December 1985

Hcport No. STAN-CS-86-1093
*Also numbered KSL-85-54*

# A General Reading List for Artificial Intelligence

by

Devika Subramanian and Bruce G. Buchanan

**Department of Computer Science**

Stanford University
Stanford, CA 94305

**Knowledge Systems Laboratory**
**Report No. KSL-85-54**

**December 1985**

# A General Reading List

# for

# Artificial Intelligence

**Devika Subramanian and Bruce G. Buchanan**

**Department of Computer Science**
**Stanford University**
**Stanford, CA 94305**

This work was funded in part by the following contracts and grants: DARPA N00039-83-C-0136, NIH/SUMEX RR-00785-12, NASA-AMES NCC-2-274, Boeing Computer Services W266875, NSF IST-8312148, and a gift from Lockheed. Devika Subramanian is supported by an IBM Graduate Fellowship.

KSL-85-54

revised December 1985
first draft April 1985

# A General Reading List

# for

# Artificial Intelligence

by

Devika Subramanian and Bruce G. Buchanan


This reading list is based on the syllabus for the course CS229b offered in Winter 1985. This course was an intensive 10 week survey intended as preparation for the 1984-85 qualifying examination in Artificial Intelligence at Stanford University.

# Preface

In the 1984-85 academic year, we offered a seminar to Ph.D. students studying for the qualifying examination in AI in the Computer Science Department at Stanford. Since the intent was to survey nearly all of AI and highlight key issues, the annotated reading list may be helpful to others who are getting started in AI. We organized the readings in ten topics, corresponding to the ten weeks the seminar ran.

The first section, Introduction to AI, is a list of books and articles that will help in obtaining an understanding of the enterprise of AI. Topics 2 through 5 form part of a theoretical core for AI. They set the stage for understanding issues in the subsequent, areas. Topics 6 through 9 are specialized topics in which AI research has been driven by specialized applications. The final section, Advanced Topics, covers many important current issues which are not properly included in the earlier, more basic sections.

This reading list is graded and annotated. For each topic, we present basic reading drawn for the most part from the AI Handbook. Required papers are taken from the Webber-Nilsson collection of AI readings. Recommended readings serve both to introduce new research (post 1982) and also reinforce material covered sketchily in the basic reading.

This compilation of the reading list is an ongoing effort. Comments and suggestions on this are very welcome for future editions.

# Contents

# 1 Introduction to AI

This set of readings will help obtain an understanding of what Artificial Intelligence is and provide a historical perspective. These may be read along with the reading on the other topics. These are not pre-requisites for any of the other readings in this syllabus. For current discussion on the nature of AI see the reading list on Advanced Topics in this syllabus. The opening chapters of AI textbooks by Winston, Rich and McDermott & Charniak, Nilsson & Genesereth are also recommended.

- ***Computers and Thought: Feigenbaum and Feldman, McGraw-Hill, 1963***

  Though a bit dated, this book is still an important collection of ideas in Artificial Intelligence. Turing's epoch making article "Computing Machinery and Intelligence" is reprinted here. Minsky's "Steps towards AI" contains a research program that asks questions that are important(and are unanswered!) even today. Gelernter's geometry machine and Samuel's checker player are also described here. The dichotomy in AI between those who wish to duplicate human intelligence (and use AI as a vehicle for studying human intelligence) and those who wish to create a machine intelligence (without regard to whether or not it is human) is evident in the organization of this book.

- ***Semantic Information Processing: Minsky, MIT Press, 1968***

  This collection mostly contains articles that represent pioneering work on AI, done at MIT in the early and mid-sixties. Included are Raphael's SIR, Evans' geometry analogy program and Bobrow's STUDENT. McCarthy's "Programs with Common Sense" is reprinted here as well as Minsky's thought-provoking "Matter, Minds and Models". The preface is exceedingly well written and is of historical interest.

- ***Artificial Intelligence and Natural Man: Boden, Basic Books, 1977***

  A classic. Has an extensive annotated bibliography.

- ***Machines who think:*** *McCorduck*, ***Freeman, 1979***

  A very interesting (and entertaining) book which traces the history of AI in the United States.

# 2 Search and Heuristics

Basic Reading

- ***Artificial Intelligence (2nd ed.), Patrick Winston, 1984, Addison- Wesley, Chapter 4***

  This is a very readable introduction to basic search strategies. The taxonomy of search strategies on Page 88 should help organize the material learnt from Section C of Chapter 2 of the Handbook. The examples presented here are good. It is recommended that you skim through this before you read the presentation in the Handbook and in Nilsson's book.

- ***The Handbook of Artificial Intelligence, 1981, William Kaufmann Inc., Volume 1, Chapter 2***

  This is a reasonably well written introduction to the vast literature in search. It gives pointers to most of the important papers in this area. Sections B and C contain basic concepts that should be learnt well. A lighter introduction to material in Section C is in Winston (above). Section C.3 and Chapter 2 of Nilsson's text are mutually redundant, so are Section C.4 and Chapter 3 of Nilsson's text. Details of systems in Section D are unimportant but you do need to know which search strategy was used and why.

- ***Artificial Intelligence, Elaine Rich, 1983, McGraw*** Hill, ***Chapters 2,3 and 4***

  Sections 2.1, 3.6 through 3.7, and all sections in Chapter 4 are relevant.

- ***Principles of Artificial Intelligence, N.J. Nilsson***, 1980, Tioga ***Publishing Co., Chapters 2 and 3***

  There is a clear, detailed discussion of search strategies here. It is about 50 pages long, but familiarity with these methods is a prerequisite for planning, theorem proving, expert systems etc. You may skip Section 3.3 of Chapter 3. See remark on Handbook Chapter 2 above regarding overlap with this book.

*Search and Heuristics*

Required papers .

***Readings*** in ***artificial intelligence, Nilsson and Webber, Chapter 1***

. The five papers in this chapter introduce some of the important research issues in search.

- ***On Representations of Problems of*** Reasoning ***about Actions : Saul Amarel, 1968***

  Amarel's paper is a classic study on how shifts in problem representation can drastically reduce the size of the search space. Recent work at CMU (Korf 80) and Rutgers (Riddle 84) attempt to pursue the intriguing ideas presented here. This paper **occurs** again in this reading list under Knowledge Representation since this discusses the automation of representation shifts. For now, you should be familiar with examples of successive reformulations which reduce search effort in problem solving.

- ***A Problem*** Similarity ***Approach to Devising*** Heuristics **: *John* Gaschnig, *1979***

  The use of heuristic estimating functions for controlling search raises the question **of** how to obtain these functions. John Gaschnig's paper suggests an interesting approach laying a nice foundation for future work in this area. Short and well-written paper. Compare with the learning of heuristics ***a la*** Lenat.

- ***Optimal Search Strategies : William Woods, 1982***

  Woods's paper views recognition as search. Instead of searching for a minimal **cost** path to a goal state, Woods seeks the final state with the highest score (regardless of the the cost of the path to the state). The shortfall method for scoring states . is an instance of the **A*** algorithm. The density method, which is also optimal, is interesting because it is not an instance of **A*.** This paper ends with a comparison with the strategies used in other speech understanding systems, which **can** be **read** when we address the topic of Speech Understanding Research. The basic idea behind the shortfall and density methods must be learnt but the details are unimportant.

- ***Consistency*** in ***Networks of Relations : Alan Mackworth, 1977***

  Mackworth's exceedingly clear paper provides a very good introduction to constraint

satisfaction and network consistency algorithms. The style of presentation is worth noting as well as the contents, which are fundamental to understanding more recent work in this area. The three maladies of backtracking and their proposed remedies should be reasonably well understood. Think also about the use of constraints in AI systems[Steele 78].

- ***The B* Tree Search Algorithm : a Best First*** *Proof* ***Procedure : Hans Berliner,*** **1979**

  Much of the early work in search (and AI) was done in the context of game playing programs. Chess has posed particularly challenging problems here. Berliner's paper proposes the B* algorithm for searching game and proof trees. The main idea in this algorithm is the use of two bounds to cut off search - the optimistic bound used by the A* algorithm as well as a pessimistic bound. You should be able to explain this algorithm with an example.

Recommended Reading

This is a compendium of more recent work in search. For work before 1982, the handbook has all the pointers. It may not be necessary to read all that material if you are familiar enough with the presentation in the handbook.

- ***Nature of Heuristics I : Doug Lend***

  ***AI Journal, Vol 19, 1982***

  This is the first article in the three part series on heuristics. It introduces the field of *heuretics* and forms the basis for the results reported in the second and third articles referenced below.

- ***Special AI Journal issue on Search and Heuristics***

  ***Vol 21, Nos 1-2, March 1983***

  The stout-hearted can begin directly with the 10 articles here that are representative of current research in this area. The general theme that knits these papers together is the quest for understanding the workings of heuristic knowledge. The following

papers in this collection are recommended. The rest of the papers in this collection are on performance analysis of various search algorithms and game playing strategies. It is recommended that you read the abstracts or skim through the contents.

***Search and Reasoning in Problem Solving : Herb Simon, 1985***

This paper contrasts the search metaphor and reasoning metaphor in problem solving. Worth reading in full.

***Nature of Heuristics II and III : Doug Lenat, 1983***

These describe research aimed at automating the process of learning by discovery in various fields, including heuristics. Skim through them with special attention to the examples.

- ***Knowledge vs Search : A Quantitative Analysis using A* : J. Pearl***

  ***AI Journal, Vol. 20, No.1 ,1983***

  Knowledge and search are two major commodities that fuel and propel AI programs. We have a qualitative understanding of the interaction between the two (cf. Dendral). This paper is an attempt to quantify the knowledge-search tradeoff in the context of heuristic search algorithms. It is a mathematical exposition of the dependence on the average number of nodes expanded by A* on the accuracy of its heuristic estimate. Think about the problems of doing this sort of analysis in the context of a heuristic program, say Dendral.

- ***Strategies in Heuristic Search : M. Georgeff***

  ***AI Journal, Vol 20, No.4, 1983***

  Real valued heuristic functions have been extensively used as a means of constraining search in combinatorially large problem spaces. An alternative approach called strategic search is examined, in which heuristic information is expressed as problem specific strategies. These are intended to guide one toward a goal state, but there is no guarantee for success. Admissible algorithms using strategy search are presented.

Strategy search in hierarchically organized problem spaces and meta level strategies to guide application of base level strategies are also considered. Extremely interesting, look at some examples in the paper.

- ***1976 ACM*** Turing ***Award Lecture by*** Simon ***and Newell***

  ***CACM Vol 19, pages 113-126***

  "Computer Science as empirical inquiry – symbols and search".

  Emphasizes the primacy of search in AI. Historical interest.

- ***Constraints : Guy Steele, Jr. and Gerry Sussman***

  ***MIT-AI Memo 502, 1978***

  Constraint propagation is introduced here. Read and compare this with problem solving by search.

- ***Generalization as Search : T. Mitchell, 1982***

  ***Nilsson and Webber collection : Chapter 5, Article 8***

  The generalization problem in systems which learn from examples is presented in the search framework. Observe the mileage we get when we cast problems into the search perspective. This should start one thinking about search in relation to other areas in AI.

# 3 Knowledge Representation

Basic Reading

- ***Principles of Artificial Intelligence, N.J. Nilsson, 1980,*** Tioga ***Publishing Co., Chapters 1,4,9,10.3***

  Chapter 1 introduces production systems with lots of examples. Chapter 4 introduces predicate calculus as a knowledge representation language in AI. The following concepts should be learnt well : unification, pattern matching, converting to CNF, resolution (just the definition, for now), validity and satisfiability of a wff, completeness and soundness of a set of axioms and inference rules. Read section 4.3 carefully to get a sense of how Nilsson intends predicate calculus to be used in AI (also section 10.3 is relevant here).

  Chapter 9 introduces units and has a very extensive treatment of semantic nets and operations on them. The concepts of property inheritance and procedural attachments are very important. Read Section 9.6 to get a historical perspective on the semantic net formalism and also a comparison with other formalisms. Do exercise 9.5 at the end of the chapter with respect to every subarea of AI.

- ***The Handbook of Artificial Intelligence, 1981, William Kaujmann Inc., Volume 1, Chapter 3***

  Skim through Section A. For the procedural/declarative controversy see [Winograd 75] below. Skim also through Section B. Read Section C carefully to get a sense of the following for each representation formalism.

  - Example of a use of the knowledge representation formalism.
  - The operations that can be performed on it.
  - Disadvantages and advantages of the formalism with an example of a case where it would be hopeless and where it is extremely useful.

– Current research issues.

Supplement the material on semantic nets with the more comprehensive treatment in Nilsson. Production systems are well explained in the Davis and King article that is recommended. Semantic primitives are not a representational formalism ***per se***, but could form the basis of one. For more on this read Chapter 4 of the Handbook. The original paper on frames by Minsky is worth looking over, it has ideas that could be pursued for further research. The main utility of the Handbook is the large number of pointers it provides to the rather extensive literature on knowledge representation. It is not necessary to follow up on all these, except for the ones listed under recommended reading.

- ***Artificial Intelligence (2nd ed.), Patrick Winston, 1984, Addison-Wesley, Ch 2 and 8***

  Read pages 21-24 and pages 41-42 of Chapter 2 only if you are hard pressed for time. Look at the desiderata for good representation on Page 23. The rest of the chapter consists of examples highlighting the issues raised there. Read Chapter 8 after you have covered Section C of the Handbook. This will help in organizing the material that is covered there with some very good examples.

Required papers

***Readings in Artificial Intelligence: Nilsson and Webber***

- ***On Representations of Problems of Reasoning about Actions : Saul Amarel, 1968***

  Several successive reformulations of the familiar missionaries and cannibals problem lead to improved problem solving efficiency. The kinds of reformulations hinted at here have not yet been automated. What sorts of knowledge do we need to do reformulation in general? Do we need a theory of representations to so this? This is still an active area of research. Read the conclusions of this article carefully, almost every paragraph contains an idea worth exploring as a dissertation in AI. Think of an example of incremental reformulations in a domain other than the highly artificial one of the

missionaries and cannibals. Current work by Riddle [Riddle 84], Korf [Korf 82], Lowry [Lowry 84] attempts to solve some of the problems in this area of representation shifts.

- ***The Logic of Frames : Patrick Hayes, 1979***

  Several alternatives to logic as a representation language have been **proposed at various** times. Most have turned out to be syntactic variants of first order **logic rather than** fundamentally different systems. Many of these elevate important implementation details like indexing to the level of syntax of the language. This **paper shows how** a frame based system can be interpreted as a system of first order logic. **Contrast** this with the section on Page 21 of Winston's book (2nd ed) entitled *: **Theoretical equivalence is different from practical equivalence.***

## Recommended Reading

- ***Extending the Expressive Power of Semantic Nets : Schubert***

  ***AI Journal, Vol 7, 1976***

  The semantic net notation is extended for the representation of logical **connectives,** quantifiers, time and modal operators. Look at the examples in this paper. As **an** exercise : Represent the transitivity axiom using semantic nets.

- ***Some Problems and Non-problems of Representation Theory : P.J. Hayes***

  ***British Computer* Society, *AISB summer conference, 1974***

  This is Hayes' perspective of what is and is not important in **Knowledge Representation.** Recommended only for people who think they might **do a thesis in Knowledge** Representation.

- ***Programs with Common Sense : John McCarthy***

  ***In Minsky's Semantic Information Processing, MIT Press, 1968***

  Describes the Advice-Taker. Historically important. Asks the question: what are the rcprcscntational and reasoning requirements for a common **sense problem solver?**

- ***The Second Naive Physics Manifesto : P. J. Hayes***

  ***In Formal theories of the Commonsense World, edited by Jerry Hobbs and Robert Moore, Ablex Press, 1985***

  Hayes suggests a program of research for the construction of a formal theory of the commonsense world : he identifies 'clusters' in our commonsense knowledge and indicates an order in which to tackle them. Very valuable for those intending to do research in this area.

- ***KRYPTON : Integrating Terminology and Assertion : Brachman, Fikes and Levesque***

  ***In the proceedings of AAAI-88***

  A hybrid representation system is presented that combines in a completely integrated fashion a frame based description language and an assertional component that uses a first order resolution theorem prover. An update of this paper occurs in the proceedings of IJCAI-85.

- ***A Framework jot Representing Knowledge : Minsky***

  ***MIT-AI memo 594, 1974, also in the Psychology of Computer Vision, edited by Patrick Winston, MIT Press, 1975***

  The classic paper which started off work on frames. Skim through this paper, it has several interesting ideas. The appendix has a critique of the logic approach to Knowledge Representation.

- ***Frame Representations and the Procedural/Declarative controversy : Winograd***

  ***In Representation and Understanding, edited by Bobrow and Collins, Academic Press, 1975***

  Highly recommended for a clear understanding of this famous controversy.

- ***In Defense of Logic : P. J. Hayes***

  ***In the Proceedings of IJCAI5, 1977***

Recommended only for those with a deep interest in Knowledge Representation. Indicates that logic is the only knowledge representational formalism with a very well specified semantics in spite of its spartan syntax.

- ***Distinctions and Confusions : A Catalogue Raisonne : Israel, Brachman***
  ***In the Proceedings of IJCAI7, 1981***
  A tongue in cheek paper clarifying the issues on the semantic nets vs predicate calculus debate. Worth reading in full.

- ***An*** Overview ***of Production Systems*** **: Davis *and* King**
  ***in Machine Intelligence* 8, *1977,* edited *by edited by Elcock and* Michie**
  ***Edited and Reprinted in Buchanan and Shortliffe(eds), Rule Based Expert Systems***
  The definitive work on production systems. Explains what they are, and highlights when they are useful and when they are not.

- ***Semantic Net Representations* in *Rule Based Inference Systems: Duda, Hart and Nilsson***
  **in *Pattern Directed Inference Systems***
  This explains the knowledge representation mechanism in PROSPECTOR.

- ***SIGA RT special issue on KR, Vol. 70, 1980,* edited *by Brachman and Smith***
  **This** was a Knowledge Representation questionnaire sent out to AI practitioners of the day. The hope was that the editors would compile a perspective(s) which emerged from that survey. What resulted was more like the Tower of Babel. The questions are worth noting. Skim judiciously through the responses!

- ***Vision : Marr***
  Pages 19-29 have Marr's perspective on knowledge representation which is quite different from the conventional view in AI. These few pages are highly recommended.

- ***Reflection and* Semantics** in ***a Procedural Language : Brian Smith***

  ***MIT- TR-272, 1982***

  **This** is Brian Smith's PhD thesis describing the construction of a reflective Lisp called 3-Lisp. It is extremely long. The introduction and the first two chapters are highly recommended.

- ***Formal Theories of the Commonsense World : Hobbs and Moore, Ablex, 1985***

  This is a collection of work that focuses on ***what*** an intelligent agent needs to know to make its way in the real world. All the papers here are worth reading. Some that are especially interesting are

  - Naive Physics I : Ontology of liquids : P.J. Hayes

    The tricky problems encountered in representing and reasoning with our common sense knowledge of liquids are described here.

  - A Qualitative Physics based on Confluences : de Kleer and Brown

    Qualitative differential equations are used to model reasoning about the behavior of complex physical systems. This will allow a robot in **the real world to make** quick decisions about the outcome of events on the basis of incomplete **qualitative** information.

- ***Readings in Knowledge Representation : Brachman and Levesque, Morgan Kaujmann Publishing Co., 1985***

  An excellent collection of readings in knowledge representation, many of **the selections** there are in this **list.**

# 4 Planning, Problem Solving and Automatic Programming

Basic Reading

- ***Artificial Intelligence (2nd ed.), Patrick Winston, Addison- Wesley, 1984, Chapter 3,5,6,7***

  A recommended order for going through these chapters is 6,7,5,3. Chapter 6 is on problem solving paradigms. The most important section is the one on Generate and Test. Chapter 7 introduces resolution proofs, planning in the blocks world and problem solving by constraint propagation. This exposition is not at a very high level, but you have to know at least this much before you tackle the rest of the readings, also the examples given here are invaluable for understanding the material. Chapter 5 introduces control metaphors in problem solving. The section on means-end analysis and GPS is very well written. You may skip that section in the Handbook if you read this treatment. Chapter 3 is on constraint propagation. This book does not cover Automatic Programming. It has a limited coverage of planning as indicated above. Good exposition of problem solving paradigms.

- ***The Handbook of Artificial Intelligence, Volume 3, William Kaufmann Inc., 1982, Chapter 15, 10, 11 B and C***

  Chapter 15 is on planning and problem solving. The four systems STRIPS, AB-STRIPS, NOAH and MOLGEN should be learnt well. In case you get lost in the details of these systems, read Earl Sacerdoti's (see below) excellent review of problem solving tactics, which will help form a unified perspective in which to view these systems. GPS is covered in Chapter 11. section B. Skim through this if you have read Winston's exposition. Chapter 11, section C is on the Hayes-Roth and Hayes-Roth opportunistic planner. Chapter 10, sections A, B, C.l, C.4, C.5 and C.8 are relevant here. We will cover all Automatic programming work which can be viewed from the planning perspective - i.e transforming high level specifications into a program in a

given target language. Synthesizing programs from examples (this involves induction) will be covered under Learning (but note Manna and Waldinger's synthesis by deduction) .

- ***Principles of Artificial Intelligence, N.J. Nilsson, Tioga Publishing Co., 1980, Chapters 7 and 8***

  These two chapters cover STRIPS, ABSTRIPS and RSTRIPS in gory detail. Read them after the Handbook sections. Do the exercises at the end. They will stimulate thinking on the various issues in planning.

Required papers

***From*** Readings **in** Artificial ***Intelligence : Webber and Nilsson, 1982***

- ***Application of Theorem Proving to Problem Solving : Green, 1969***

  Indicates how a purely deductive approach based on first order logic can be used to generate robot plans. Can be seen as an attempt to build McCarthy's Advice Taker (see readings for Knowledge Representation). The answer extraction method is important. Nilsson's text explains this also. The examples in the Green paper are to be noted. Think about the problems with this approach to problem solving.

- ***The Frame Problem and Related Problems* in *AI : Hayes, 1978***

  Studies the frame problem which arises in the context of representation of actions. This paper is somewhat hard to read. It presents all the solutions to the frame problem that have been proposed in literature in a unified framework.

- ***Learning and Executing* Generalized *Robot Plans: Fikes, Hart and Nilsson, 1972***

  This is one of earliest systems that dealt with execution monitoring. Also note how STRIPS deals with the frame problem.

- ***Achieving Several Goals Simultaneously : Waldinger, 1977***

  This well written paper talks about goal regression, a method similar to passing a

condition back over an operator in program verification. Chapter 8 of **Nilsson's** text is on RSTRIPS which uses goal regression.

- ***Planning and Meta-Planning : Stefik, 1981***

  This explains hierarchical control of planning. Contrast the ideas here **with those in** the meta-level architecture proposed by Genesereth and Smith (reading on **Advanced** Topics). Stefik has another paper (see below) where he expounds **the other key idea** in MOLGEN – explicit representation of the interaction between **subproblems as con-**straints and constraint posting as a method of communication between the differ& levels **of planning.**

- ***An Experiment in Knowledge-Baaed Automatic Programming : Barstow, 1979***

  What you should get out of this paper, for now, is a sense of how the program **synthesis** problem can be tackled in the production system framework. This **work has been quite** successful in the real world.

- ***Patterns and Plans in Chess : Wilkins, 1980***

  The program PARADISE which generates chess plans using rule **based representations** of expert knowledge in chess tactics is described here. Read section 12 **(comparison** with robot planning) carefully.

Recommended Reading

- ***Problem Solving Tactics : Earl Sacerdoti***

  ***In AI magazine, 1981***

  This is a nice overview of planning and problem solving techniques. **Supplementary** reading to the presentation in the Handbook.

- ***Sciences of the Artificial : Simon***

  ***MIT Press, 1969***

  This little book has four of Simon's most beautiful essays. Simon's ant is **in the second**

essay. The piece entitled "The architecture of complexity" is highly recommended. Simon also speculates on how research in representations should **proceed.**

- ***Planning with Constraints – Molgen:Part 1 : Stefik***

  ***in AI Journal Vol 16, No 2, 1980***

  This is the constraint posting article referred to previously in this document.

- ***Design of a Programmer 'a Apprentice : Rich and Shrobe, 1979***

  ***In AI: An MIT Perspective, edited by Winston and Brown, 1979***

  This is an early description which indicates the input-output behavior of a Programmer's Apprentice and the kinds of reasoning mechanisms and programming knowledge needed to support it.

- ***Planning for Conjunctive Goals : Dave Chapman, 1985***

  ***MIT-AI- TR-802***

  This is a rational reconstruction of previous work on domain-independent **conjunctive** planning : the state of the art is distilled into the algorithm TWEAK that has been proved to be correct and complete.

- ***Reasoning about plans : Drew McDermott, 1985***

  ***In Formal Theories of the Commonsense World, edited by Hobbs and Moore, Ablex, 1985***

  A theory of actions in constructed in the framework of a previously developed **ax**iomatization of time. Planning is viewed as a kind of deduction. Critical aspects of actions like subtask, success, failure and feasibility are formalized.

- ***A Model of Naive Temporal Reasoning : Allen and Kautz, 1985***

  ***In Formal Theories of the Commonsense World, edited by Hobbs and Moore, Ablex, 1985***

  A formal account of time based on intervals is described. This allows reasoning **about**

time at various levels of granularity for problem solving and story understanding.

- **Qualitative** ***Process Theory : Forbus, 1985***

  ***MIT AI Lab PhD thesis, 1985***

  The knowledge and reasoning mechanisms needed to support commonsense process understanding is presented here.

- ***A*** **Perspective** ***on Planning : Stun Rosenschein***

  ***In the Proceedings of AAAI 84***

  The **BDI** model is described here. Planning becomes intention management in this **model.**

Also see papers by Balzer, Dershowitz, Smith, **Manna, Kant and Barstow in the proceedings** of IJCAI-85 to get a feel for recent research. Steier, Kant **and Newell [Kant and Steier 85]** work on algorithm design methods. Barstow's papers address practical automatic **programming issues.**

# 5 Deduction and Inference

## Basic Reading

- ***The Handbook of Artificial Intelligence, Volume 3, William Kaufmann Inc., 1982, Chapter 12***

  The overview is well written and should be read carefully. For details on **the historical** evolution, read the Loveland article in the recommended readings below. The **resolution** rule of inference is discussed very briefly in Section B. Supplement it **with** the material in Section 5.2 of Nilsson's text. Non-resolution theorem **proving should be** familiar to those who have read Manna's MTC text. The interesting thing to note here are the heuristics built into Automatic Theorem Provers (ATPs). A **more complete** list is in the Bledsoe article (below). The Boyer-Moore theorem prover is **important.** The brief overview on non-monotonic logics is a good introduction to newer research in the area(covered in the Recommended Reading below). It is a hot topic of research now. Logic programming is better covered in Kowalski's book (below).

- ***Principles of Artificial Intelligence, N.J. Nilsson, 1980, Tioga Publishing Co.,*** *Chapters 5 and 6*

  Relevant sections here are 5.2 [Resolution strategies with examples], 5.3 **[Simplification** – note especially procedural attachments which are discussed further in **Wehyrauch's** article], 5.4 [Answer extraction]

  Chapter 6 is long and the sections on forward and backward deduction systems can be skimmed. It is useful to go through the control strategies. The bibliographical remarks at the end of Chapter 6 are very important for a historical perspective.

## Required papers

from ***Webber and Nilsson's Readings in Artificial Intelligence***

- ***Non-Resolution Theorem Proving : Bledsoe, 1977***

  Bledsoe (a convert from resolution theorem proving) gives alternatives to resolution in ATP. For each note the main idea along with an example. Get a sense as to why each of the problems mentioned in Table 1 constitute challenges for present day ATP's. Read the bibliography to learn names of people in this enterprise.

- **Using *Rewriting Rules* for *Connection Graphs to Prove Theorems : Chang and Slagle, 1979***

  This is an attempt to speed up resolution theorem proving. Chang and Slagle pre-compile potential matches between clauses in their connection graph. Reading pages 219-222 of Nilsson before this might help. The exposition here is very clear, however.

- ***On Closed World Databases : Ray Reiter, 1978***

  The consequences of the closed world assumption in query evaluation in data bases is examined here. Examples and statements of theorems are important. Proofs are not necessary.

- ***A Deductive Approach to Program* Synthesis *: Manna and Waldinger, 1980***

  Program synthesis cast in a theorem proving framework. Combines techniques of unification, induction (cf. Boyer Moore theorem prover) and transformation rules (cf. PECOS) into a unified framework. See how they synthesize recursive programs.

- ***Prolegomena to a Theory of Mechanized Formal Reasoning : Weyhrauch, 1980***

  The important concepts here are semantic attachments and reflection. Semantic attachments provide a way of realizing the benefits of procedural and declarative representations in the same framework. Reflection principles and meta-theories are covered in greater detail in Wehyrauch's paper in'the 5th Conference on Automated Deduction Systems (Springer-Vcrlag Lecture notes in CS, no. 138).

- Subjective Bayesian ***methods*** for ***Rule Bused Inference Systems : Duda, Hurt and Nilsson, 1976***

  This is a short paper and the details here are important. Read the relevant sections of the **Buchanan/Duda** article (below) before you read this.

Recommended Reading

The readings below attempt to cover default and probabilistic reasoning, circumscription, NM **logics**, resource limited reasoning, logic programming. Additions to this set are welcome.

- ***Automated Theorem Proving: a Quarter Century Review : Loveland***

  ***in AMS Contemporary Mathematics v.29 "Automated Theorem Proving: After 25 Years"* edited *by Bledsoe and Loveland***

  This is the Loveland overview referred to above.

- ***Principles of Rule Bused Expert Systems : Buchanan and Duda***

  ***HPP-82-14***

  The sections on the handling of uncertainty in expert systems are relevant for now. It might be useful to read this before the **Duda** article in the required **reading.**

- ***On* Reasoning *by Default : Reiter***

  ***in TINLAP-C***

  Very short and highly readable introduction to default reasoning. Highly **recommended.**

- ***Computation and Deduction : P.J. Hayes***

  **in *the Czech. Acad. of* Sciences *proc., 1978***

  Argues for consideration of all computation as controlled deduction.

- ***Logic for Problem Solving : Kowalski, Elsevier, 1979***

  A well written, though wordy introduction to logic programming. **There** are **examples**

and issues to be gleaned from here. Compare this with PLANNER.

- ***Algorithm = Logic + Control : Kowalski***

  ***in CACM, Vol 22, No. 7, pages 424-436***

  Perspective : separates control information from logic information in programming (i.e separate the ***what*** **from the** ***how).***

- ***AI journal special issue on Nonmonotonic Logic, Vol 13, nos 1:2***

  This describes current efforts on non monotonic **reasoning. McCarthy's circumscription** paper (the follow up paper is listed below), **McDermott and Doyle's "Non-**monotonic logic I" (Moore's correction to **it is below), Reiter's paper on default logic,** Terry Winograd's article on resource-limited **reasoning as well as Wehyrauch's article** here are worth **reading.**

- ***Semantical Considerations on Nonmonotonic Logic : Moore***

  ***SRI Tech note 284, June 83***

  Corrections to the McDermott and Doyle version of Nonmonotonic **logic.**

- ***Circumscription : McCarthy***

  ***SAIL tech report, October 1989***

  Generalization of the 1980 paper, has a cleaned up version of the basic **formalism.**

- ***Default Reasoning as Likelihood Reasoning : Rich***

  ***in the Proceedings of the AAAI-83***

  Integration of default reasoning with probabilistic **reasoning.**

- ***Probabilistic Logic: N.J. Nilsson***

  ***In the AI Journal*** **(in** ***press)***

- ***Automated Reasoning:Introduction and Applications, Wos, Overbeek, Lusk and Boyle,***

  ***Prentice- Hall, 1984***

The preface has a reading guide. Chapters **1** thru 5 are introductory and can be skipped. Resolution strategies are presented with examples in Chapters 6 and **7.** Chapters 9 and 10 are highly recommended : they indicate how automated reasoning can be used to do discovery in mathematics. Chapter 16 indicates open problems in this area.

# 6 Expert Systems

The objective here is to learn about expert systems and the principles **they are based on** before studying the details of individual systems. It is recommended that the basic reading be done in the order given.

Basic Reading

- ***Principles of Rule Based Expert Systems : Buchanan and Duda, 1982***

  ***HPP-82-14***

  This paper discusses key issues in expert system design: representation, inference **and** uncertainty management. It has a large number of pointers to specific expert systems which one should follow up in the Handbook or in the book ***Building Expert Systems.*** It is primarily about rule based expert systems, though other representational frameworks are briefly discussed. The interested should look at the recommended readings from the MYCIN book which cover this exhaustively. Section 5 is extremely important for a well rounded understanding of expert systems. It lists all the key concepts and indicates the range of problems for which expert systems are useful.'

- ***New research on expert systems : Buchanan, 1981***

  ***HPP-81-l***

  The current state of the art and directions for future research are presented here. Lessons from early work and characteristics of domains for which present-day expert systems are suited are also indicated. There are many pointers to recent work **on** extending the capabilities of expert systems.

- ***Expert Systems : Paul Harmon and David King, Wiley, 1985, Chapter 9, 10***

  Chapter 9 is an excellent review of the early expert systems and complements the AI Handbook treatment(below). Chapter 10 covers the more recent systems R1, MOLGEN, The DIPMETER Advisor and EXPERT(some of these are not discussed in the AI Handbook).

- ***Building Expert systems, Hayes-Roth, Waterman and Lenat, Addison- Wesley, 1988, Chapters 1,2,3,4,5,7,8***

  This book is a how-to book and is quite easy to read. Chapters **4,5,7,8 are rec**ommended in particular. Lots of open problems remain in the areas addressed **by** Chapters 7 and 8.

- ***The Handbook of Artificial Intelligence, Volume 2, William Kaufmann Inc., 1982, Chapters 7,8,9***

  All the overviews should be read carefully. The overview section of **Chapter 7 could be** supplemented with the first chapter of "Building Expert Systems" **especially for the** historical evolution of expert systems. The details of the individual expert **systems in** these chapters should be read with a view to answering some of the questions below.

  - Associational questions - who, when, where, why.
  - Main ideas - representation, control, validation etc.
  - Concept-wise indexing - which systems explored concept X ?
  - A canonical example of the system's functioning which brings out its **limitations** and strengths.
  - Current status of the system.

The original reference for Dendral is in the Required papers below. The **Meta- Dendral** reference is in the reading list for Learning. The original reference for Crysalis is **Alan** Terry's thesis (HPP-83-19). For Macsyma, you could read Moses' Macsyma primer [Mathlab Memo 2, MIT]. To get more about symbolic integration read [Moses **1971:** Symbolic integration, the stormy decade: ACM **14,** pp 548-560]. There **are two papers** on Prospector in the Webber-Nilsson collection - one detailing the model design **aspect** and the other explaining the probabilistic belief updating scheme **used. Rl is not** covered here, so read the AI magazine articles on this subject. It is important to **not**

lose sight of the issues when wading through the details of the individual systems. This caveat extends to the whole of this reading list. Preparation of time lines and short summaries of systems while reading about them is highlyrecommended.

Required papers

***Readings in AI : Webber and Nilsson***

- ***Consultation Systems for Physicians : Ted Shortliffe, 1980***

  This proposes design criteria for clinical expert systems (like MYCIN) to increase their acceptability among physicians. Human engineering issues are important in the design of an expert system because the end users are often humans. This articles talks about the specific problems (mechanical, epistemological and psychological) that arise in the design of medical computing systems. A revised version of this article occurs as Chapter 5 of [Clancey84].

- ***An Experiment in Knowledge-Based Automatic Programming : David Barstow, 1979***

  PECOS forms a component of a full fledged and working AP environment at Kestrel Institute. Read this to get some details on the kinds of knowledge encoded in PECOS. Also note how PECOS derives the well-known trick of computing the intersection of two linked lists in linear time.

- ***Dendral and Meta-Dendral : Their Applications Dimension : Buchanan and Feigenbaum, 1978***

  Very important paper. Has succinct descriptions of Dendral and Meta-Dendral.

- ***Model Design in the PROSPECTOR Consultant System for Mineral Exploration : Duda, Hart and Nilsson, 1979***

  This is a description of the PROSPECTOR system which uses inference networks of geological assertions and the Bayesian propagation formalism to model the judgmental reasoning of economic geologists. The model design process illustrates the transfer of knowledge from human experts to computational formalisms.

- ***The Hearsay II Speech- Understanding System : Erman et al., 1980***

  This is a good example of application driven AI research that led to the **development** of a new problem solving architecture. The BB1 architecture [Hayes-Roth84] is a substantially extended and generalized version **of** this system.

- ***Interactive Transfer of Expertise : Acquisition of New Inference Rules : Randall Davis, 1979***

  One approach to the KA bottleneck is explored here : this is to partially **automate the** role of the intermediary knowledge engineer, by giving the system knowledge about the form of the rules and making it acquire knowledge from the expert in the context of an error in the reasoning made by the **system.**

Recommended Reading

- ***Expert Systems: Working Systems and the Research Literature: Bruce Buchanan KSL-85-87***

  This is a compilation of work on expert systems upto 1985, with an emphasis on systems actually working. The defining characteristics of expert systems **are important.** Also scan the Expert System papers published in the proceedings of **IJCAI and AAAI** in recent years.

- ***Expert Systems : Where are we, where are we going? : Randall Davis in AI magazine 1982***

  This article talks about the current state of the art in expert systems and **indicates** the deficiencies in them. It then proposes ways of overcoming them by the use of causal models and reasoning from first principles. This is extremely well written **and** it is highly recommended.

- ***ROGET - A KB consultant for acquiring the conceptual structure of expert systems : Bennet***

*HPP-83-24*

This attempts to attack the most difficult part in the design **of an expert system** – obtaining the vocabulary and inference structure. The approach **is to treat this** problem as a classification problem (Clancey 83) and there are 9 problem classes identified. Only one (for diagnosis) of them is implemented.

- ***Rule Baaed Expert Systems : Buchanan and Shortliffe***

  **Addison** *Wesley, 1984*

  This is the MYCIN book. Chapters 21-24 cover other representational **frameworks for** expert systems. **A** good introduction to Knowledge Engineering **is in Chapter 7. All** you ever wanted to know about uncertainty management in MYCIN is in Chapters 10-13. Chapter 17 expounds on explanation and its role in AI research.

- ***The Epistemology of a Rule-Based Expert System : a Framework for Explanation :*** *W. J. Clancey*

  *HPP-81-17*

  Indicates the shortcomings of the production rule framework for teaching **and expla-** nation. Indicates the kind of knowledge that is embedded in rules that **have to be** made explicit for the purpose of tutoring or explanation.

- ***Maxims for Knowledge*** **Engineering :** ***David Barstow and Bruce Buchanan***

  ***HPP-81-4***

  This lists some principles (and folklore) in the art and craft of Knowledge **Engineering.**

- ***The advantages of abstract control knowledge in expert system design : William*** *Clancey*

  ***HPP-83-17***

  Abstract control knowledge makes the design of an expert system **more transparent** and explainable. It also provides a generic framework for constructing KB's for **related** problems in other domains. It is a useful starting point for the study of **strategy. This**

paper has examples of abstract control knowledge from NEOMYCIN. The difficulties in gathering such knowledge is stressed.

- ***The Blackboard Architecture : a general framework ⇗□□ Problem solving ? : Barbara Hayes-Roth***

  ***HPP-88-38***

  This is an attempt to define and evaluate the blackboard architecture independent of a particular application. It clearly defines the components of a blackboard system and enumerates the assumptions behind these components. This is of relevance here because it provides a framework for bringing many knowledge sources to bear on **a** problem. This is more flexible than a rule based expert system, but for computational efficiency we need compilation (cf HARPY).

- ***Rule based Understanding of Signals:* Nii *and Feigenbaum***

  ***In Pattern Directed Inference Systems, edited by Waterman and Hayes-Roth***

  This is a signal processing application that uses the blackboard architecture. It is the root of the AGE project at Stanford (Nii).

- ***Rl : A rule based configure? of computer systems: J. McDermott***

  ***AI Journal, September 82***

  The R1 paper. R1 is an expert system that configures VAXen for DEC. It is **a** production system that uses forward chaining and is probably the most widely used expert system in existence.

- ***An* overview *of Meta-Level Architecture : Genesereth and Smith***

  MRS is a realization of the Advice Taker (McCarthy58) with explicit control of **reason-**ing. It is a logic based system that provides a variety of representation and inference methods for the implementation of expert systems.

- ***Heuristic Methods for Imposing Structure on Ill-Structured problems : Pople***

***In AI* in** Medicine, Szolovits ***(ed), AAAS Selected Symposium 51, Westview* Press, *1982***

This addresses the nature of the clinical reasoning process and then offers a description, critique and a status report on INTERNIST(CADUCEUS).

- ***AI Journal, Special* issue *on applications in the* Sciences *and* medicine**

  ***Volume 11, No. l-2, Aug 1978***

  There are many interesting papers in this collection. Skim through this to get **an idea** of the kind of work going on and the main research issues. The following papers are recommended.

  1. The plan recognition problem : an intersection of psychology and **AI.** By Schmidt, Sridharan and Goodson

  This describes the BELIEVER system which encompasses **a** psychological theory **of** how humans understand actions of experts.

  **2.** A model based method for computer aided medical decision making : **Weiss,** Kulikowski, Amarel and Safir

  This is the CASNET paper.

  3. Categorical and probabilistic reasoning in medical diagnosis : Peter Szolovits, **Pauker**

  PIP, INTERNIST, CASNET and MYCIN are compared here.

# 7 Learning

Basic Reading

- ***The Handbook of Artificial Intelligence, Volume 3, William Kaufmann Inc., 1982, Chapter 14***

  This is a comprehensive introduction to work on learning. It is a pre-requisite to understanding any of the subsequent material. This survey covers pre-1982 research only. For more recent work read the Machine Learning book referenced below.

Required papers

***from the Webber and Nilsson collection of AI readings***

- ***Generalization as Search : Mitchell, 1981***

  Casts the generalization problem as a search problem. It provides **a framework for** comparing the various generalization strategies in learning literature.

- ***Interactive Transfer of Expertise : Acquisition of New Inference Rules : Randy Davis, 1979***

  **TEIRESIAS** is described here. The most important contribution **is the acquisition of** knowledge in the context of a bug and the use of rule models obtained by induction on the existing set of rules. These are well explained in Section 7. Section 9 highlights the limitations of this approach and indicates open research questions.

Recommended reading

- ***The Computers and Thought Award Lecture : IJCAI 81 : Mitchell***

  ***in the Proceedings of IJCAI 81***

  Mitchell presents the problem of bias in generalization and the **need** for justified generalizations, with examples from LEX. Open problems in learning **are indicated.** Highly recommended.

- ***Models of Learning Systems : Buchanan, Mitchell, Smith, Johnson***

  ***STAN-CS-79-692, also in Encyclopedia of Computer Science and Technology, Dekker, Vol 11,1978***

  This model of learning systems has proved extremely useful for understanding as well as designing learning programs. One of the components in this model (the Critic) was extremely well analyzed by Dietterich and Buchanan( see reference below).

- ***An overview of Machine learning : Carbonell, Michalaki and Mitchell***

  ***in Machine learning : an AI approach, 1983, Tioga Publishing Co.***

  **A** short sketch of current work in machine learning. Also includes a historical survey and a guide to the work presented in the later chapters in the book. Highly recommended.

- ***Why should machines learn? : Simon***

  ***in Machine learning : an AI approach, 1983, Tioga Publishing Co.***

  **A** short article giving Simon's views on machine learning. Every serious **machine** learning researcher should attempt to answer the questions that Simon raises. He **plays** devil's advocate here and concludes that with the exception of cognitive modeling, some rethinking of long term objectives in machine learning are in order.

- ***Learning by Experimentation : Acquiring and Refining Problem Solving Heuristics : Mitchell, Utgoff, Baneji***

  ***in Machine Learning : an AI approach, 1983, Tioga Publishing Co.***

  The LEX system is described. The most interesting section is 6.4 – where new term generation is discussed.

- ***The role of the critic in learning systems : Dietterich and Buchanan, 1981***

  ***STAN-U-81-891***

  ***Also in Adaptive Control of Ill-defined Systems: Selfridge, Rissland and Arbib (eds)***

This is a nice analysis of the task performed by the critic in a learning system. The styie of analysis is as important as the contents here.

- ***Model directed* learning *of production rules : Buchanan and Mitchell***
  **in *Pattern Directed Inference Systems, ed. Waterman and Hayes-Roth, 1978***
  The Meta-Dendral paper.

- ***The role of heuristics* in *learning by* discovery : *Lenat***
  **in Machine *learning : an AI approach,* 1983, Tioga Publishing Co.**
  Three case studies are presented here : AM, EURISKO and a speculation that nature **adopts a heuristic learning *by* discovery** approach to evolution. .

- ***Toward chunking as a general learning mechanism : Laird, Rosenbloom, Newell***
  ***in AAAI-84***
  Chunking is presented as a general learning mechanism within a general problem solving architecture, SOAR. It is shown to exhibit both the practice effect as well as **strategy acquisition.**

Learning Apprentices[Mitchell84] form the major thrust of current work on machine learning. Analogy is a hot topic too: work by [Winston84], [Kling71], [Anderson83], [Gentner83], [Carbonell83] and [Greiner85] are worth reading. Genetic algorithms[Holland75], and work on incremental learning and experimentation[Dietterich84] are also being explored. The most recent work in machine learning is to be found in the proceedings of the Machine Learning Workshops (1983 and 1985). A new Machine learning journal will be published starting 1986 and will contain current research in the field.

# 8 Natural Language Understanding

This collection of readings covers both natural language understanding and speech recognition. Thanks go to John Lamping for helping me set up the syllabus for the **Natural** Language understanding **section.**

Basic Reading

- ***The Handbook of* Artificial *Intelligence, Volume 1, William Kaujmann Inc., 1982, Chapter 4***

  This chapter should be read fairly carefully. Those interested in **details of grammars** and parsing should read Winograd's book **on the syntax of natural language (see** below). This chapter does not cover semantics and should **be augmented with Wino-**grad's paper "What does it mean to understand **language?". To get a picture of what a** present day Natural Language front-end looks like, read the Martin **paper below.**

- ***The Handbook of Artificial Intelligence, Volume 1, William Kaujmann Inc., 1982, Chapter 5***

  This is a fairly detailed account of the major speech understanding **projects of the** 70's. Compiling a comparison of these systems is instructive and is **suggested as an exercise.**

Required papers

- ***The Hearsay II Speech Understanding system : Erman et al., 1980***
  ***in the Webber-Nilsson collection***

  This describes the HEARSAY architecture - a general framework **for coordinating** independent processes to achieve cooperative problem solving behavior in the **face** of uncertainty. The details of the different levels are unimportant, so skim **through the** extended example appropriately. The comparison with other **systems is interesting.** The section on performance is extremely important. Read the conclusions **carefully.**

- ***What does* it *mean to understand language? : Winograd***

  ***in* Cognitive *Science, Vol 4, 1980, pp 209-241***

  Winograd explains the difficulties in understanding natural language. It is historically organized and represents Winograd's shifts in his thinking about natural language.

- TEAM : Paul Martin et. **al.**

  in the proceedings of IJCAI 83

  This should be read thoroughly for an idea of **a** modern natural language front end for data bases.

Recommended Reading

- ***A Semantic Process for* Syntactic *Disambiguation : Hirst***

  ***in the Proceedings of AAAI 84***

  This should be scanned for understanding how semantic processing can work and how it can be combined with syntactic processing.

- ***Computational models of discourse : Brady and Berwick***

  ***MIT Press***

  The first part of the introduction to this book (pages 27-37) should be read to **get a** feel for research issues in discourse.

- ***Language as a* cognitive *process : Vol I : Syntaz : Winograd***

  ***Addison- Wesley, 1989***

  Read Chapter 1 and Chapter 7 and skim through the rest.

- ***Telling your computer to recognize speech not wreck a nice beach : Baker***

  ***IBM Research report RC 6935, No. 29665***

  This is probably the 'cutest' title for an AI paper. The contents should be skimmed over.

- ***Utterance and* Objective : *Grosz***

  **in *AI magazine : Vol 1 , No.1 , Spring 1980***

  The central issues of natural language from a communication perspective are described very clearly here. An updated version of this is to be found in the proceedings of IJCAI-85.

- ***Elements of a Plan-Based Theory of Speech Acts : Cohen and Perrault, 1979***

  ***In the Webber and Nilsson collection***

  Language use is modeled by viewing speech acts as operators in a planning system thus allowing both physical and speech acts to be integrated into plans.

- ***Computers and* Cognition : *Winograd and Flores, Ablex, 1986***

  An interesting perspective on Artificial Intelligence in general and language understanding in particular. The critique of the expert systems approach (Chapter **10)** is food for thought.

# 9 Vision and Robotics

Basic Reading

- *The Handbook of* **Artificial** *Intelligence, 1982, William Kaujmann Inc., Volume 3, Chapter 13*

  The overview should be read carefully. The blocks world work in **Section B can be** skimmed. The main ideas in these systems are indicated **in the handout on vision** systems in the course notes for CS229b. Edge detecting and line finding as **well** as analysis of texture are important research issues. The Brady **paper gives more** details on this. The various shape-from methods in Section **D are important. The** presentation is better in Brady. Read the section on relaxation algorithms **carefully.** Vision systems should be understood well, particularly ACRONYM.

- *Artificial Intelligence (2nd ed.), Patrick Winston, 1984, Addison- Wesley, Chapter 10*

  This is a nice account of some of the more recent work in vision **and is a pre-requisite** to understanding the Brady paper.

- ***Computational Approaches to Image Understanding : Brady***

  in ***Computing Surveys, Vol. 14, No. 1, March 1982***

  A survey of the recent developments in Image Understanding. The **first part of the** paper identifies the field of image understanding by delineating it from pattern **recognition,** image processing and computer graphics. This is followed by a description **of** the common themes in IU research that have crystallized over the **past decade. Then** there is a very nice review of work in geometrically simple **microworlds (Roberts,** Huffman and Clowes, Mackworth, Kanade, Turner, Barrow and **Tenenbaum, Waltz).** Operations on the image : edge detection, shape from shading, **segmentation and** texture are described. Shape from stereo, contour and motion are surveyed next. **The** paper ends with a survey of work on representing 3D objects.

## Recommended Reading

- ***Task Planning : Tomas Lozano- Perez***

  ***In Robot Motion : Planning and Control, eds Brady et al., MIT Press, 1982***

  This is a primer on task planning. Task planning is divided into three phases : modelling, task specification and manipulator program synthesis. Note similarities and differences in the approaches to specification and synthesis with respect to automatic programming and planning.

- ***Survey of Model Based Image Analysis Systems : Binford***

  ***SAIL Tech. report***

  This paper surveys and critiques the state of the art in model based image analysis systems. It also describes principles of design of general vision systems.

- ***Perceptual Organization as a Basis for Visual* Recognition *: Lowe and Binjord***

  in ***the Proceedings of AAAI 83.***

  Evidence is presented that bottom-up grouping of image features is usually prerequisite to the recognition and interpretation of images. Several principles are hypothesized for determining which image relations are to be formed. Using this, **a curve** segment ation algorithm is presented.

- ***Automated Visual Inspection : Chin et al.***

  ***in the IEEE Transactions of Pattern Analysis and Machine Intelligence, .***

  Skim through this paper to get a sense of the kinds of applications in visual inspection in industry.

- ***AI Journal, Special Issue on Computer Vision : Aug 1981***

  The articles worth reading here are

  - The Preface by Michael Brady. This was probably what developed into the paper described above.

- Inferring Surfaces from Images : **Binford**
- Symbolic Reasoning among 3-D models and 2-D images : Brooks

Read the introductions (or skim through the contents of) the remaining articles which should be taken as representative of current vision research.

# 10 Advanced Topics

This is a pot-pourri of topics which are necessary for a well-rounded understanding of AI.

- Advanced reasoning/planning
  - Circumscription : a form of Nonmonotonic Reasoning, McCarthy, 1980

    *In Readings in AI, Webber and Nilsson (eds), 1982*
  - Reasoning about Knowledge and Action : Moore, 1977

    *In Readings in AI, Webber and Nilsson (eds), 198%*
  - A Truth Maintenance System : Doyle, **1979**

    *In Readings in AI, Webber and Nilsson (eds), 1982*
  - Assumption Based Truth Maintenance : deKleer, **1984**

    *Xerox PARC Technical Note*
  - Computation and Deduction : Hayes

    *Proceedings of the Czech. Academy of Sciences, 1973*
  - Design of a Programmer's Apprentice : Rich and Schrobe, **1979**

    *In AI : An MIT Perspective, Winston and Brown (eds), MIT Press, 1979*
  - Qualitative and quantitative reasoning in classical mechanics : deKleer, 1979

    *In AI : An MIT Perspective, Winston and Brown (eds), MIT Press, 1979*
  - On Inheritance Hierarchies with Exceptions : Etherington **and Reiter, 1984**

    *In the* **Proceedings** *of the AAAI-84*
  - Nonmonotonic reasoning : Genesereth and Nilsson, **1985**

    *In their forthcoming book on AI*
  - Learning by Experimentation : Acquiring and Refining Problem Solving Heuristics : Mitchell, Utgoff, Banerji, 1983

    **in** *Machine Learning : an AI approach, 1983,* **Tioga** *Publishing Co.*

- Common knowledge : Halpem and Moses, 1984

  ***IBM Tech. Report, 1984***

- Advanced **KR/Software** architectures

  - The Logic of **Frames** : Hayes, 1979

    ***In Readings in AI, Web ber and Nilsson (eds), 1982***

  - Meta-level architecture : Genesereth and Smith, 1983

    ***In the proceedings of AAAI-83***

  - Towards Chunkine as a General Learning Mechanism: Laird, Rosenbloom and Newell, 1984

    ***In the Proceedings of AAAI-84***

  - A Blackboard Model of Control : Barbara Hayes-Roth, 1985

    ***AI Journal, 1985***

  - OPS, A Domain Independent Production System Language, 1977

    ***In the Proceedings of IJCAI-77***

- Advanced architectures for AI

  Thanks to Vineet Singh for extending and annotating this list. Notation used below: I : introduction to issues, R : resource allocation, L : language, H : hardware and architecture.

  - What should AI want from supercomputers : Doyle. (I)

    ***AI Magazine, Vol 4, No. 4, Winter 1983***

  - The Connection Machine : Dan **Hillis** (H)

    ***MIT AI Lab memo, AIM 646***

  - DADO: a tree structured architecture for production systems : Stolfo and Shaw **(H)**

    ***In the Proceedings of IJCAI-83***

- A Variable Supply Model for Distributing Deductions :· Singh and Genesereth (R)

  ***In the Proceedings of IJCAI-85, also HPP-84-14***

- &LAMBDA : Gabriel and McCarthy. (L)

  ***STAN-CS-84-1007***

- A subset of Concurrent Prolog and its Interpreter : Shapiro (L)

  ***TR-003, ICOT, Japan***

- Hardware and Software Architectures for AI : Deering (I)

  ***In the Proceedings of AAAI-84***

- The Architecture of the FAIM-1 Symbolic Multiprocessing System : Davis and Robison (H,L,R)

  ***In the Proceedings of IJCAI-85***

- NETL : A System for Representing and using Real-World Knowledge: Fahlman, 1979

  ***MIT PhD thesis, MIT Press, 1979***

- **Philosophical/Epistemological** issues

  - Some Philosophical Problems from the Standpoint of AI : McCarthy, 1969

    ***In Readings in AI, Webber and Nilsson (eds), 1982***

  - Epistemological problems of AI : McCarthy, 1977

    ***In* Readings in *AI, Webber and Nilsson (eds), 1982***

  - Programs with Common Sense (The Advice Taker): McCarthy

    ***In Semantic Information Processing,, Minsky (ed), 1968***

  - The Second Naive Physics Manifesto : Hayes, 1985

    ***In Formal Theories of the Commonsense World, Hobbs and Moore (eds), Ablex, 1985***

- Ascribing Mental Qualities to Machines : McCarthy, 1979

  ***STAN-CS-79-725***

- A perspective on planning : Stan Rosenschein

  ***in the proceedings of AAAI-84***

- AI meets Natural Stupidity : Drew McDermott, 1976

  ***SIGPLAN Newsletter, Number 57, pages 4-9***

  ***also in Mind Design edited by J. Haugeland, MIT Press, 1981***

- Minds, Brains and Programs : John Searle, 1981

  ***in Mind Design edited by J. Haugeland, MIT Press, 1981***

- What Computers Can't Do: A Critique of Artificial Reason : **Dreyfus,** 1972

  ***Harper and Row, 1972***

- ***Minds and Machines: Anderson***

  ***Prentice- Hall, 1964***

- Computer Power and Human Reason: Weizenbaum, 1976

  ***W. H. Freeman, 1976***

• Directions for AI

- AI : Engineering, Science or Slogan? : **Nilsson,** 1981

  ***AI Magazine, Vol 3, No 1, Winter 1981/1 982***

- AI prepares for 2001 AD : **Nilsson,** 1983

  ***AI Magazine, Vol 4, No 4, 1983***

- The Current State of AI: One Man's Opinion : Schank, 1983

  ***AI Magazine, Vol 4, No 1, Winter-Spring 1983***

- The Nature of AI: A Reply to Schank : Bundy, 1983 ***AI Magazine, Vo14, No 4, Winter 1983***

# 11 Acknowledgements

We would like to thank Prof. Nils Nilsson who patiently read through innumerable drafts of this reading list and made constructive suggestions. Thanks also to Chris Fraley, Vineet Singh and Stuart Russell for their comments on a very early version of this document. Members of CS229b in this and in the previous year (Winter 84, TA'ed by Ben Grosof), provided considerable input.

# Bibliography


[Allen 85) Allen, J. and Kautz.
A Model of Naive Temporal Reasoning.
In *Formal* Theories *of the Commonsense* World, . Ablex Press, 1985.

[Amarel 82] Amarel, Saul.
On Representation of Problems of Reasoning about Actions.
In Webber, B.L. and Nilsson, N.J. (editors), Readings in *Artificial Intelligence, .* Tioga Press, Palo Alto, CA, 1982.

[Anderson 83] Anderson, JR.
Acquisition of Proof Skills in Geometry.
In Michalski, R., Carbonell, J. and Mitchell, T. (editors), Machine *Learning,* . Tioga Press, 1983.

[Anderson, Alan 64]
Anderson.
*Minds and Machines.*
Prentice-Hall, 1964.

[Baker 80] Baker, Janet M.
*How to tell your computer to recognize speech, not wreck a nice beach.*
Technical Report RC 6935, IBM Yorktown Heights, 1980.

[Barr, et al. 82] Barr, Avron, Paul R. Cohen and Edward A. Feigenbaum.
*The Handbook of Artificial Intelligence, Volumes I, II, and III.*
William Kaufmann, Inc., Los Altos, CA, 1981 and 1982.

[Barrow 81] Barrow, H. G. and J. M. Tenenbaum.
Computational Vision.
In *IEEE.* IEEE, 1981.

[Barstow 81] Barstow, David R. and Bruce G Buchanan.
*Maxims for Know/edge Engineering.*
Technical Report HPP-81-4, HPP, Stanford University, May, 1981.
Reprinted as Stanford Comp Sci Memo.

[Barstow 82] Barstow, D.
A Experiment in Knowledge-Based Automatic Programming.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artficial Intelligence, .* Tioga Press, CA, 1982.

[Bennett 83] Bennett, J.S.
*ROGET: A knowledge based system for acquiring the conceptual structure of an expert system.*
Technical Report HPP-83-24, Heuristic Programming Project, Stanford University, May, 1983.

[Berliner 82] Berliner, Hans.
The B* Tree Search Algorithm: A best first proof procedure.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence, .* Tioga Press, Palo Alto, CA, 1982.

[Berwick 83] Berwick, Robert C.
Introduction: Computational Aspects of Discourse.
In J. Michael Brady and Robert C. Berwick (editors), *Computational Models of Discourse, .* MIT Press, Cambridge, MA, 1983.

[Binford 81] Binford, Thomas 0.
Inferring surfaces from images.
*Artificial Intelligence*, August, 1981.
Special Issue on Computer Vision.

[Binford 82] Binford, Thomas 0.
*A survey of model based image systems.*
Technical Report, Stanford Al Lab., 1982,

[Bledsoe 82] Bledsoe, W.W.
Non-resolution Theorem Proving.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence, .* Tioga Press, CA, 1982.

[Boden 77] Boden, Margaret A.
*Artificial Intelligence and Natural Man.*
Basic Books, Inc., New York, 1977.

[Bolles 77] Bolles, R.
Verification Vision for Programmable Assembly.
In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence,* pages 569-575. 1977.

[Brachman 80] Brachman, R.J., and B. Smith.
Special Issue on Knowledge Representation.
*SIGART 50,* 1980.

[Brachman 83] Brachman, Ronald J., Hector J. Levesque, and Richard E. Fikes.
KRYPTON: Integrating Terminology and Assertion.
In *Proceedings of the National Conference on AI.* AAAI, Washington, DC., August, 1983.

[Brady 82] Brady, M.
Computational Approaches to Image Understanding.
*Computing Surveys* 14(1):2-78, March, 1982.

[Brooks 81] Brooks, Rodney.
Symbolic Reasoning among 3D models and 2D images.
*Artificial Intelligence* , August, 1981.
Special Issue on Computer Vision.

[Buchanan 78a] Buchanan, B. G., Mitchell, T. M., Smith, R. G. and Johnson, C. R. Jr.
Models of Learning Systems.
*Encyclopedia of Computer Science and Technology* 11:24-51, 1978.
also Stanford report STAN-CS-79-692.

[Buchanan 78b] Buchanan, B. G. and Mitchell, T. M.
Model-Directed Learning of Production Rules.
In Waterman, D. A. and Hayes-Roth, F. (editors), *Pattern-Directed Inference Systems,* pages 297-312. Academic Press, New York, 1978.

[Buchanan 82a] Buchanan, B.G.
New research on expert systems.
In J.E. Hayes, Donald Michie, and Y-H Pao (editors), *Machine Intelligence 10,* pages 269-299. Ellis Horwood Limited, Chichester, England, 1982.
Originally published as a Techreport in 1981.

[Buchanan 82b] Buchanan, B. and R. Duda.
*Principles of Rule-Based Expert Systems.*
Technical Report HPP-82-14, Department ofComputer Science, Stanford University, 1982.
Also in Yovits, M. (ed): Advances in Computers, Vol 22, Academic Press, 1983.

[Buchanan 82c] Buchanan, B.G. and Feigenbaum, E.A.
Dendral and Meta-Dendral : Their Applications Dimension.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence, .* Tioga Press, CA, 1982.

[Buchanan 84] Buchanan, BG. and E.H. Shortliffe.
*Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project.*
Addison-Wesley, Reading, MA, 1984.

[Buchanan 85] Buchanan, Bruce G.
*Expert Systems: Working Systems and the Research Literature.*
Technical Report KSL-85-37, Knowledge Systems Laboratory, Stanford University, 1985.
To appear in Expert Systems.

[Carbonell 83] Carbonell, J. G.
Learning by Analogy: Formulating and Generalizing Plans from Past Experience.
In Michalski, Ryszard S., Jaime G. Carbonell, and Tom M. Mitchell (editors), *Machine Learning, .* Tioga, Palo Alto, CA, 1983.

[Chang 82] Chang, C. and Slagle, J.
Using rewrite rules for connection graphs to prove theorems.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence,.* Tioga Press, CA, 1982.

[Chapman 85] Chapman, D.
*Planning for Conjunctive Goals.*
Technical Report MIT-AI-TR-802, MIT, 1985.

[Chin 82] Chin et. al.
Automated Visual Inspection.
*IEEE Transactions of Pattern Analysis and Machine Intelligence ,* 1982.

[Clancey 83a] Clancey, William J.
The Advantages of Abstract Control Knowledge in Expert System Design.
In *Proceedings of the 3rd National Conference on A/,* pages 74-78. Washington, DC., August, 1983.

[Clancey 83b] Clancey, W. J.
The Epistemology of a Rule-Based Expert System: A Framework for Explanation.
*Artificial Intelligence* 20(3):215-251, May, 1983.

[Clancey and Shortliffe 84]
Clancey, W.J., and E.H. Shortliffe.
*Readings in Medical Artificial Intelligence: The First Decade.*
Addison-Wesley, Reading, MA, 1984.

[Cohen 82] Cohen, Phil and Perrault, R.
Elements of a Plan-Based Theory of Speech Acts.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence, .* Tioga Press, CA, 1982.

[Davis 77] Davis, R. and J. King.
An Overview of Production Systems.
In Edward W. Elcock and Donald Michie (editor), *Machine Intelligence 8: Machine Representation of Know/edge,* pages 300332. Ellis Horwood Ltd., Chichester, England, 1977.
Also reprinted with changes in Buchanan and Shortliffe 1984.

[Davis 82a] Davis, R.
Expert systems: Where are we? And where do we go from here?
*The A/ Magazine* 3(2):1-22, spring, 1982.

[Davis 82b] Davis, R.
Interactive Transfer of Expertise : Acquistion of New Inference Rules.
In *Readings in Artificial Intelligence,* . Tioga Press, CA, , 1982.

[Davis 85] Davis, A. and Robison.
The Architecture of the FAIM-1 Symbolic Multiprocessing System.
In *Proceedings of IJCAI-85.* 1985.

[de Kleer 79] de Kleer, J.
Qualitative and quantitative reasoning in classical mechanics.
In P. H. Winston and R. H. Brown (editors), Artificial *Intelligence: An MIT Perspective,* pages S-30. The MIT Press, Cambridge, 1979.

[Deering 84] Deering.
Hardware and Software Architectures for Al.
In *Proceedings of the AAAI-84.* 1984.

[deKleer 85a] deKleer, J. and Seely Brown, J.
A Qualitative Physics based on Confluences.
In Hobbs, J. and Moore, R.C. (editors), *Formal Theories of the Commonsense World,* . Ablex Press, 1985.

[deKleer 85b] deKleer, J.
*Assumption-based truth maintenance.*
Technical Report, Xerox PARC, 1985.

[Dietterich 81] Dietterich, T. G. and Buchanan, B. G.
*The Role of the* Critic *in Learning Systems.*
Report HPP-81-19, Stanford University, December, 1981.
also STAN-CS-81.891.

[Dietterich 84] Dietterich, T.G.
*Theory Driven Data Interpretation.*
PhD thesis, Stanford University, 1984.

[Doyle 82] Doyle, J.
A Truth Maintenance System.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence, .* Tioga Press, CA, 1982.

[Doyle 83] Doyle, Jon.
What Should Al Want From the Supercomputers?
Al *Mag.* 4(4), winter, 1983.

[Dreyfus 72] Dreyfus, H. L.
*What Computers Can't Do: A Critique of Artificial Reason.*
Harper & Row, New York, 1972.

[Duda 78] Duda, R. O., Hart, P. E., Nilsson, N. J., and Sutherland, G. L.
Semantic Network Representations in Rule-Based Inference Systems.
In Donald A. Waterman and Frederick Hayes-Roth (editors), *Pattern-Directed Inference Systems,* pages 203-221. Academic Press, New York, 1978.

[Duda 82a] Duda, R., Hart, P. and Nilsson, N.J.
Subjective Bayesian Methods for Rule-Based Inference Systems.
In Webber, B.L. and Nilsson. N.J. (editors), *Readings in Artificial Intelligence, .* Tioga Press, CA, 1982.

[Duda 82b] Duda, R., Gaschnig, J. and Hart, P.
Model Design in the Prospector Consultatnt System for Mineral Exploration.
In *Readings in Artificial Intelligence, .* Tioga Press, CA, , 1982.

[Erman 82] Erman, L.D., Hayes-Roth, F., Lesser, V.R., and Reddy, D.R.
The Hearsay-II speech-understanding system: Integrating knowledge to resolve uncertainty.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence.* Tioga Press, CA, 1982.

[Etherington and Reiter 83]
Etherington, David W., and Raymond Reiter.
On Inheritance Hierarchies With Exceptions,
In *Proceedings of the National Conference on Al.* AAAI, Washington, D.C., August, 1983.

[Fahlman 79] Fahlman, S.E.
*A system for representing and using real world knowledge.*
PhD thesis, MIT, 1979.

[Feigenbaum and Feldman 63]
Feigenbaum, Edward A. and Julian Feldman. (editors).
*Computers and Thought.*
McGraw-Hill Book Company, New York, New York, 1983.

[Fikes 82] Fikes, R., Hart, P. and Nilsson, N.J.
Learning and Executing Generalized Robot Plans.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence, .*
Tioga Press, CA, 1982.

[Forbus 85] Forbus, K.D.
*Qualitative Process Theory.*
PhD thesis, MIT, 1985.

[Forgy 79] Forgy, Charles L.
*On the Efficient Implementation of Production Systems.*
PhD thesis, Carnegie-Mellon University, Dept. of Computer Science, January, 1979.

[Forgy and McDermott 77]
Forgy, C., and McDermott, J.
OPS, A Domain-Independent Production System Language.
*Proceedings of the Fifth International Joint Conference on Artificial Intelligence*
1:933-939,1977.

[Gabriel 84] Gabriel, R. and McCarthy, J.
*QLAMBDA.*
Technical Report STAN-CS-844007, Stanford CS Department, 1984.

[Gaschnig 82] Gaschnig, John.
A problem similarity approach to devising heuristics.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence, .*
Tioga Press, Palo Alto, CA, 1982.

[Genesereth 83] Genesereth, Michael R.
An Overview of Meta-Level Architecture.
In *Proceedings of the AAAI.* AAAI, Washington, DC., August, 1983.

[Genesereth 85] Genesereth, MR. and Nilsson, N.J.
*Fundamentals of Artificial Intelligence.*
Tioga Press, CA, 1985.

[Gentner 83a] Gentner, Dedre.
Structure- mapping: A Theoretical Framework for Analogy.
*Cognitive Science* 7(2), *1983.*

[Gentner 83b] Gentner, Dedre, and Albert L. Stevens (eds.).
*Mental Models.*
Lawrence Erlbaum Associates, Hillsdale, NJ, 1983.

[Georgeff 83] Georgeff, M.P.
Strategies in Heuristic Search.
*Artificial intelligence* 20(4):393-425, July, 1983.

[Green 82) Green, C.
Application of Theorem Proving to Problem Solving.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial intelligence, ,*
Tioga Press, 1982.

[Grosz 80) Grosz, B.J.
Utterance and Objective: Issues in Natural Language Communication,
*AI Magazine* 1(1):11-20, spring, 1980.

[Halpern 84] Halpern, J. and Moses, Y.
*Knowledge and Common Knowledge in a Distributed Environment.*
Technical Report, IBM Research, San Jose, 1984.

[Harmon and King 85]
Harmon, P. and King D.
*Expert Systems: Artificial intelligence in Business.*
John Wiley & Sons, New York, 1985.

[Haugeland 81] Haugeland, John (ed.).
*Mind Design: Philosophy, Psychology, Artificial intelligence.*
Massachusetts Institute Technology Press, 1981.

[Hayes 73] Hayes, P. J.
Computation and Deduction.
In *Proceedings of the Second Symposium on Mathematical Foundations of Computer Science,* pages 105-l 16. Czechoslovakian Academy of Sciences, 1973.

[Hayes 74] Hayes, P.J.
Some problems and non-problems of representation theory.
In *Proceedings of AISB summer conference, University of Sussex,* pages 63-69.
British Computer Society, AISB, 1974.

[Hayes 77] Hayes, P. J.
In Defence of Logic.
In *Proceedings of the Fifth IJCAI,* pages 559-565. August, 1977.

[Hayes 79] Hayes, Patrick J.
The Naive Physics Manifesto.
In Donald Michie (editor), *Expert Systems in the Micro-Electronic Age,* . Edinburgh University Press, Edinburgh, Scotland, 1979.

[Hayes 82a] Hayes, P. J.
The Logic of Frames.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Arrificiai intelligence, .* Tioga Press, Palo Alto, CA, 1982.

[Hayes 82b] Hayes, P. J.
The Frame Problem and Related Problems in Al.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Arrificiai intelligence, .* Tioga Press, CA, 1982.

[Hayes 85] Hayes, P. J.
The Second Naive Physics Manifesto.
In Hobbs, J. and Moore, R.C. (editors), *formal Theories of the Commonsense World,* . Ablex Press, 1985.

[Hayes-Roth 83] Hayes-Roth, B.
*The Blackboard Architecture: A General Framework for Problem Solving?.*
Report HPP-8330, Department of Computer Science, Stanford University, 1983.

[Hayes-Roth 85] Hayes-Roth, Barbara.
A Blackboard Architecture for Control.
*Artificial intelligence ,* 1985.

[Hayes-Roth, et al. 83]
Hayes-Roth, Frederick, Donald A. Waterman, and Douglas B. Lenat(eds.).
*Building Expert Systems.*
Addison-Wesley, Reading, MA, 1983.

[Hillis 81] Hillis, W. Daniel.
*The Connection Machine.*
AIM 646, Al Lab, MIT, Cambridge, MA, 1981.

[Hirst 84] Hirst, Graeme.
A Semantic Process for Syntactic Disambiguation.
In *Proceedings of the AAAi-84 National Conference.* AAAI, University of Texas at Austin, TX, August, 1984.

[Hobbs 85] Hobbs, J.R. and Moore, R. C.
*Formal Theories of the Commonsense World.*
Ablex Press, 1985.

[Holland 75] Holland, J. H.
*Adaptation in Natural and Artificiai Systems.*
University of Michigan Press, Ann Arbor, 1975.

[Israel 81] Israel, D. and Brachman, R.
Distinctions and Confusions: A Catalogue Raisonne.
In *Proceedings of IJCAI-7*. AAAI, Vancouver, B.C, 1981.

[Kling 71] Kling, Robert E.
A Paradigm for Reasoning by Analogy.
*Artificial intelligence* 2(2):147-178, fall, 1971.
Also IJCAI-71, British Computer Society, London, 1971.

[Kowalski 79a] Kowalski, R.
Algorithm = logic + control.
*CACM* 22(7):424-436, 1979.

[Kowalski 79b] Kowalski, Robert.
*Logic for Problem Solving.*
North-Holland/Elsevier, Amsterdam, 1979.

[Laird, et al. 84] Laird, John E., Paul S. Rosenbloom, and Allen Newell.
Towards Chunking as a General Learning Mechanism.
In *Proceedings of the AAAI-84 National Conference.* AAAI, University of Texas at Austin, TX, August, 1984.

[Lenat 82a] Lenat, D.B.
Nature of Heuristics I.
*Artificial intelligence* 19, 1982.

[Lenat 82b] Lenat, D. B.
*The Role of Heuristics in Learning by Discovery: Three Case Studies.*
Tioga, Palo Alto, 1982, .

[Lenat 83] Lenat, D.B.
The Nature of Heuristics II and III.
*Artificial intelligence* 21(1), March, 1983.

[Loveland 83] Loveland.
Automated Theorem proving : a quarter century review.
*American Mathematical Soceity 29,* 1983.

[Lowe 83] Lowe, David and Binford, Thomas 0.
Perceptual organization as a basis for visual recognition.
In *Proceedings of the AAAI.* AAAI, AAAI, 1983.

[Lozano-Perez 82]
Lozano-Perez, Tomas.
Task Planning.
In Brady, Michael et. al. (editors), Robot Motion : Planning *and Control, .*
Massachusetts Institute Technology Press, 1982.

[Mackworth 82] Mackworth, Alan.
Consistency in Networks of Relations.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence, .*
Tioga Press, Palo Alto, CA, 1982.

[Manna 82] Manna, Z. and Waldinger, R.
A Deductive Approach to Program Synthesis.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artficial Intelligence, .*
Tioga Press, CA, 1982.

[Marr 82] Marr, David.
Vision.
W.H. Freeman, San Francisco, CA, 1982.

[Martin, et al. 83] Martin, Paul, Douglas Appelt, and Fernando Pereira.
Transportability and Generality in a Natural-Language Interface System.
In Alan Bundy (editor), *Proceedings* of the Eighth *IJCAI*, pages 573-581. IJCAI, Inc.,
Karlsruhe, West Germany, August, 1983.
Vol. 1.

[McCarthy 68] McCarthy, J.
Programs with Common Sense.
In Minsky, M. (editor), *Semantic Information Processing,* pages 403-417. MIT
Press, Cambridge, Mass., 1968.
Also in Mechanisation of Thought Processes, Proc. Sympos. National Physical Lab.
vol. I, pp.77-84. London: Her Majesty's Stationary Office. November, 1958.

[McCarthy 79] McCarthy, John.
*Ascribing Mental Qualities to Machines.*
Technical Report STAN-CS-79-725 AIM-326, Stanford University, March, 1979.

[McCarthy 82a] McCarthy, John.
Circumscription - A Form of Non-Monotonic Reasoning.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence, .*
Tioga Press, CA, 1982.

[McCarthy 82b] McCarthy, J. and Hayes, P.J.
Some philosophical problems from the standpoint of Al.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Inteligence,* , Tioga Press, CA, 1982.

[McCarthy 82c] McCarthy, J.
Epistemological Problems in Al.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Inteligence,* . Tioga Press, CA, 1982.

[McCorduck 79] McCorduck, P.
*Machines Who Think.*
W. H. Freeman and Co., San Francisco, 1979.

[McDermott 76] McDermott, D.
Artificial intelligence meets natural stupidity.
*S/G PLA N Newsletter* (57):4-9, 1976.

[McDermott 82] McDermott, J.
R1 : A rule-based configurer of computer systems.
*Artificial Intelligence* 19(1):39-88, September, 1982,
Also CMU Tech Report CMU-CS-80- 119, April 1980.

[McDermott 85] McDermott, Drew.
Reasoning about plans.
In *Formal Theories* of *the Commonsense World,* . Ablex Press, 1985.

[McDermott and Doyle 80]
McDermott, Drew V. and Doyle, Jon..
Non-monotonic logic I.
*Artificial Intelligence* 13( l&2):41 -72, April, 1980.
Special issue on Non-Monotonic Logic.

[Minsky 75] Minsky, Marvin.
A Framework for Representing Knowledge.
In Patrick H. Winston (editor), *The Psychology of Computer Vision,* pages 21 l-277. McGraw-Hill, New York, 1975.

[Mitchell 81] Mitchell, T.M.
Bias in Generalization.
In *Proceedings* of *IJCAI-8* I, . AAAI, 1981.
Computers and Thought Award Lecture.

[Mitchell 82] Mitchell, T.M.
Generalization as Search.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence, .* Tioga Press, CA, 1982.

[Mitchell 84] Mitchell, T., Steinberg, L., and S. Amarel.
*A Learning Apprentice for Knowledge Acquisition in an Expert System.*
Rutgers Digital Design Project Working Paper 16, Rutgers University, 1984.

[Mitchell, et al. 81]
Mitchell, T. M., Utgoff, P. E., Nudel, B. and Banerji, R.
Learning Problem-Solving Heuristics through Practice.
In *Proceedings of the Seventh IJCAI*, pages 127- 134. Vancouver, August, 1981.

[Mitchell, et al. 82]
Mitchell, T. M., Utgoff, P. E. and Banerji, R. B.
Learning by Experimentation: Acquiring and Refining Problem-Solving Heuristics.
In Michalski, R. S., Carbonell, J. G. and Mitchell, T. M. (editors), *Machine Learning, .* Tioga, 1982.

[Moore 82] Moore, R.C.
Reasoning About Knowledge and Action.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence, .* Tioga Press, CA, 1982.

[Moore 83] Moore, Robert C.
Semantical Considerations on Nonmonotonic Logic.
In Alan Bundy (editor), *Proceedings of the Eighth IJCAI*, pages 272-279. IJCAI, Karlsruhe, West Germany, August, 1983.
Volume 1.

[Moravec 80] Moravec, H. P.
*Obstacle Avoidance and Naviagtion in the Real World by a Seeing Robot Rover.*
Technical Report AIM-340, Stanford Artificial Intelligence Laboratory, September, 1980.

[Moses 71] Moses, J.
Symbolic integration: The stormy decade.
*Communications ACM* 14(8):548-560, 1971.

[Nevatia 82] Nevatia, Ramakant.
*Machine Perception.*
Prentice-Hall, Inc., Englewood Cliffs, NJ, 1982.

[Newell 76] Newell, A. and Simon, H. A.
Computer Science as empirical inquiry : symbols and search.
*Communications of the ACM* :113-126, 1976.

[Nii and Feigenbaum 78]
Nii, H. P. and E. A. Feigenbaum.
Rule-Based Understanding of Signals.
In D.A. Waterman and F. Hayes-Roth (editors), *Pattern-Directed Inference Systems,* pages 483-501. Academic Press, New York, 1978.

[Nilsson 80] Nilsson, Nils J.
*Principles of Artificial Intelligence.*
Tioga Press, Palo Alto, CA, 1980.

[Nilsson 82] Nilsson, Nils J.
Artificial Intelligence: Engineering, Science, or Slogan?
*Al Magazine* 3(1), winter, 198111982.

[Nilsson 83] Nilsson, Nils J.
Artificial Intelligence Prepares for 2001.
*A/ Magazine* 4(4), winter, 1983.

[Pearl 83] Pearl, Judea.
Knowledge versus Search: A Quantitative Analysis Using A*.
*Artificial Intelligence* 20(1):1-13, January, 1983.

[Pople 82] Pople, H. Jr.
Heuristic Methods for Imposing Structure on Ill-structured Problems.
In Szolovits, P. (editor), *A/ in Medicine,* . Westview Press, 1982.
AAAS Selected Symposium.

[Reiter 78] Reiter, R.
On reasoning by default.
In Waltz (editor), *TINLAP-2 (Theoretical Issues in Natural Language Processing),* pages 210-218. University of Illinois, July, 1978.

[Reiter 80] Reiter, R.
A Logic for Default Reasoning.
A/13(1&2):81-132, April, 1980.
Special issues on Non-Monotonic Logic.

[Reiter 82] Reiter, R.
On Closed World Databases.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence, .*
Tioga Press, CA, 1982.

[Rich 83] Rich, Elaine.
*Artificial Intelligence.*
McGraw-Hill Book Co., New York, 1983.

[Rich 84] Rich, Elaine.
Default reasoning as likelihood reasoning.
In *Proceedings of AAAI-84.* 1984.

[Rich and Shrobe 79]
Rich, C., and H.E. Shrobe.
Design of a programmer's apprentice.
In Winston and Brown (editors), *AI: An MIT Perspective,* pages 137-173. MIT Press, Cambridge, MA, 1979.

[Rosenschein 84] Rosenschein, Stan.
A Perspective on Planning.
In *Proceedings of the AAAI-84* . AAAI, University of Texas at Austin, TX, August, 1984.

[Sacerdoti 81] Sacerdoti, Earl D.
Problem Solving Tactics.
*A/ Magazine* 2(1), winter, 1980/1981.

[Schank 83a] Schank, Roger C.
The Current State of AI: One Man's Opinion.
*A/ Magazine* IV(I), 1983.

[Schank 83b] Bundy, Alan.
The Nature of AI: A Reply to Schank.
*AI Magazine* IV(4), 1983.

[Schmidt 78] Schmidt, Sridharan and Goodson.
The plan recognition problem: an intersection of psychology and AI.
*Artificial Intelligence*, August, 1978.

[Schubert 76] Schubert, L.K.
Extending the Expressive Power of Semantic Networks.
*Artificial Intelligence* 7(2):163-198, 1976.

[Shapiro 84] Shapiro, Ehud.
*A subset of Concurrent Prolog and its Interpreter.*
Technical Report TR-003, ICOT, Japan, 1984.

[Shortliffe 82] Shortliffe, E.
Consultation Systems for Physicians.
In *Readings in Artificial Intelligence,* . Tioga Press, CA, , 1982.

[Simon 69] Simon, Herbert A.
*The Sciences of the Artificial.*
MIT Press, Cambridge, MA, 1969.

[Simon 82] Simon, H. A.
Why Should Machines Learn?
In Michalski, R. S., Carbonell, J. G. and Mitchell, T. M. (editors), *Machine Learning, ,*
Tioga, 1982.

[Simon 83] Simon, H.A.
Search and Reasoning in Problem Solving.
*Artificial Intelligence* 21:7-29, March, 1983.

[Singh 85] Singh, V. and Genesereth, M.R.
A Variable Supply Model for Distributing Deductions.
In *Proceedings of IJCAI-85, Los Angeles.* 1985.
also HPP-84-14.

[Smith 82] Smith, B.
*Reflection and semantics in a procedural language.*
Technical Report MIT-TR-272, Cambridge, MA, M.I.T., 1982.

[Steele 78] Steele, G.L., and G.J. Sussman.
*Constraints.*
Al Memo 502, MIT, Cambridge, MA, November, 1978.

[Stefik 80] Stefik, M.J.
*Planning with Constraints.*
PhD thesis, Computer Science Department, Stanford University, January, 1980.
Stanford Rep. Nos. HPP-80-2, STAN-CS-80-784.

[Stefik 82] Stefik, M.
Planning and Meta-Planning.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence,* .
Tioga Press, CA, 1982.

[Stolfo, et al. 83] Stolfo, Salvatore J., Daniel Miranker, and David Elliot Shaw.
Architecture and Applications of DADO: A Large-Scale Parallel Computer for Artificial Intelligence.
In Alan Bundy (editor), *Proceedings of the Eighth IJCAI*, pages 850-854. IJCAI, Inc., Karlsru he, West Germany, August, 1983.
Vol. 2.

[Szolovits 78] Szolovits and Pauker.
Categorical and Probabilistic Reasoning in Medical Diagnosis.
*Artificial Intelligence* , August, 1978.

[Waldinger 82] Waldinger, R.
Achieving Several Goal Simultaneously.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artficial Intelligence,*. Tioga Press, CA, 1982.

[Webber 82] Webber, B.L. and Nilsson, N.J.
*Readings in Artificial Intelligence.*
Tioga Press, 1982.

[Weiss 78] Weiss, Kulikowski, Amarel and Safir.
A model based method for computer aided medical decision making.
*Artificial Intelligence* , August, 1978.

[Weizenbaum 76] Weizenbaum, J.
*Computer Power and Human Reason: From Judgment to Calculation.*
W. H. Freeman, San Francisco, 1976.
Comments to the book by Kuipers and McCarthy, in SIGART Newsletter No. 58, June 1976.

[Weyhrauch 82] Weyhrauch, R.
Prolegomena to a Theory of Mechanized Formal Reasoning.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artficial Intelligence, .* Tioga Press, CA, 1982.

[Wilkins 82] Wilkins, D.
Patterns and Plans in Chess.
In Webber, B.L. and Nilsson, N.J. (editors), *Readings in Artificial Intelligence, .* Tioga Press, CA, 1982.

[Winograd 753 Winograd, T.
Frame representations and the procedural/declarative controversy.
In D. G. Bobrow and A. Collins (editors), *Representation and Understanding: Studies in Cognitive Science,* pages 185-210. Academic Press, New York, 1975.

[Winograd 80a] Winograd, T.
Extended inference modes in reasoning by computer systems.
*A/13:5-26*, April, 1960.
Special issue on Non-Monotonic Logic.

[Winograd 80b] Winograd, T.
What does it mean to understand language?
***Cognitive Science*** 4:209-241, 1960.

[Winograd 83] Winograd, Terry.
***Language as a Cognitive Process, Vol. I: Syntax.***
Addison- Wesley, Reading, MA, 1983.

[Winograd 86] Winograd, Terry and Flores, Fernando.
***Computers and Cognition.***
Ablex, 1966.

[Winston 84] Winston, Patrick Henry.
***Artificial Intelligence (2nd ed.).***
Addison-Wesley, Menlo Park, CA, 1964.

[Woods 82] Woods, W.A.
Optimal Search Strategies for Speech-Understanding Control.
In Webber, B.L. and Nilsson, N.J. (editors), ***Readings in Artificial Intelligence,.***
Tioga Press, Palo Alto, CA, 1962.

[Wos 84] Wos, Overbeek, Lusk and Boyle.
***Automated Reasoning: Introduction and Applications.***
Prentice-Hall, 1964,