Massachusetts Institute of Technology
Artificial Intelligence Laboratory

Working Paper 239 November 1982

What to Read

A Biased Guide to AI Literacy for the Beginner

Philip E. Agre

Abstract: This note tries to provide a quick guide to AI literacy for the beginning AI hacker and for the experienced AI hacker or two whose scholarship isn't what it should be. Most will recognize it as the same old list of classic papers, give or take a few that I feel to be under- or over-rated. It is not guaranteed to be thorough or balanced or anything like that.

Acknowledgements. It was Ken Forbus' idea, and he, Howie Shrobe, Dan Weld, and John Batali read various drafts. Dan Huttenlocher and Tom Knight helped with the speech recognition section. The science fiction section was prepared with the aid of my SF/AI editorial board, consisting of Carl Feynman and David Wallace, and of the ArpaNet SF-Lovers community. Even so, all responsibility rests with me.

Introduction

I discovered the hard way that it's a good idea for a beginning AI hacker to set about reading everything. But where to start? The purpose of this note is to provide the beginning AI hacker (and the experienced AI hacker or two whose scholarship isn't what it should be) with some starting places that I've found useful in my own reading. If you already know what's going on, you'll recognize the same old list of classic papers plus or minus a few things that I consider to be under- or over-rated. If most of this is new to you, you've got a lot of bedtime reading ahead.

This document is to be taken lightly. It is heavily biased in various ways. My own interests and experience are partly at fault. Some areas of AI are better served by survey papers than others, and in some cases I will give a reference to a survey rather than do one myself. There is a bias toward newer work, which I rationalize on the grounds that it is easier to work backward through the references than forward. No doubt you will also notice an extreme bias toward MIT ideas and outlooks, to the detriment of those of other institutions where perfectly fine work is being done. All of the obvious explanations of this phenomenon are correct.

There is a school of thought that holds that it's better *not* to be well-read, the argument being that studying the mistakes of the past will lead one's mind into a rut. This being a religious matter, I will satisfy myself with the suggestion that it's much preferable to acquire the capacity not to be led and the capacity of learning from those with whom one disagrees.

A few general guidelines are in order for reading the papers I'll mention.. First, keep in mind that AI is a fast-moving field. Researchers change their minds rapidly and subtly and often without mentioning it. Try to get a feel for the development of particular authors' thinking, as well as that of the field generally. There are any number of themes that run through all AI work; I'll mention some of them but you'll have to find most of them yourself. Second, when an author describes a program, read closely to find out if the program really worked. Frequently a program existed but only worked properly on the published examples. Shocking, yes, but true. Third, keep close track of "buzzwords", words that sound like they have precise technical meanings but may not. Buzzwords, though evil, are a necessity, and every AI researcher has a large mental library of them, each associated with the authors who use them. I'll double-quote any number of them later on. Finally, realize that this document is going to be obsolete by the time you read it.

Here's the plan: first I'll make some lists of books, conferences, journals, people, and places you should know about. Then I'll go through the field area by area making suggestions as to what one should read to get a beginner's knowledge of what's going on in each area.

Books

If you have to have a rough outline of AI by dawn, then read Winston's text "Artificial Intelligence" (1977). Boden's "Artificial Intelligence and Natural Man" (1977) is drier and more detailed. Intended for the general reader, it gives more attention to issues of how AI relates to philosophy, psychology, and social issues. Neither should be seen as more than a quick sketch of the field.

There is a standard set of anthologies that are frequently referred to. It's good know where you can find a copy of each. Feigenbaum and Feldman's "Computers and Thought" (1963), Minsky's "Semantic Information Processing" (1968a), Schank and Colby's "Computer Models of Thought and Language" (1973), and Bobrow and Collins' "Representation and Understanding" (1975) are the standard references for most AI work before about 1973. Winston's "The Psychology of Computer Vision" (1975a) describes early vision work. The "Machine Intelligence" anthology series, edited by Michie and various others, has been published roughly annually since 1967. It is oriented toward more formal approaches to AI. The papers in Waterman and Hayes-Roth's "Pattern-Directed Inference Systems" (1978) are unified by the idea of a system made up of a large number of computational elements which are triggered when some condition description matches a model of the world. Winston and Brown's "Artificial Intelligence: An MIT Perspective" (1979) is a two-volume collection of papers from MIT. The "Handbook of Artificial Intelligence" (1981,1982), edited by Barr and Feigenbaum, is a three-volume encyclopedia of AI, with articles surveying the field by area. I find it useful as an enumeration of projects but short on synthesis. (Brady and Berwick 1982 forthcoming) deals with discourse processes. Herbert Simon's collected papers appear in "Models of Thought" (1979). Johnson-Laird and Wason's "Thinking" (1977) is oriented toward cognitive psychology. The 10th anniversary (1981) issue of the journal Cognition is a source of pointers to the latest in cognitive psychology. Haugeland's "Mind Design" (1981) is philosophically oriented.

Perhaps this is the place to slip in the observation that a remarkably large part of the important work in AI has appeared in PhD theses. Among the most important are those of Winston (1970), Winograd (1971), Charniak (1972), Sussman (1973), Rieger (1974), Sacerdoti (1977), Brachman (1978b), de Kleer (1979), Doyle (1980), Clinger (1981), Brooks (1981a), and Smith (1982).

Conferences

The proceedings of the AAAI (American Association for AI) and IJCAI (International Joint Conference on AI), issued in alternating years, provide a sort of yellow pages of AI research. Everything's there, but the presentations are highly compressed and the quality is, um, uneven. TINLAP (Theoretical Issues in Natural Language Processing) covers various approaches to natural language. When the alternative is having you read a thesis of several hundred pages, I'll frequently also provide a reference to a short description of the work in a conference proceedings. These only provide sketches, though, and you should have a look at the originals eventually.

Journals

Journals do not serve as communications media or accreditors of good work in AI nearly so much as in most academic fields. Reasons for this include the relative youth and extreme wealth of the field, the speed with which it's moving, the fact that most work is done in a few large laboratories, and the availability of relatively high-bandwidth data communication facilities such as the Arpanet. Consequently, journals in AI are to be treated as archives; the way to find out what's happening this month is to read the internal reports of the various research centers as they come out.

In any event, the main journal of AI is "Artificial Intelligence". Have a look through it to see what Establishment research looks like. Likewise "Cognitive Science", a newer journal aimed at a synthesis of AI and its neighbors. The "International Journal of Robotics Research" is brand-new and covers technical aspects of robotics. The "AI Magazine", put out by AAAI, and the "Sigart Newsletter", put out by the ACM's Special Interest Group on AI, contain abstracts, conference announcements, workshop reports, outlines of AI work at various sites, bibliographies, new addresses, and various sorts of unrefereed papers, and are reasonable ways to keep up. "ACM Computing Surveys" has review articles on various topics in Computer Science. "Cognition" is a fine journal of cognitive psychology. Especially recommended is "Behavioral and Brain Sciences", which covers as wide an area as its title suggests. Its salient feature is its "open peer commentary": each article is accompanied by shorter commentaries by various people and a response by the author. The quality is uneven, but there's no better way to get an idea of the current issues in a broad range of relevant fields. The "Computer Music Journal" covers all aspects of computer music, from acoustics to composition to AI models of music skills.

People

Here are some lists of names you should know. I would very much appreciate it if nobody got mad at me about either this list or the one about Places. They are emphatically not intended to be complete, and the categories are very rough indeed. Note that the names in this list are in alphabetical order.

Founders of the field: McCarthy, Minsky, Newell, Simon

Representation: Bobrow, Brachman, Bundy, Collins, Forbus, Hayes, Israel, Lehnert, Martin, Quillian, Rieger, Schank, Stevens, Szolovits, Woods

Vision: Ballard, Binford, Feldman, Frisby and Mayhew, Grimson, Hildreth, Horn, Hubel and Wiesel, Julesz, Marr, Poggio, Richards, Rosenfeld, Tenenbaum, Ullman, Waltz

Robotics: Albus, Bizzi, Brady, Hollerbach, Lozano-Perez, Luh, Mason, Moravec, Paul, Popplestone, Raibert, Scheinman, Shimano, Taylor, Whitney

Problem-solving and learning: Andreae, Angluin, Berliner, Carbonell, de Kleer, Doyle, Fahlman, Fikes, Hart, Hayes-Roth, Langley, Lenat, D McDermott, Nilsson, Robinson, Sacerdoti, Stefik, Sussman, Weyhrauch, Winston

Natural language and speech recognition: Allen, Berwick, Erman, Grosz, Klatt, Lesser, Lyon, Marcus, McDonald, Reddy, Riesbeck, Sidner, Webber, Wilensky, Wilks, Winograd, Zue

Expert systems and design aids: Anderson, Balzer, Buchanan, Davis, Feigenbaum, Genesereth, Kant, J McDermott, Rich, Shrobe, Waters

Education: H Abelson, Kay, O'Shea, Papert, Young

Psychology: R Abelson, Carey, Chomsky, Gentner, Neisser, Norman, Osherson

Philosophy: Boden, Dennett, Dreyfus, Fodor, Haugeland, Putnam, Pylyshyn, Searle, Sloman, Smith, Weizenbaum

Places

Here is a list of some places where AI is being done these days. Most of these places have internal research report series, and it's generally worthwhile trying to keep up with these. This list isn't meant to be exhaustive either.

American universities: MIT, CMU, Stanford, Yale, UMass-Amherst, Rochester, Illinois, Pennsylvania, Rutgers, SUNY-Buffalo, Berkeley, Maryland, Texas

More or less private research institutions: Bolt Beranek and Newman (BBN, Cambridge MA), Information Sciences Institute (ISI, affiliated with USC, Marina del Ray CA), Stanford Research Institute (SRI, Menlo Park CA)

Industry: Schlumberger, Fairchild (Palo Alto CA), Xerox (Palo Alto CA), Atari (Sunnyvale CA and Cambridge MA), DEC (Tewksbury MA)

Abroad: Edinburgh, Sussex, Centre pour les Etudes Semantiques et Cognitives (Geneva), McGill (Quebec, Canada), Canterbury (Christchurch, NZ), Kyoto University and various other Japanese laboratories

Suggestions by area

For convenience in exposition, I have broken AI down into a number of areas. Needless to say, this is an unfortunate thing to have to do, and you shouldn't take the divisions very seriously. Not that there aren't divisions in AI, but these aren't they.

AI programming

Learn to program in Lisp. The way to do this is to read the course notes for Sussman and Abelson's 6.001 course, which will teach you a dialect of Lisp known as Scheme. Find an implementation of Scheme (they exist for 20's and Vaxes) and run their examples for yourself. (See (Stacy 1982) for an introduction to the computing facilities at the MIT AI Lab.) The transition from Scheme to other dialects of Lisp, such as MacLisp, Lisp Machine Lisp, and InterLisp, will be relatively painless. There are other good Lisp texts, such as (Winston and Horn 1981), and (Friedman 1974), but the 6.001 notes are better. See (Winston 1977) for some more AI-oriented Lisp lessons. There are two texts for advanced AI programming techniques, (Charniak 1980) and (Schank and Riesbeck 1981).

There is a long tradition of AI programming languages. The main lines of this tradition can be traced back to McCarthy's invention of Lisp (see McCarthy 1978 for the early history of Lisp), one of whose contributions was to take one routine aspect of AI programming, storage management, and put it in the background, much as other AI programming languages would later attempt to do with other common activities. Most subsequent AI languages have been implemented in some version of Lisp. There is a large Lisp implementation literature, but this will only be of peripheral interest to those studying AI per se. See, for example, (Moses 1970), (Steele 1978), and (Clark and Green 1977).

Planner (Hewitt 1971,1972) was an early attempt to incorporate a notion of AI control structure ("chronological backtracking") into a programming language in the much same way as Lisp incorporated storage management. Conniver (McDermott and Sussman 1974) was an attempt to address Planner's shortcomings. (Sussman and McDermott 1972) contains an excellent analysis of these shortcomings and Conniver's answers to them; see if you can get ahold of a copy of the first, unexpurgated, edition. See (Fahlman 1973,1974) for a description of a large program written in Conniver. Actors (Hewitt 1977) are an attempt to reformalize computation in terms of active "objects" that communicate by passing messages. There's also Amord (de Kleer et al 1977,1978), which attempted to incorporate dependency maintenance (of which much later) into a rule-based language (cf the discussion of production systems below). (Shrobe 1979a,b) describes a large program written in Amord.

Prehistory

(Minsky 1961) is a good survey of AI's prehistory. Also read Minsky and Papert's book "Perceptrons" (1972). This is a beautiful piece of applied mathematics that exhaustively analyzes one sort of small computing device, the linear perceptron. This device was the source of much unfocussed theoretical interest during the sixties. As with many of the great works in AI (and in science generally), the most important aspect of the Perceptrons book is its view of the methodology one should bring to bear in the analysis of computational elements, a view which greatly improved on the one it replaced. Unfortunately, Minsky and Papert solved all the easy problems in perceptron analysis and so managed to kill off the field of analytical investigation of small computing devices, quite contrary to their intentions. The introduction of large computers and the consequent ability of researchers to experiment with large computing devices may also have had something to do with it. Two more classics from the days of AI's origins are Wiener's "Cybernetics" (1961) and Simon's "The Sciences of the Artificial" (1970). Both are inspiring works describing many fundamental ideas of the field.

Knowledge representation

Much effort has gone into trying to build a general-purpose scheme for representing knowledge. This began with the earliest "semantic networks" (Quillian 1968). Schank and his school have developed a representation scheme based on "semantic primitives" known as "conceptual dependency". Look at (Schank 1973b) and (Schank 1980) and try to understand the evolution of the philosophy behind conceptual dependency. (Rieger 1974, Schank and Rieger 1974) describe an early model of memory and language comprehension based on conceptual dependency. (Woods 1975) began the systematic philosophical analysis of representation constructs. This line led through (Brachman 1978a,b) to the KL-ONE project at BBN (Brachman 1978c). The FRL representation language (Roberts and Goldstein 1977) came of an attempt to formalize in a representation language Minsky's idea of a "frame" (Minsky·1974). Try to understand what got lost in the translation and then read (Minsky 1982 forthcoming). Three central and related issues in knowledge representation are the nature of the connection between representations of apples and apples themselves, the role of symbolic logic in representations, and the nature of the relationship between language and knowledge representation. Try especially to understand the positions of Minsky and the creators of KRL (Bobrow and Winograd 1976) on these issues. For an unusually sophisticated defense of logic in knowledge representation see (Hayes 1977). Another defender of logic is McCarthy (see, e.g., 1977). See (Moore 1977,1979) for an interesting and elaborate attack on the implementational problems of reasoning with modal logic. (See also the discussion of non-monotonic logic below.) For another view, see (Israel and Brachman 1981). The February 1980 Sigart Newsletter is a very useful "Special issue on knowledge representation" edited by Ron Brachman and Brian Smith.

Representation of physical processes

A sub-tradition in knowledge representation especially interesting for the degree of progress it seems to be making is the attempt to build representations of physical mechanisms and processes. Start with (de Kleer and Brown 1982), which is an excellent exposition of some of the basic considerations. To get an idea of the historical development of this area start with Rieger's "flush toilet" paper (1976), noting the problems in his representation that de Kleer and Brown address in theirs. Proceed to Hayes' "ontology for liquids" (1979b; see also 1979a), where begins the theme of explicit representation of histories. Then read deKleer's PhD thesis (1979) (a fine thesis; it really did work on hundreds of examples). The latest word is Forbus' ongoing Qualitative Process Theory project (1982). Also have a look at Allen's (1981) ideas about reasoning about time.

Problem solving

Let's call the work reported in Newell and Simon's remarkable book "Human Problem Solving" (1972) the beginning of the tradition of the analysis of human problem solving (a.k.a. "planning") behavior and its simulation by computer. Read especially the final chapter, which presents a theory of the mind based on production systems (of which more later) and a model called GPS (General Problem Solver). (Newell and Simon 1963, Newell Shaw and Simon 1960) describe GPS and many other important early ideas in detail. See (Minsky 1982 forthcoming) for more about GPS-like algorithms. GPS was an instance of an early approach to problem solving based on so-called "weak methods". Nilsson's "Problem Solving Methods in Artificial Intelligence" (1971) is an early text centering on heuristic search and other weak methods. (Berliner 1981) discusses the limits of the weak methods.

The next important chapter in problem solving research was Strips (Fikes and Nilsson 1971). Compare the Strips notion of "difference" to the analogous notion in GPS. (Fikes Hart and Nilsson 1972) extended Strips to allow it to learn generalized forms of the plans it constructed. (Sacerdoti 1974) made another extension to cut down the search space of plans by planning in a sequence of abstract plan spaces. (Compare this algorithm to the stereo matching algorithm reported in (Marr 1982) and (Grimson 1981a).) Sacerdoti's thesis (1975,1977) viewed plans as partially ordered structures and the planning process as a hierarchical expansion and refinement of such plans.

A number of important problem-solving ideas have come from MIT in recent years. Start with Sussman's thesis (1973), which described a theory of problem solving and learning based on the construction of new plans by debugging plans built from plans already in a library of previously constructed and debugged plans. McDermott's thesis (1977a,b) describes a problem solving program with a representation of its own state. (Stallman and Sussman 1977) began a series of works developing the idea of "dependency-directed backtracking", an improvement on Planner's chronological backtracking. (Indeed, notice that the phrase "dependency-directed backtracking" can now easily be read between the lines of (Sussman and McDermott 1972), which predated the idea by several years.) TMS (Doyle 1978) was an attempt to provide dependency-based consistency maintenance as a utility for the programmer. Non-monotonic logic arose as an attempt to formalize the mechanism of TMS and the formal notions of assumption and consistency that such programs suggested. An issue of "Artificial Intelligence" (Volume 13, numbers 1 and 2, April 1980) was devoted to the issues raised by non-monotonic logic. See (Israel 1980) for another view. Amord (mentioned above) was a programming language that incorporated dependency maintenance. (Doyle and London 1980) is a bibliography of work on "belief revision" ideas such as these. Read (Minsky 1979) and convince yourself that a k-line is a dependency tree.

(Sussman and Steele 1981) and (Borning 1979) describe two closely related implementations of the idea of "constraints" between various quantities in the representation of a structure. (Waltz 1972,1975) and (Brooks 1981a,b) represent variations on quite a different approach, in the context of the analysis of visual images, and (Stefik 1980,1981) represents yet another approach. "Constraint" is a buzzword to watch out for, since it is frequently used imprecisely. It is also used in an unrelated methodological sense by cognitive psychologists, especially at MIT.

One common theme in problem-solving work is reflexivity: allowing a reasoning program to reason about its own workings in the same way as it reasons about the world. The first concrete application of this idea was (Newell Shaw and Simon 1960). See (Doyle 1980) and (Smith 1982) for two excellent but quite different recent approaches to the idea. (Weyhrauch and Thomas 1974) pioneered the mechanics of self-referential representations. Also, compare the notions of self-knowledge in these works to the one in (Minsky 1968b).

Learning

The idea of learning has a long history in AI. There is a long tradition of theoretical approaches to learning, surveyed in (Angluin and Smith 1982 forthcoming). The modern history of AI-oriented learning research begins with the PhD theses of Winston (1970) and Sussman (1973), wherein the notions of near-misses and debugging of almost-right plans were first formalized. Recently much interest has developed in learning and reasoning by analogy (current *flows* in a wire). See (Evans 1968) for the first word and (Winston 1980,1982; cf Minsky 1980) for the last. Also, see (Gentner 1980,1982) for a psychological perspective. (Minsky 1982 forthcoming) is brand-new and great fun. (Mitchell Carbonell and Michalski 1981) collects reports on learning projects at various sites, and (Nudel and Utgoff 1981) is a bibliography of learning work. These collections should not be used as primary resources by the discriminating beginner, though.

The most common idea in AI learning work, though it usually appears implicitly, is that one can take advantage of the fact that many representation schemes are natural lattices under the operations of minimal common generalization (mcg) and maximal common specialization (mcs, also known as unification). (Early appearances of this observation include (Reynolds 1970; see also Plotkin 1970). It also figures in the large literature on "resolution theorem proving", at which you should have a quick look. See, e.g., (Robinson 1967) or (Plaisted 1981). See (Birkhoff 1967) for a good introduction to lattice theory.) The idea in learning is that one can constrain the "concept" one is learning from "below" by forming the mcg of one's "positive" data and from "above" by forming the mcs of one's "negative" data. It is a good practice to look for "lattice climbing" ideas underlying any learning algorithm one comes across in the literature. (Michalski 1980) describes a particularly impressive lattice-theoretic algorithm based on a symbolic generalization of traditional statistical clustering techniques.

Vision

Vision too has a long history in AI, but it is better served than most areas by survey articles and the like, so I won't go into great depth here. Start with Marr's book "Vision" (1982), which describes most of a decade of work at MIT by Marr and his group. (Then read (Grimson 1981b) and (Terzopoulos 1982) for the latest on the surface interpolation problem.) Brady's survey article (1981) covers much of the same ground but has a better coverage of other AI work on vision. Useful pointers into the early vision literature can be found in (Winston 1975a). (Brooks 1981a,b) discusses the boundary between low-level vision and higher-level representation and problem solving. For an especially pretty piece of applied math in vision, see (Longuet-Higgins and Prazdny 1980) on optical flow.

Robotics

The robotics literature can be roughly divided into the more romantic work and the more technical work, the latter being the current wave. The outstanding example of the former school is (Moravec 1980), an amusing account of a mobile robot project. The newer work has been until just recently mostly concerned with the precise control of robot arms, and there is a large literature concerned with the problem of predicting the dynamics of a given arm and using that information to produce control signals to its motors in real time. Theoretically speaking, that effort has essentially been completed with (Hollerbach 1979), which presents a computationally tractable formulation of arm dynamics. There is also a growing literature on compliance and other force-feedback techniques; see, e.g., (Mason 1981). (Paul 1981) is a new book covering most of this ground; the beginning reader will find its unified notational conventions a relief from the chaos of the literature. See (Bizzi and Abend 1982) for a review of some of the relevant physiological literature. Thought is now being given to the control of more-or-less anthropomorphic robot hands; see (Hollerbach 1982) for a sketch of some of the issues and (Salisbury 1982a,b) for the latest word in robot hands.

Speech recognition

The major body of work on speech recognition in AI was done under the five-year ARPA Speech Project. The most interesting of these systems from an AI point of view was Hearsay-II. It incorporated multiple knowledge sources which communicated with one another through a heterarchical "blackboard" structure (Erman et al 1980). There was much speculation on the applicability of this control structure to other problems. Hearsay suffered from fairly poor acoustic and phonetic level processing. The HWIM system at BBN (Woods et al 1976) paid more attention to the the literature on acoustics and psychoacoustics, but failed to incorporate that information into a fully operational system. The highest performance system was Harpy

(Lowerre 1976), a successor to Hearsay. This system precompiled all of its knowledge about acoustics, phonetics, syntax and semantics into one large network, and then employed the "beam search" method to search this space. Although beam search is thought to be important for AI, both network searching and the ideas about speech inherent in Harpy's network are now considered naive as speech recognition methods. A good general review of the whole ARPA five-year project can be found in (Klatt 1977).

Existing systems for recognizing continuous speech are dependent on the use of a small vocabulary (usually less than 1000 words), restricted syntax, and a single speaker. Though many systems can be retrained for different speakers, they are not truly speaker independent. The most important trend in recent years has been in the incorporation of new acoustic phonetic knowledge into computational theories. This is the approach taken by Zue (1980, 1981), where the idea is to take advantage of low level knowledge about the characteristics of (English) speech to develop a speaker independent, large vocabulary speech recognizer. There has also been recent work on modelling the characteristics of the human auditory system. Searle's (1979) auditory modelling, for example, has been motivated by the physiology and psychophysics of hearing. Lyon's (1982) signal-processing algorithms for speech recognition are similarly motivated by the physiological and psychophysical literature. Think about Marr's methodology in vision research while reading the speech literature (even if the authors don't).

Natural language

The literature on mechanized natural language processing becomes more informed by the linguistic literature as time goes on. (Winograd 1971,1973) describes an early system for language understanding, based on the idea of procedural semantics. (Charniak 1972,1975) was an early attempt at story comprehension. The debate in (Dresher and Hornstein 1976, Winograd 1977, Dresher and Hornstein 1977) serves to illustrate some of the issues separating the early AI approach to language from that of Chomskyan linguistics. Much of the parsing literature is centered on the idea of an augmented transition network (ATN); see (Woods 1970) for the origins of ATNs. Other parsing work has been motivated by generative linguistics (for which see Chomsky 1975). Marcus' (1980) remarkably simple parsing mechanism has been extended by Berwick (1981) to acquire grammar from attempts to parse naturally occurring utterances. (Winograd 1983) provides a detailed survey of current parsing research. (McDonald 1982) describes a practical and well-motivated program for generating text from arbitrary source representations. There is a large literature on the processes underlying coherent discourse, for which (Brady and Berwick 1982 forthcoming) is a good starting point.

Expert systems

An expert system is a program meant for practical application that uses a large body of knowledge about a particular field. Expert system technology has been largely based on production systems. Production systems started out as the basis of a model of human problem-solving (Newell and Simon 1972) but have since come to be simply a procedural formalism for expressing knowledge about domains; Nilsson's (1980) text formalizes much of AI within a production system framework. The earliest important expert system was Dendral (Lindsay et al 1981); other famous ones include Mycin (Shortliffe 1976) and R1 (McDermott 1980). One of the main slogans of the expert systems school is "in the knowledge lies the power", and there is much concern for knowledge acquisition by expert systems. (Davis 1979) describes an approach to human interfaces for knowledge acquisition. Meta-Dendral (Buchanan and Mitchell 1978) is a neat addition to Dendral that induces new rules for its data base; its salient characteristic is that it has a fairly deep understanding of its domain (mass spectroscopy) and can use this to analyze candidate rules. The Programmer's Apprentice project at MIT (Rich and Shrobe 1978, Waters 1982) also attempts to work from a "deep" theory of its domain. Recent expert system work has concentrated on developing more general control structures and "deeper" knowledge representations; see (Stefik et al 1982) for a good review of ways of doing this.

Education

(Papert 1980) describes the MIT Logo project, which sought to teach and understand children by letting them program a computer in a graphical language called Logo. Though the kids clearly loved it, you should read carefully to see exactly what claims are made for Logo and how these claims are argued. See (diSessa 1982) and (Stevens Collins and Goldin 1979) for analyses in computational terms of students' mistakes in physics. One approach to computer-aided learning starts with a computational model of the child's mental processes and uses it to analyze the child's mistakes and, eventually, to fix the "bugs" in the child's reasoning. (Young and O'Shea 1981), (Brown and vanLehn 1979), and (Sleeman and Brown 1981) describe three projects which seek to represent a child's knowledge about the subtraction algorithm as a set of productions in a production system, and to explain the child's subtraction mistakes in terms of bugs in her productions.

Methodology

Now read (Marr 1982) again. Marr's contribution to AI was as much in methodology as in the particular domain of vision. He distinguishes computational theory, algorithm, and implementation in the investigation of a human capacity, and castigates other AI researchers for not making these distinctions and for working in areas where it is not yet possible to formulate a computational theory (Marr 1977). Compare Marr's methodology in vision with Chomsky's in language (see, e.g., 1980), and then read (Berwick and Weinberg 1982 forthcoming), in which the analogy is made explicitly. In particular, try to understand their respective notions of "constraint" and the role of modularity in their respective methodologies. There aren't very many statements around of the traditional AI methodology, if there can be said to be such a thing. (Winograd 1980) is one methodological discussion; I disagree with it but it's an important paper anyway. (McDermott 1976) makes a number of methodological criticisms of AI work. I advise thinking very hard about this sort of thing while working on AI.

Theory

There are important and subtle intuitions to be gotten from the theory of computation, and even if there aren't one should know the basic results just for culture's sake. (Minsky 1967) is a good text covering traditional theoretical ideas. (Clinger 1981) contains a more general formulation of the notion of computation that does not require the assumption of a discrete global clock; it deserves to be much better known. It might help to know some "Scottery"; (Stoy 1977) is a good text for the Scott-Strachey approach to formal semantics.

Architecture

Most computers over the past thirty years have been "von Neumann" (serial) machines, which make a strict separation between process and data and thus have a "von Neumann bottleneck" (Backus 1978) between processor and memory. Remedies for this situation have been pleaded for and attempted quite frequently but have met with little success for lack of appropriate high-level non-serial computational concepts. The latest series of such attempts began with Fahlman's (1979) NETL machine, a parallel database machine for AI applications. (Hillis 1981) describes the beginnings of the "Connection Machine" project, which uses a large network of small general processors which are capable of sending messages to one another. Connection Machine algorithms are being developed for a wide variety of applications. Both machines are about to begin construction.

Cognitive psychology

Cognitive psychology has become newly invigorated in recent times as ideas and approaches derived from artificial intelligence and Chomskyan generative linguistics have come to find wider application. (Lindsay and Norman 1972), (Newell and Simon 1972), (Johnson-Laird and Wason 1977) are three collections of works which show computational influences. See also the work of Gentner (e.g., 1980) and Norman (e.g., 1981). For a review of influences from generative linguistics, see (Keil 1981). (Berwick and Weinberg 1982 forthcoming) should help make the Chomskyan approach intelligible to the computationally oriented. Many psychologists could make good use of computational concepts and intuitions; see (Boden 1979) for an analysis of Piaget as one such. The psychology of music is a promising area for computational research; see (Levitt 1981) for a computational model of jazz improvisation and (Minsky 1981) for the beginnings of a computational approach to music theory. A competing approach to cognitive psychology is the "ecological", which is derived largely from the work of Gibson (e.g., 1979). (Ullman 1980) is a critique of the ecological approach to perception from a computational viewpoint. (Neisser 1976) is an interesting synthesis of computational and ecological ideas. (Neisser 1982) looks at memory in light of this synthesis (it's great fun in any event); (Baddeley 1976) is a good review of the traditional memory literature. Read (Bartlett 1932) for culture.

Neurophysiology

Detailed studies of the nervous system no longer have any significant influence on AI work on the higher cognitive processes. However, the abstract idea of the brain as a vast "network" of small processors with something like spatial locality still exerts a strong force in AI models of both representation and processing. Just about anything written by Minsky or Sussman (for example) can be looked at in this way. AI, on the other hand, does have something to say about neurophysiology. Neurophysiology is Marr's "implementation" level, and (Marr 1982) describes how his methodological framework has enabled neurophysiologists to get a handle on the mechanisms underlying visual processing. (Marr 1969,1970) are earlier neurophysiological works by Marr; try to understand what methodological criticisms he later came to make of them.

Philosophy

Although AI is of great interest to philosophers, even good AI people generally don't make very good philosophers themselves. This is largely because of their tendency to think of philosophical problems as trivial given a proper understanding of AI. See (Johnson-Laird 1977, Fodor 1978, Johnson-Laird 1978, Fodor 1979) for an instructive example concerned with "procedural semantics". (Hofstadter 1979) is an amusing popular account of some of the philosophical issues surrounding AI. Haugeland's "Mind Design" anthology (1981) contains most of the standard references for philosophy that is more or less friendly to the idea of AI. The two main buzzwords are "functionalism" and "intentionality". The first is understood and the second isn't. Dennett is my favorite of this lot; his papers are collected in (Dennett 1978). Compare his outlook to Minsky's. Read anything by Brian Smith that you can get your hands on. Fodor's papers are collected in (Fodor 1981). (Sloman 1978) is uneven but useful even so. (Turing 1950) started the modern "can a computer think?" debate; it is still interesting though dated.

AI has an opposition. It is to be taken seriously. Start with (Miller 1978), which is an intelligent attempt at philosophical refereeing. Then try to figure out what's wrong with (Searle 1980); Searle says it can't be done. Dreyfus (1979) says it can't be done too, but for more fundamental reasons. He's probably wrong too, of course, but you should figure out why for yourself, and then read Papert's unpublished reply (1968). (Weizenbaum 1976) says it shouldn't be done, and the argument is sufficiently cogent that you have a moral obligation to read it.

Science fiction

It is important to keep in mind that AI is very hard and that we haven't yet had one percent of the good ideas that will be required to do a proper job of it. One good way to do this is to sample the relevant science fiction now and again. Start with Hofstadter and Dennett's anthology "The Mind's I" (1981), which has a number of excellent pieces by people like Stanislaw Lem and Jorge Luis Borges. (Although it is advertised as semi-serious it is best regarded as intellectual cotton-candy.) Here is a list of some other things to look at, roughly in order: Robert Heinlein's "The Moon is a Harsh Mistress", Stanislaw Lem's "The Cyberiad", Isaac Asimov's "I, Robot" and "The Rest of the Robots", Fred Pohl's "Beyond the Blue Event Horizon", John Varley's "Overdrawn at the Memory Bank", and Vernor Vinge's "True Names". (This last story is rather hard to find; it was published by Dell in #5 of its "Binary Star" series. It has a cult following among computer types.)

Bibliography

A technical note for people who worry about such things: A number of the works cited in this bibliography have never been properly published, and so you shouldn't conclude that you can cite a work formally just because it appears here.

James F Allen, An interval-based representation of temporal knowledge, Proc IJCAI-81, pp 221-226

Dana Angluin and Carl H Smith, A survey of inductive inference: Theory and methods, Yale Univ Computer Science Dept, 1982, forthcoming

John Backus, Can programming be liberated from the von Neumann style?: A functional style and its calculus of programs, Comm ACM, v21 n8 August 1978, pp 613-641

Alan D Baddeley, The psychology of memory, Harper and Row, 1976

Avron Barr and Edward A Feigenbaum, eds, The handbook of artificial intelligence, vol 1 Kaufmann 1981, vols 2 and 3 Kaufmann 1982

F C Bartlett, Remembering: A study in experimental and social psychology, Cambridge Univ Press, 1932

Hans J Berliner, An examination of brute force intelligence, Proc IJCAI-81, pp 581-587

Robert C Berwick, Locality principles and the acquisition of syntactic knowledge, PhD Thesis, MIT Dept of Electrical Engineering and Computer Science, 1981

Robert C Berwick and Amy S Weinberg, The role of grammars in models of language use, Cognition, 1982, forthcoming

Garrett Birkhoff, Lattice Theory, Amer Math Soc Colloquium Publications, 25, 1967

Emilio Bizzi and William Abend, Posture control and trajectory formation in single and multiple joint arm movements, in J E Desmedt, ed, Brain and spinal mechanisms of movement control in man: New developments and clinical applications, Raven Press, 1982

Daniel G Bobrow and Allan M Collins, eds, Representation and understanding: Studies in cognitive science, Academic Press, 1975

Daniel G Bobrow and Terry Winograd, An overview of KRL: A knowledge representation language, Xerox PARC report CSL-76-4, July 1976

Margaret Boden, Artificial intelligence and natural man, Basic Books, 1977

Margaret Boden, Piaget, Harvester, 1979

Alan Borning, Thinglab: A constraint-oriented simulation laboratory, Stanford Univ Computer Science Dept Report 79-746, also Xerox PARC report SSL-79-3, 1979

Ronald J Brachman, On the epistemological status of semantic networks, BBN Report 3807, April 1978 (a)

Ronald J Brachman, A structural paradigm for representing knowledge, BBN Report 3605, May 1978 (b)

Ronald J Brachman, KLONE reference manual, BBN Report 3848, July 1978 (c)

Michael Brady, Computational approaches to image understanding, MIT AI Lab Memo 653, October 1981, also in ACM Computing Surveys, v14 n1 March 1982, pp 3-71

Michael Brady and Robert C Berwick, eds, Computational models of discourse, MIT Press, 1982, forthcoming

Rodney A Brooks, Symbolic reasoning among 3-d models and 2-d images, PhD Thesis, Stanford Dept of Computer Science, June 1981 (a)

Rodney A Brooks, Model-based three dimensional interpretations of two dimensional images, Proc IJCAI-81, pp 619-624 (b)

John Seely Brown and Kurt VanLehn, Toward a generative theory of "bugs", Xerox PARC report CIS-2, December 1979

Bruce G Buchanan and Tom M Mitchell, Model-directed learning of production rules, in (Waterman and Hayes-Roth 1978), pp 297-312

Eugene Charniak, Toward a model of children's story comprehension, MIT AI Lab TR 266, December 1972

Eugene Charniak, A partial taxonomy of knowledge about actions, Proc IJCAI-75, pp 91-97

Eugene Charniak, Christopher K Riesbeck, and Drew V McDermott, Artificial intelligence programming, Erlbaum, 1980

Noam Chomsky, Reflections on language, Pantheon, 1975

Noam Chomsky, Rules and representations, Columbia Univ Press, 1980

Douglas W Clark and C Cordell Green, An empirical study of list structure in Lisp, Comm ACM, v20 n2 February 1977, pp 78-87

William D Clinger, Foundations of actor semantics, MIT AI Lab TR 633, May 1981

Randy Davis, Interactive transfer of expertise: Acquisition of new inference rules, Artificial Intelligence, v 12, 1979, pp 121-157

Johan de Kleer, Jon Doyle, Guy L Steele Jr, and Gerald Jay Sussman, Explicit control of reasoning, MIT AI

Lab Memo 427, June 1977, also in Proc Symposium on AI and Programming Languages, in ACM Sigplan Notices/Sigart Newsletter, August 1977

Johan de Kleer, Jon Doyle, Charles Rich, Guy L Steele Jr, and Gerald Jay Sussman, AMORD: A deductive procedure system, MIT AI Lab Memo 435, January 1978

Johan de Kleer and Gerald Jay Sussman, Propagation of constraints applied to circuit synthesis, MIT AI Lab Memo 485, September 1978, also in Intl J Circuit Theory, v8 n2 pp 127-144, April 1980

Johan de Kleer, Causal and teleological reasoning in circuit recognition, MIT AI Lab TR 529, September 1979

Johan de Kleer and John Seely Brown, Assumptions and ambiguities in mechanistic mental models, Xerox PARC report CIS-9, March 1982

Daniel C Dennett, Brainstorms: Philosophical essays on mind and psychology, Bradford, 1978

Andrea diSessa, Unlearning Aristotelian physics: A study of knowledge-based learning, Cognitive Science, 1982, pp 37-75

Jon Doyle, Truth maintenance systems for problem solving, MIT AI Lab TR 419, January 1978

Jon Doyle and Philip London, A selected descriptor-indexed bibliography to the literature on belief revision, ACM Sigart Newsletter, April 1980, pp 7-23

Jon Doyle, A model for deliberation, action, and introspection, MIT AI Lab TR 581, May 1980

B Elan Dresher and Norbert Hornstein, On some supposed contributions of artificial intelligence to the scientific study of language, Cognition, v4 n4 December 1976, pp 321-398; Terry Winograd's reply is (Winograd 1977)

B Elan Dresher and Norbert Hornstein, Reply to Winograd, Cognition, v5 n4 December 1977, pp 379-392; this is a reply to (Winograd 1977)

Hubert Dreyfus, What computers can't do, 2nd ed, New York, 1979; the introduction is also in (Haugeland 1981) pp 161-204

Lee D Erman, Frederick Hayes-Roth, Victor R Lesser, and D Raj Reddy, The Hearsay-II speech-understanding system: Integrating knowledge to resolve uncertainty, Computing Surveys, v12 n2 June 1980, pp 213-253

Thomas G Evans, A program for the solution of geometric-analogy intelligence test questions, in (Minsky 1968a)

Scott E Fahlman, A planning system for robot construction tasks, MIT AI Lab TR 283, May 1973; see also (Fahlman 1974)

Scott E Fahlman, A planning system for robot construction tasks, Artificial Intelligence, v5 n1 Spring 1974, pp 1-49

Scott E Fahlman, NETL: A system for representing and using real-world knowledge, MIT Press, 1979

Edward A Feigenbaum and Julian Feldman, eds, Computers and thought, McGraw-Hill, 1963

Richard E Fikes and Nils J Nilsson, Strips: A new approach to the application of theorem proving to problem solving, Artificial Intelligence, v2 n3 1971, pp 189-208

Richard E Fikes, Peter E Hart, and Nils J Nilsson, Learning and executing generalized robot plans, Artificial Intelligence, v3 n4 Winter 1972, pp 251-288

Jerry Fodor, Tom Swift and his procedural grandmother, Cognition, v6 n3 September 1978, pp 229-247; this is a reply to (Johnson-Laird 1977), and Johnson-Laird's reply is in (Johnson-Laird 1978)

Jerry Fodor, In reply to Philip Johnson-Laird, Cognition, v7 n1 March 1979, pp 93-95; this is a reply to (Johnson-Laird 1978)

Jerry Fodor, Representations, MIT Press, 1981

Kenneth D Forbus, A study of qualitative and geometric knowledge in reasoning about motion, MIT AI Lab TR 615, February 1981

Kenneth D Forbus, Qualitative process theory, MIT AI Lab Memo 664, February 1982

Daniel P Friedman, The Little Lisper, Science Research Associates, 1974

Dedre Gentner, The structure of analogical models in science, BBN Report 4451, 1980

Dedre Gentner, Are scientific analogies metaphors?, in D S Miall, ed, Metaphor: Problems and perspectives, Harvester, 1982

James J Gibson, The ecological approach to visual perception, Houghton Mifflin, 1979

W Eric L Grimson, From images to surfaces, MIT Press, 1981 (a)

W Eric L Grimson, The implicit constraints of the primal sketch, MIT AI Lab Memo 663, October 1981 (b)

John Haugeland, ed, Mind design, Bradford, 1981

Patrick J Hayes, In defence of logic, Proc IJCAI-77, pp 559-565

Patrick J Hayes, The naive physics manifesto, in Donald Michie, ed, Expert systems in the micro-electronic age, Edinburgh Univ Press, 1979 (a)

Patrick J Hayes, Naive physics I: Ontology for liquids, memo, Centre pour les Etudes Semantiques et Cognitives, Geneva, 1979 (b)

Carl Hewitt, Procedural embedding of knowledge in Planner, Proc IJCAI-71, pp 167-182

Carl Hewitt, Description and theoretical analysis (using schemata) of Planner: A language for proving theorems and manipulating models in a robot, MIT AI Lab TR 258, April 1972

Carl Hewitt, Viewing control structures as patterns of passing messages, Artificial Intelligence, v8 n3 June 1977, pp 323-364

W Daniel Hillis, The connection machine, MIT AI Lab Memo 646, September 1981

Douglas R Hofstadter, Godel, Escher, Bach: An eternal golden braid, Basic Books, 1979

Douglas B Hofstadter and Daniel C Dennett, eds, The mind's I: Fantasies and reflections on self and soul, Basic Books, 1981

John M Hollerbach, A recursive lagrangian formulation of manipulator dynamics and a comparative study of dynamics formulation complexity, MIT AI Lab Memo 533, June 1979, also in IEEE Trans on Systems, Man, and Cybernetics, vol SMC-10 n11 November 1980, pp 730-736

John M Hollerbach, Dexterous hand design and control workshop, MIT AI Lab Memo 661, January 1982

David J Israel, What's wrong with non-monotonic logic?, Proc AAAI-80, pp 99-101

David J Israel and Ronald J Brachman, Distinctions and confusions: A catalogue raisonne, Proc IJCAI-81, pp 452-459

Philip N Johnson-Laird and P C Wason, Thinking: Readings in cognitive science, Cambridge Univ Press, 1977

Philip N Johnson-Laird, Procedural Semantics, Cognition, v5 n3 September 1977, pp 189-214; Jerry Fodor's reply is in (Fodor 1978)

Philip N Johnson-Laird, What's wrong with Grandma's guide to procedural semantics: A reply to Jerry Fodor, Cognition, v6 n3 September 1978, pp 249-261; this is a reply to (Fodor 1978), and Fodor's reply is in (Fodor 1979)

Frank C Keil, Constraints on knowledge and cognitive development, Psychological Review, v88 n3 May 1981, pp 197-227

Dennis H Klatt, Review of the ARPA Speech Understanding Project, J Acoust Soc Amer, vol 62 no 6, 1978

David A Levitt, A melody description system for jazz improvisation, MS thesis, MIT Dept of Electrical Engineering and Computer Science, 1981

Peter H Lindsay and Donald A Norman, Human information processing: An introduction to psychology, Academic Press, 1972

Robert K Lindsay, Bruce G Buchanan, Edward A Feigenbaum, and Joshua Lederberg, Applications of AI for organic chemistry: The Dendral project, McGraw-Hill, 1981

H C Longuet-Higgins and K Prazdny, The interpretation of moving retinal images, Proc R Soc Lond B, v208, pp 385-387, 1980

Bruce T Lowerre, The Harpy speech recognition system, PhD thesis, CMU Dept of Computer Science, April 1976

Richard F Lyon, A computational model of filtering, detection, and compression in the cochlea, Proc ICASSP, 3-5 May 1982, Paris

Mitchell P Marcus, A theory of syntactic recognition for natural language, MIT Press, 1980

David Marr, A theory of cerebellar cortex, J Physiol (London), v202, pp 437-470, 1969

David Marr, A theory for cerebral neocortex, Proc R Soc Lond B, v176, pp 161-234, 1970

David Marr, Artificial intelligence: A personal view, Artificial Intelligence, v9 n1 August 1977, pp 37-48, also in (Haugeland 1981), pp 129-142

David Marr, Vision, Freeman, 1982

Matthew T Mason, Compliance and force control for computer controlled manipulators, MIT AI Lab TR 515, April 1979, also in IEEE Trans on Systems, Man, and Cybernetics, vol SMC-11 n6 June 1981, pp 418-432

John McCarthy, History of Lisp, ACM Sigplan Notices, v13 n8 August 1978, pp 217-223

John McCarthy, Epistemological problems of artificial intelligence, Proc IJCAI-77, pp 1038-1044

Drew McDermott and Gerald Sussman, The conniver reference manual, MIT AI Lab Memo 259A, January 1974

Drew McDermott, Artificial intelligence meets natural stupidity, ACM Sigart Newsletter, April 1976, also in (Haugeland 1981), pp 143-160

Drew McDermott, Flexibility and efficiency in a computer program for designing circuits, MIT AI Lab TR 402, June 1977 (a)

Drew McDermott, Vocabularies for problem solver state descriptions, Proc IJCAI-77, pp 229-234 (b)

John McDermott, R1: A rule-based configurer of computer systems, CMU Computer Science Dept report

CS-80-119, April 1980

David D McDonald, Natural language generation as a computation problem: An introduction, in (Brady and Berwick 1982 forthcoming), pp 209-265

Bernard Melzer and Donald Michie, eds, Machine intelligence 5, Elsevier, 1970

Ryszard S Michalski, Knowledge acquisition through conceptual clustering: A theoretical framework and an algorithm for partitioning data into conjunctive concepts, Policy Analysis and Information Systems, Special Issue on Knowledge Acquisition and Induction, v4 n3 1980 .

Laurence Miller, Has artificial intelligence contributed to an understanding of the human mind?: A critique of arguments for and against, Cognitive Science, v2 n2 April-June 1978, pp 111-127

Marvin Minsky, Steps toward artificial intelligence, Proc IRE, v49 n1, January 1961, also in (Feigenbaum and Feldman 1963)

Marvin Minsky, Computation: Finite and infinite machines, Prentice-Hall, 1967

Marvin Minsky, ed, Semantic information processing, MIT Press, 1968 (a)

Marvin Minsky, Mind, matter, and models, in (Minsky 1968a) (b)

Marvin Minsky and Seymour Papert, Perceptrons, MIT Press, first edition 1969, second edition 1972

Marvin Minsky, A framework for representing knowledge, MIT AI Lab Memo 306, June 1974, also in part in (Winston 1975a) pp 211-277 and in (Haugeland 1981) pp 95-128; another version appears in Proc TINLAP 1975 and in (Johnson-Laird and Wason 1977) pp 355-376

Marvin Minsky, K-Lines: A theory of memory, MIT AI Lab Memo 516, June 1979, also in Cognitive Science, 1980, pp 117-133

Marvin Minsky, Jokes and the logic of the cognitive unconscious, MIT AI Lab Memo 603, November 1980

Marvin Minsky, Music, mind, and meaning, MIT AI Lab Memo 616, February 1981; a revised version appears in the Computer Music Journal, v5 n3 Fall 1981

Marvin Minsky, Learning meaning, forthcoming MIT AI Lab Memo, 1982

Tom Mitchell, Jaime Carbonell, and Ryszard Michalski, eds, Sigart Special Issue on Machine Learning, April 1981, pp 25-54

Robert C Moore, Reasoning about knowledge and action, Proc IJCAI-77, pp 223-227

Robert C Moore, Reasoning about knowledge and action, PhD Thesis, MIT Dept of Electrical Engineering and Computer Science, February 1979

Hans P Moravec, Obstacle avoidance and navigation in the real world by a seeing robot rover, Stanford AI Lab Memo AIM-340, September 1980, also CMU Robotics Institute TR-3

Joel Moses, The function of Function in Lisp, or, Why the funarg problem should be called the environment problem, MIT Project MAC Memo AI-199, June 1970

Ulric Neisser, Cognition and reality, Freeman, 1976

Ulric Neisser, Memory observed: Remembering in natural contexts, Freeman, 1982

Allen Newell and Herbert A Simon, GPS, a program that simulates human thought, in (Feigenbaum and Feldman 1963)

Allen Newell, J C Shaw, and Herbert A Simon, A variety of intelligent learning in a general problem solver, in Marshall C Yovits and Scott Cameron, eds, Self-organizing systems, Pergamon, 1960, pp 153-189

Allen Newell and Herbert A Simon, Human problem solving, Prentice-Hall, 1972

Nils J Nilsson, Problem solving methods in artificial intelligence, McGraw-Hill, 1971

Nils J Nilsson, Principles of artificial intelligence, Tioga, 1980

Donald A Norman, Categorization of action slips, Psychological Review, v88 1981, pp 1-15

Bernard Nudel and Paul E Utgoff, A bibliography on machine learning, ACM Sigart Newsletter, April 1981, pp 54-64

Seymour Papert, The artificial intelligence of Hubert L Dreyfus, unpublished manuscript, 1968

Seymour Papert, Mindstorms, Basic Books, 1980

Richard P Paul, Robot manipulators: Mathematics, programming, and control, MIT Press, 1981

David A Plaisted, Theorem proving with abstraction, Artificial Intelligence, v16 n1 January 1981, pp 47-108

Gordon D Plotkin, A note on inductive generalization, in (Melzer and Michie 1970), pp 153-163

M Ross Quillian, Semantic memory, in (Minsky 1968a), pp 227-270

John C Reynolds, Transformational systems and the algebraic structure of atomic formulas, in (Melzer and Michie 1970), pp 135-151

Charles Rich and Howard Shrobe, Initial report on a Lisp programmer's apprentice, IEEE Trans on Software Engineering, v4 n5 November 1978

Charles J Rieger III, Conceptual memory: A theory and computer program for processing the meaning content of natural language utterances, Stanford AI Lab Memo AIM-233, July 1974

Charles J Rieger III, An organization of knowledge for problem solving and language comprehension, Artificial Intelligence, v7 n2 Summer 1976, pp 89-127

R Bruce Roberts and Ira P Goldstein, The FRL manual, MIT AI Lab Memo 409, June 1977

John A Robinson, A review of automatic theorem-proving, Proc Symp on Appl Math of the Amer Math Soc, 19, 1967, pp 1-18

Earl D Sacerdoti, Planning in a hierarchy of abstraction spaces, Artificial Intelligence, v5 n2 Summer 1974, pp 115-135

Earl D Sacerdoti, The nonlinear nature of plans, IJCAI-75, pp 206-214

Earl D Sacerdoti, A structure for plans and behavior, Elsevier, 1977

J Kenneth Salisbury, Kinematic and force analysis of articulated hands, PhD Thesis, Stanford Mechanical Engineering Dept, May 1982 (a)

J Kenneth Salisbury, Articulated hands: Force control and kinematic issues, Intl J Robotics Research, v1 n1, Spring 1982, pp 4-17 (b)

Roger C Schank and Kenneth M Colby, Computer models of thought and language, Freeman, 1973 (a)

Roger C Schank, Identification of conceptualizations underlying natural language, in (Schank and Colby 1973a), pp 187-247 (b)

Roger C Schank and Charles J Rieger III, Inference and the computer understanding of natural language, Artificial Intelligence, v5 n4 Winter 1974, pp 393-412

Roger C Schank, Language and memory, Cognitive Science, v4 n3 July-September 1980, pp 243-284

Roger C Schank and Christopher K Riesbeck, Inside computer understanding, Erlbaum, 1981

Campbell L Searle, J Zachary Jacobson, and S G Rayment, Stop consonant discrimination based on human audition, J Acoust Soc Am, v65 n3 March 1979, pp 799-809

John Searle, Minds, brains, and programs, with peer commentary and author's response, Brain and Behavioral Sciences, v3 n3 September 1980, pp 417-424 (with peer commentary pp 424-450 and author's response pp 450-457), also in (Haugeland 1981) pp 282-306, also in (Hofstadter and Dennett 1981) pp 353-373 (with commentary by Hofstadter pp 373-382)

Edward H Shortliffe, Computer-based medical consultations: Mycin, Elsevier, 1976

Howard E Shrobe, Dependency-directed reasoning for complex program understanding, MIT AI Lab TR

503, April 1979 (a)

Howard E Shrobe, Dependency-directed reasoning in the analysis of programs which modify complex data structures, Proc IJCAI-79, pp 829-835 (b)

Herbert A Simon, The sciences of the artificial, MIT Press, 1970

Herbert A Simon, Models of thought, Yale U Press, 1979

Derek H Sleeman and M J Smith, Modelling students' problem solving, Artificial Intelligence, v16 n2 May 1981, pp 171-187

Aaron Sloman, The computer revolution in philosophy, Harvester, 1978

Brian C Smith, Reflection and semantics in a procedural language, MIT Lab for Computer Science Report TR-272, 1982

Christopher C Stacy, Getting started computing at the AI Lab, MIT AI Lab Working Paper 235, 1982

Richard M Stallman and Gerald Jay Sussman, Forward reasoning and dependency-directed backtracking in a system for computer-aided circuit analysis, MIT AI Lab Memo 380, October 1977, also in Artificial Intelligence, v9 n2 October 1977, pp 135-196

Guy L Steele Jr, Rabbit: A compiler for Scheme, MIT AI Lab TR 474, May 1978

Mark Stefik, Planning with constraints, Stanford Heuristic Programming Project Memo HPP-80-2, January 1980

Mark Stefik, Planning with constraints (Molgen: Part 1), Artificial Intelligence, v16 n2 May 1981, pp 111-139

Mark Stefik, Jan Aikins, Robert Balzer, John Benoit, Lawrence Birnbaum, Frederick Hayes-Roth, and Earl Sacerdoti, The organization of expert systems: A tutorial, Artificial Intelligence, v18 n2 March 1982, pp 135-173

Albert Stevens, Allan Collins, and Sarah E Goldin, Misconceptions in students' understanding, Int J Man-Machine Studies, v11, 1979, pp 145-156

Joseph E Stoy, Denotational semantics: The Scott-Strachey approach to programming language theory, MIT Press, 1977

Gerald Jay Sussman and Drew McDermott, Why conniving is better than planning, MIT AI Lab Memo 255A, April 1972, also in Proc FJCC v41 pp 1171-1179 AFIPS Press 1972; the unexpurgated version is AIM 255

Gerald Jay Sussman, A computer model of skill acquisition, MIT AI Lab TR-297, August 1973, also Elsevier, 1975

Gerald Jay Sussman and Guy L Steele Jr. Constraints: A language for expressing almost-hierarchical descriptions, MIT AI Lab Memo 502A, August 1981, also in Artificial Intelligence, v14 n1 August 1980, pp 1-39

Demetri Terzopoulos, Multi-level reconstruction of visual surfaces, MIT AI Lab Memo 671, April 1982

Alan M Turing, Can a machine think?, Mind, October 1950, pp 433-460, also in (Feigenbaum and Feldman 1963)

Shimon Ullman, Against direct perception, MIT AI Lab Memo 574, March 1980, also in Behavioral and Brain Sciences, v3 n3 September 1980, pp 373-381 (with peer commentary pp 381-408 and author's response pp 408-415)

David L Waltz, Generating semantic descriptions from drawings of scenes with shadows, MIT AI Lab TR 271, November 1972

David L Waltz, Understanding line drawings of scenes with shadows, in (Winston 1975a), pp 19-91

Donald A Waterman and Frederick Hayes-Roth, eds, Pattern-directed inference systems, Academic Press, 1978

Richard C Waters, The programmer's apprentice: Knowledge based program editing, IEEE Trans on Software Engineering, Vol SE-8 n1 January 1982

Norbert Wiener, Cybernetics, MIT Press and Wiley, 1961

Joseph Weizenbaum, Computer power and human reason, Freeman, 1976

Richard W Weyhrauch and Arthur J Thomas, FOL: A proof checker for first-order logic, Stanford AI Lab Memo AIM-235, September 1974

Terry Winograd, Procedures as a representation for data in a computer program for understanding natural language, MIT AI Lab TR 17, February 1971

Terry Winograd, A procedural model of language understanding, in (Schank and Colby 1973a), pp 152-186

Terry Winograd, On some contested suppositions of generative linguistics about the scientific study of language, Stanford AI Lab Memo 300, May 1977, also in Cognition, v5 n4 December 1977, pp 379-392; this is a reply to (Dresher and Hornstein 1976), and their reply is in (Dresher and Hornstein 1977)

Terry Winograd, What does it mean to understand language?, Cognitive Science, 1980, pp 209-241

Terry Winograd, Language as a cognitive process, volume 1: Syntax, Addison-Wesley, 1983

Patrick H Winston, Learning structural descriptions from examples, MIT AI Lab TR 231, September 1970, summarized in (Winston 1975b)

Patrick H Winston, The psychology of computer vision, McGraw-Hill, 1975 (a)

Patrick H Winston, Learning structural descriptions from examples, in (Winston 1975a), pp 157-209 (b)

Patrick H Winston, Artificial intelligence, Addison-Wesley, 1977

Patrick H Winston and Richard H Brown, eds, Artificial intelligence: An MIT perspective, 2 vols, MIT Press, 1979

Patrick H Winston, Learning and reasoning by analogy: The details, MIT AI Lab Memo 520, May 1980

Patrick H Winston and Berthold K P Horn, Lisp, Addison-Wesley, 1981

Patrick H Winston, Learning by augmenting rules and accumulating censors, MIT AI Lab Memo 678, May 1982

William A Woods, Transition network grammars for natural language analysis, Comm ACM, v13 n10 October 1970, pp 591-606

William A Woods, What's in a link?: Foundations for semantic networks, in (Bobrow and Collins 1975), pp 35-82

William A Woods et al, Speech understanding systems: Final report (November 1974 - October 1976), BBN Report 3438, 5 volumes, 1976

Richard M Young and Tim O'Shea, Errors in children's subtraction, Cognitive Science, v5 n2 April-June 1981, pp 153-177

Victor W Zue and Richard M Schwartz, Acoustic processing and phonetic analysis, in W A Lea, ed, Trends in speech recognition, 1980

Victor W Zue, Acoustic phonetic knowledge representation: Implications from spectrogram reading experiments, MIT Speech Group Working Paper, 1981