustin, R., ed. (1973), Natural Language Processing, Algorithmic Press, New York, 1973. (LANG-G)
- Ryder, J. L. (1971), "Heuristic Analysis of Large Trees as Generated in the Game of GO," AI Memo 155, Artificial Intelligence Project, Stanford Univ., Stanford, CA, 1971. (GAME)
- Sacerdoti, Earl D. (1973), "Planning in a Hierarchy of Abstraction Spaces," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. To appear in Artificial Intelligence, 1974. (DED)
- Sakai, T., Nagao, M., and Knobe, T. (1972), "Computer Analysis and Classification of Photographs of Human Faces," First USA-Japan Comp. Conf. Proc., Oct. 1972. (VIS)
- Samuel, A. (1959), "Some Studies in Machine Learning Using the Game of CHECKERS," IBM J. Research Develop., Vol. 3, 211-229, 1959. Reprinted in Computers and Thought, E. Feigenbaum and J. Feldman, eds., 71-105, McGraw-Hill Book Company, New York, 1963. (GAME)
- Samuel, A. L. (1967), "Some Studies in Machine Learning Using the Game of CHECKERS II--Recent Progress," IBM J. Res. Dev., Vol. 11, 601-617, 1967. (GAME)
- Sandewall, E. J. (1971), "Representing Natural Language Information in Predicate Calculus," Machine Intelligence, Vol. 6, B. Meltzer and D. Michie, eds., 255-277, Edinburgh Univ. Press, Edinburgh. (REP)
- Sandewall, E. J. (1972a), "Formal Methods in the Design of Question-Answering System," J. Artificial Intelligence, 237, 1972. (REP)
- Sandewall, E. J. (1972b), "PCF-2, A First-Order Calculus for Expressing Conceptual Information," Computer Science Report, Comp. Sci. Dept., Uppsala Univ., Uppsala, Sweden. (REP)
- Schank, R. C. (1972), "Conceptual Dependency: A Theory of Natural Language Understanding," Cognitive Psychology, Vol. 3, 552-631, 1972. (PSYCH, LANG)
- Schank, R. C. (1973), "The Fourteen Primitive Actions and Their Inferences," Stanford AIM-183, Comp. Sci. Dept., Stanford Univ., Stanford, CA, 1973. (REP, LANG, PSYC)

- Schank, R. C. et al. (1972), "Primitive Concepts Underlying Verbs of Thought," Stanford AIM-162, Comp. Sci. Dept., Stanford Univ., Stanford, CA, 1972. (REP, LANG, PSYC)
- Schank, R. C. and Colby, K. (1973), Computer Models of Thought and Language, W. H. Freeman & Co., San Francisco, 1973. (LANG-G, PSYC-G)
- Schank, R. C., Goldman, Neil, Rieger, Charles J., and Riesbeck, Chris (1973), "MARGIE: Memory, Analysis, Response Generation, and Inference on English," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (LANG)
- Schank, R. C., Tosler, L., and Weber, S. (1970), "SPIROZA 11: Conceptual Case-Based Natural Language Analysis," Memo AIM-109, Stanford Artificial Intelligence Project, Stanford Univ., Stanford, CA, 1970. (LANG)
- Scheinman, V. D. (1969), "Design of a Computer-Controlled Manipulator," thesis, Dept. of M.E., Stanford Univ., Stanford, CA. Available as Stanford AIM-92, June 1969. (ROB)
- Shannon, C. (1950), "Programming a Digital Computer for Playing Chess," Philosophy Magazine, Vol. 41, 356-375, Mar. 1950. Reprinted in The World of Mathematics, Vol. 4, J. R. Newman, ed., Simon and Schuster, New York, 1954. (GAME)
- Shirai, Y. (1972), "A Heterarchical Program for Recognition of Polyhedra," Memo No. 263, Artificial Intelligence Laboratory, MIT, Cambridge, MA. (VIS)
- Shirai, Y. (1973), "A Context Sensitive Line Finder for Recognition of Polyhedra," Artificial Intelligence, Vol. 4, No. 2, 95-119, Summer 1973. (VIS)
- Shortliffe, E. H. et al. (1973), "An Artificial Intelligence Program to Advise Physicians Regarding Antimicrobial Therapy," Computers and Biomedical Research, Vol. 6, 544-560, 1973. (AIDS)
- Simmons, R. F. (1965), "Answering English Questions by Computer: A Survey," Comm. ACM, Vol. 8, 53-70, 1965. (LANG-G)
- Simmons, R. F. (1969), "Natural Language Question-Answering Systems: 1969," Comm. ACM, Vol. 13, 15-30, 1970. (LANG-G)
- Simmons, Robert F. (1973), "Semantic Networks: Their Computation and Use for Understanding English Sentences," Computer Models of Thought and Language, K. M. Colby and Roger Schank, eds., W. H. Freeman and Company, San Francisco, CA, 1973. (LANG)
- Simmons, Robert F., Klein, S., and McConlogue, D. (1964), "Indexing and Dependency Logic for Answering English Questions," American Documentation, Vol. 15, No. 3, 196-204, 1964. (LANG)
- Simon, H. A. (1963), "Experiments with a Heuristic Compiler," J. ACM, Vol. 10, No. 4, Oct. 1963. (PROG)
- Simon, H. A. (1969), The Sciences of the Artificial, MIT Press, Cambridge, MA. (GEN)
- Simon, H. A. (1972), "The Heuristic Compiler," Representation and Meaning, H. A. Simon and L. Siklossy, eds., Prentice-Hall, Englewood Cliffs, NJ, 1972. (PROG)
- Simon, H. A. and Feigenbaum, E. A. (1964), "An Information-Processing Theory of Some Effects of Similarity, Familiarization, and Meaningfulness in Verbal Learning," J. Verbal Learning and Verbal Behavior, Vol. 3, 385-396. (PSYC)
- Slagle, J. (1961), "A Computer Program for Solving Problems in Freshman Calculus (SAINT)," Lincoln Laboratory Report 5G-001, May 1961. (SEARCH, AIDS)
- Slagle, J. R. (1963), "A Heuristic Program that Solves Symbolic Integration Problems in Freshman Calculus," J. ACM, Vol. 10, No. 4, 507-520, Oct. 1963. Also in Computers and Thought, E. Feigenbaum and J. Feldman, eds., 191-203, McGraw-Hill Book Company, New York, 1963. (AIDS, SEARCH)
- Slagle, J. (1965), "Experiments with a Deductive Question-Answering Program," Comm. ACM, Vol. 8, 792-798, Dec. 1965. (DED)
- Slagle, J. R. (1967), "Automatic Theorem Proving with Renameable and Semantic Resolution," J. ACM, Vol. 14, 687-697, 1967. (TP)
- Slagle, J. R. (1971), Artificial Intelligence: The Heuristic Programming Approach, McGraw-Hill Book Company, New York, 1971. (GEN, GAME-G, SEARCH-G)
- Slagle, J. R. (1970), "Heuristic Search Programs," Theoretical Approaches to Non-Numerical Problem Solving, R. Bannerji and M. I. Mesarovic, eds., 246-273, Springer-Verlag, New York, 1970. (SEARCH)
- Slagle, J. R. and Dixon, J. (1969), "Experiments with Some Programs that Search Game Trees," J. ACM, Vol. 16, No. 2, 189-207, Apr. 1969. (GAME, SEARCH)
- Solomonoff, R. (1966), "Some Recent Work in Artificial Intelligence," Proc. IEEE, Vol. 54, No. 112, Dec. 1966. (GEN)
- Sperling, G. (1960), "The Information Available in Brief Visual Presentations," Psychological Monographs, Vol. 74. (PSYC)
- Sridharan, N. S. (1971), "An Application of Artificial Intelligence to Organic Chemical Synthesis," thesis, State Univ. of New York at Stony Brook, New York, July 1971. (AIDS)
- Sridharan, N. S. et al. (1973a), "A Heuristic Program to Discover Syntheses for Complex Organic Molecules," Stanford Artificial Intelligence Memo 205, Stanford Univ., Stanford, CA, June 1973. (AIDS)
- Sridharan, N. S. (1973b), "Search Strategies for the Task of Organic Chemical Synthesis," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (AIDS)
- Sternberg, S. (1966), "High Speed Scanning in Human Memory," Science, Vol. 153, 652-654. (PSYC)
- Sussman, G. J. (1973), "A Computational Model of Skill Acquisition," Tech. Note AI TR-297, Artificial Intelligence Laboratory, MIT, Cambridge, MA, Aug. 1973. (DED, PROG)
- Sussman, G. J. and McDermott, D. V. (1972), "From PLANNER to CONNIVER--A Genetic Approach," Proc. AFIPS FJCC, Vol. 41, 1171-1180, 1972. (SYS)
- Teitelman, W. (1969), "Toward a Programming Laboratory," Proc. 1st Intl. Joint Conf. on Artificial Intelligence, Washington, D.C., 1969. (PROG)
- Teitelman, W. (1972a), "Do What I Mean," Computers and Automation, Apr. 1971. (PROG, SYS)
- Teitelman, W. (1972b), "Automated Programming--The Programmer's Assistant," Proc. Fall Joint Comp. Conf., Dec. 1972. (PROG, SYS)
- Teitelman, W. (1973), "CLISP--Conversational LISP," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (PROG, SYS)
- Teitelman, W. (1974), INTERLISP Reference Manual, Xerox and Bolt, Beranek and Newman. Copies available from Xerox Palo Alto Research Center, Palo Alto, CA. (SYS)
- Tenenbaum, J. M. (1973), "On Locating Objects by Their Distinguishing Features in Multisensory Images," SRI Artificial Intelligence Center Tech. Note 84, Stanford Research Institute, Menlo Park, CA, Sept. 1973. To appear in Computer Graphics and Image Processing, 1974. (VIS)
- Tenenbaum, J. M. et al. (1974), "An Interactive Facility for Scene Analysis Research," SRI Artificial Intelligence Center Tech. Note. 87, Stanford Research Institute, Menlo Park, CA, 1974. (VIS)

- Thompson, F. D. (1966), "English for the Computer," Proc. AFIPS 1966 Fall Joint Comp. Conf., Vol. 29, 349-356, Spartan Books, New York. (LANG)
- Thompson, F. D. et al. (1969), "REL: A Rapidly Extensible Language System," Proc. 24th Natl. ACM Conf., 1969. (LANG)
- Thorne, J., Bratley, P., and Dewar, H. (1968), "The Syntactic Analysis of English by Machine," Machine Intelligence, Vol. 3, D. Michie, ed., American Elsevier Publishing Company, New York, 1968. (LANG)
- Tinbergen, N. (1951), The Study of Instinct, Clarendon Press, Oxford, 1951. (PSYC)
- Tonge, F. (1965), "Assembly Line Balancing Using Probabilistic Combinations of Heuristics," Management Sci., Vol. 11, No. 7, 727-735, May 1965. (AIDS)
- Turing, A. M. (1949), "Checking a Large Routine," Report of a Conference on High Speed Automatic Calculating-Machines, McLennan Laboratory, Univ. of Toronto, Canada. (PROG)
- Turing, A. M. (1950), "Computing Machinery and Intelligence," Mind, Vol. 59, 433-460, Oct. 1950. Reprinted in Computers and Thought, 11-35, E. Feigenbaum and J. Feldman, eds., McGraw-Hill Book Company, New York, 1963. (GEN)
- Vicens, P. (1969), "Aspects of Speech Recognition by Computer," Report CS-127, Ph.D. thesis, Comp. Sci. Dept., Stanford Univ., Stanford, CA. (LANG)
- Waldinger, R. J. and Lee, R. C. T. (1969), "PROW: A Step Toward Automatic Program Writing," Proc. Intl. Joint Conf. on Artificial Intelligence, 241-252, 1969. (PROG)
- Waldinger, R. J. and Levitt, K. N. (1973), "Reasoning About Programs," Tech. Note 86, SRI Artificial Intelligence Center, Oct. 1973, Stanford Research Institute, Menlo Park, CA. To appear in Artificial Intelligence, 1974. (PROG)
- Walker, Donald E., ed. (1964), English Preprocessor Manual, The Mitre Corporation, Bedford, MA, 1964 (SR-132). (LANG)
- Walker, Donald E. (1973), "Speech Understanding, Computational Linguistics, and Artificial Intelligence," Tech. Note 85, SRI Artificial Intelligence Center, Stanford Research Institute, Menlo Park, CA, Aug. 1973. (LANG)
- Waltz, D. G. (1972), "Generating Semantic Descriptions from Drawings of Scenes with Shadows," AI TR-271, Artificial Intelligence Laboratory, MIT, Aug. 1972. (VIS)
- Waterman, D. A. (1970), "Generalization Learning Techniques for Automating the Learning of Heuristics," Artificial Intelligence, Vol. 1, Nos. 1 and 2, Spring 1970. (GAME)
- Waterman, D. A. and Newell, A. (1971), "Protocol Analysis as a Task for Artificial Intelligence," Artificial Intelligence, Vol. 2, Nos. 2 and 3, 285-318, 1971. (AIDS)
- Waterman, D. A. and Newell, A. (1973), "PAS-II: An Interactive Task-Free Version of an Automatic Protocol Analysis System," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (AIDS)
- Wegbreit, Ben (1973), "Heuristic Methods for Mechanically Deriving Inductive-Assertions," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (PROG)
- Weissman, C. (1967), LISP 1.5 Primer, Dickenson Press, 1967. (SYS)
- Weizenbaum, J. (1966), "ELIZA--A Computer Program for the Study of Natural Language Communication Between Man and Machine," Comm. ACM, Vol. 9, 36-45, 1966. (LANG)
- Weizenbaum, J. (1972), "On the Impact of the Computer on Society," Science, Vol. 176, No. 609, 1972. (GEN)
- West, J. D. (1967), "A Heuristic Model for Scheduling Large Projects with Limited Resources," Management Science, Vol. 13B, 359-377. (AIDS)
- Winograd, T. (1971), "Procedures as Representation for Data in a Computer Program for Understanding Natural Language," Tech. Report AI TR-17, MIT, Cambridge, MA, 1971. Published as Understanding Natural Language, Academic Press, New York, 1972. (REP, DED, LANG)
- Winston, P. H. (1970), "Learning Structural Descriptions from Examples," Tech. Report AI TR-231, Artificial Intelligence Laboratory, MIT, Cambridge, MA, 1970. (REP, VIS)
- Winston, P. H. (1972), "The MIT Robot," Machine Intelligence, Vol. 7, 431-463, B. Meltzer and D. Michie, eds., American Elsevier Publishing Company, 1972. (ROB, VIS)
- Woods, W. A. (1970), "Transition Network Grammars for Natural Language Analysis," Comm. ACM, Vol. 13, 591-606, 1970. (LANG)
- Woods, W. A. (1973), "An Experimental Parsing System for Transition Network Grammars," Natural Language Processing, R. Rustin, ed., 111-154, Algorithmics Press, New York, 1973. (LANG)
- Woods, W. A., Kaplan, R. N., Nash-Webber, G. (1972), "The Lunar Science Natural Language Information System: Final Report," BBN Report No. 2378, Bolt, Beranek and Newman, Inc., Cambridge, MA, June 1972. (LANG)
- Woods, W. A. and Makhoul, J. (1973), "Mechanical Inference Problems in Continuous Speech Understanding," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (LANG)
- Woodriddle, D. (1963), "An Algebraic Simplify Program in LISP," Artificial Intelligence Project Memo No. 11, Stanford Univ., Stanford, CA, Dec. 1963. (AIDS)
- Wos, L. T., Carson, D. G., and Robinson, G. A. (1964), "The Unit Preference Strategy in Theorem Proving," Proc. AFIPS, Fall 1964, Vol. 25, 615-621, Spartan Books, Washington, D.C. (TP)
- Wos, L. T., Robinson, G., and Carson, D. (1965), "Efficiency and Completeness of the Set of Support Strategy in Theorem-Proving," J. ACM, Vol. 12, No. 4, 536-541, Oct. 1965. (TP)
- Yakimovsky, Yoram and Feldman, Jerome A. (1973), "A Semantics-Based Decision Theory Region Analyzer," Adv. Papers 3d Intl. Conf. on Artificial Intelligence, Stanford Univ., Stanford, CA, Aug. 1973. (VIS)
- Yates, R., Raphael, B., and Hart, T. (1970), "Resolution Graphs," Artificial Intelligence, Vol. 1, No. 4, 1970. (TP)
- Zobrist, A. (1969), "A Model of Visual Organization for the Game of GO," Proc. AFIPS, Spring 1969, 103-112. (GAME)
- Zobrist, A. and Carlson, F., Jr. (1973), "An Advice-Taking Chess Computer," Scientific American, June 1973. (GAME)
- Zwicky, Arnold M. et al. (1965), "The Mitre Syntactic Analysis Procedure for Transformational Grammars," Proc. AFIPS, Fall 1965, Vol. 27, 317-326. (LANG)