MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A.I. Memo No. 568 February, 1980

# A Selected Descriptor-Indexed Bibliography to the Literature on Belief Revision[1]

**Jon Doyle**[2]
*Artificial Intelligence Laboratory*
*Massachusetts Institute of Technology*
*Cambridge, Massachusetts*

and

**Philip London**[3]
*USC/Information Sciences Institute*
*Marina del Rey, California*

This article presents an overview of research in an area loosely called *belief revision.* Belief revision concentrates on the issue of revising systems of beliefs to reflect perceived changes in the environment or acquisition of new information. The paper includes both an essay surveying the literature and a descriptor-indexed bibliography of over 200 papers and books.

1. This paper also appears in ACM's SIGART Newsletter #71, April 1980.

2. This report describes research done in part at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643 and in part by National Science Foundation Grant MCS77-04828. J.D. also thanks the Fannie and John Hertz Foundation for their support of a graduate fellowship.

3. This research was supported in part by the Defense Advanced Research Projects Agency contract DAHC15 72 C 0308. Views and conclusions in this report are the author's and should not be interpreted as representing the official opinion or policy of DARPA, the U.S. Government, or any other person or agency connected with them.

## 1. Introduction

This article presents an overview of research in an area loosely called *belief revision*. Belief revision concentrates on the issue of revising systems of beliefs to reflect perceived changes in the environment or acquisition of new information. In addition, belief revision research includes the study of methods for representing models of environments as collections of beliefs and the development of formal theories of belief. The bulk of the article consists of a descriptor-indexed bibliography of research addressing these topics. Our intention is that this bibliography serve both to introduce the interested researcher to this literature, and to capture the current state of the field. Towards this purpose, we begin with an overview of belief revision research.

Besides more fully defining what we mean by belief revision, the overview that follows also serves to define the descriptors used to classify the entries of the bibliography. Descriptor terms appear boldfaced when they appear in the discussion. We should point out that the classification system is far from optimal; the descriptive categories that we have selected do not always permit very sharp distinctions. The reader will do well to check all reasonable categories to find full coverage of some topic.

We have attempted to include as many relevant works as possible, in particular all those in AI as well as major works in other fields. However, a topic as broad as belief revision, drawing on the literatures of artificial intelligence, logic, philosophy and psychology, makes full coverage difficult. We hope that readers finding inaccuracies in this document will suggest reclassifications or recommend additional entries to us. With such improvements, we may attempt a more thorough survey in a year's time or so. We suggest not using this bibliography as a base for other compilations without checking all the references carefully. Our purpose was to make an informal guide to the literature, and not to assure total inclusiveness or accuracy.

Finally, we wish to thank all those people who were kind and helpful enough to respond to our inquiries. We hope we have justified their efforts.

### Overview

Intelligence is often viewed as the ability to reason about and adapt to a changing environment. For this reason, most computer programs constructed by artificial intelligence researchers maintain a model of their external environment. The model is

updated to reflect changes in the environment resulting from the program's actions or indicated by its perception of external changes. AI programs frequently explore assumptions or hypotheses about these environments; this may lead to further model updating if new information conflicts with old, indicating that some of the currently held assumptions or hypotheses should be abandoned.

The need for maintaining models is widespread in AI. For example, programs which accept information in natural language must be ready to modify their beliefs or interpretations of previous inputs in light of the analysis of additional text. Programs analyzing visual images must revise their hypotheses about the content of those images when new points of view present new information. Robot programs revise their beliefs about their current environment when their own actions affect that environment. Finally, with increasing frequency, programs are being constructed to maintain and update models of some aspects of their own structure or inner environment, such as the history of their actions and inferences, their current goals, and their current abilities.

There is a strong philosophical and logical tradition in the study of rules for revising systems of beliefs. Classical logic embodies an uncontroversial sort of belief revision, that of the sound inference rules for deriving the logical consequences or theorems of a set of axioms. Inference rules such as these are called *monotonic* because the addition of axioms or rules in such a system can lead only to more theorems, never fewer. In contrast, **non-monotonic** inference rules may change the set of theorems of beliefs non-additively. The most common sorts of non-monotonic inference rules are those for resolving apparent inconsistencies in a set of beliefs. Such a rule is Sherlock Holmes' famous dictum "When you have eliminated the impossible, whatever remains, however improbable, must be the truth," which counsels us to accept a belief despite our initial lack of confidence in it if all other alternatives have been refuted.

Many systems of normative and descriptive rules for revision of beliefs have been developed. Quine, for example, has formulated the principle of "minimum mutilation," which suggests giving up only those beliefs that make the "smallest" change in our set of beliefs but still rids it of apparent inconsistency. He goes on to describe our belief systems as great "webs of belief" in which new information impinges on the boundaries of the web, only rarely requiring changes in the interior of the web. Quine's web metaphor [Quine 78] bears a remarkable similarity to the network representations used

in recent AI programs to represent systems of beliefs.

While little appears in the philosophical literature on mundane changes of belief, much appears which discusses large-scale changes in belief systems, particularly changes in scientific theories. For example, Carnap prescribes a form for the empiricist tradition, and Kuhn's classic account describes the sociological undercurrents of changes in scientific thought [Kuhn 62].

## The Frame Problem

It remained to Artificial Intelligence researchers to uncover a major problem virtually undiscussed in earlier work. This is the so-called **frame problem** of McCarthy and Hayes, the problem of how to update models to account for changes induced by actions. The basis of the problem is that even if one can succinctly specify the ways in which a system's environment might change in terms of the effects of actions, it still remains to specify some way of judging what stays unchanged in the face of these actions. In the words of McCarthy and Hayes,

> "If we had a number of actions to be performed in sequence, we would have quite a number of conditions to write down that certain actions do not change the values of certain fluents. In fact, with *N* actions and *M* fluents, we might have to write down *NM* such conditions."[1]

This problem actually includes several related subproblems:

1. How can we express the possible errors arising from fallible predictions about the effects of actions?

2. What can be done to recover from incorrect predictions about the effects of some action?

3. How should the predicted effects of actions be represented for the purposes of updating the model and interrogating the problem solving system's ability to alter its environment? That is, how are they stated and how are they retrieved?

4. How should the descriptions of causal effects of actions be applied to derive the description of the state obtained after the action?

5. What notion or representation of causality could be used to determine the

---

[1] [McCarthy 69] pg. 487.

causal effects of actions of physical mechanisms in an external environment?

The statement of the frame problem initiated research by a number of investigators and progress continues to be made on most of the problems mentioned above. As we will describe in more detail below, "defaults" and other non-monotonic inferences seem to answer (1) satisfactorily. In conjunction with data-dependency techniques they also answer (2) to some degree. (3) is usually approached via rules of the form "If P holds before the action, then Q holds after." Such rules are sometimes called "laws of motion". Indexing these rules by their actions, preconditions, and postconditions in a pattern-indexed database allows their retrieval. Dependency-based revision techniques have helped advance the solution to (4), which previously had been handled by less automatic database management techniques.. All of these problems, however, are still unsolved to varying degrees. In particular, little work has been done related to (5) (but see bibliography entries under classification **causal analysis**, <7.8>).

To the extent that the belief revision problem is a subproblem of, or at least has a large overlap with the frame problem, the progress reported below and in the references on belief revision is also progress on the frame problem. The major techniques developed to date (pattern-indexed databases, context-layered databases, data-dependencies, non-monotonic logics, and formalized meta-theories) all seem promising and deserve further study. Many important problems have not been studied to any depth, and others await adequate formulation. Some important topics include the efficiency of belief revision in large systems, the replacement of inference chains by computations, notions of causality, and the formulation of heuristic information for selecting particular inference rules.

## Concepts of Belief Revision

We now turn from describing the general problems of maintaining models of changing environments to outlining the proposed solutions to these problems discussed in the literature. These ideas can be divided roughly into two categories, "implementational" and "theoretical", but such distinctions often break down when pressed. The concepts discussed below account for much of the work in the field and are used as the basis for many of the descriptors in the indexed bibliography.

**Implementation-oriented concepts**

The implementational approaches to belief revision encompass the methods used in computer programs to overcome one or more of the problems enumerated earlier. The earliest and least sophisticated of these methods is that of **manual updates**. In programs using this technique, the entire model was represented by the values (usually numeric) of several variables. To update the model, the problem solving program was itself required to change the values of these variables. For example, the program might have a variable LOCATION-OF-BLOCK which it would alter from "ROOM-A" to "ROOM-B".

Later, globally accessible pattern-indexed databases replaced variable-based models, but manual updates were still used. Pattern-indexed databases permit the efficient retrieval of all statements in a relational database which match a pattern or statement skeleton. This allowed one to write updating procedures which changed all beliefs of a certain form. For example: Retrieve all statements matching (LOCATION BLOCK ?) and replace them with the statement (LOCATION BLOCK ROOM-B).

The **Operator application** approach to model updates represented a major step over manual updates. (The terms "operator" and "action" are often used synonomously.) The operator application process updates the model by interpreting an operator description and reflecting the update in a pattern-indexed database. **Operator descriptions** contain specifications of the expected effects of the operator. The classical example of a system using operator application is the STRIPS method, in which an operator description consists of two lists of formulas, the "add" list and the "delete" list [Fikes 71]. STRIPS' operator application procedure first deletes all database assertions matching any formula in the delete list. Then it adds to the database each of the formulas in the add list. For example, the add/delete list representation for the operator MOVE-BLOCK might be:

```
MOVE-BLOCK(block, destination)
      DELETE: (LOCATION block ?)
      ADD:    (LOCATION block destination).
```

A problem in STRIPS and related systems is that these additions and deletions were oblivious to consequential changes (i.e., changes that were logically deducible from a primitive set of action effects), so *all* effects had to be listed in the operator description. **Context-layered databases** were developed in part to allow the

association of consequential changes with the action that gave rise to them [Rulifson 72, Fahlman 74, Sussman 75a]. Context-layered databases allow the encapsulation of sections of the database, so that the consequences of an effect of an action can be grouped with a reference back to the causing action. This allows the addition or deletion of an operator effect together with its logical consequences in a single step.

Parallel with context-layering of a database, **change-triggered procedures** were employed to perform the updating appropriate to a single inference rule. The use of these procedures was necessitated by the observation that context-layering alone only approximates the concept of consequentiality. The shortfall of context-layering is its lack of a representation for the nature of the derivations of consequential effects of actions. Change-triggered procedures, (which include THANTE and THERASE in PLANNER [Hewitt 72, Sussman 71] and IF-ADDED and IF-REMOVED in CONNIVER [McDermott 74b], and other forms of pattern-directed invocation and procedural attachment), may be invoked when a given database statement or object is either added to or removed from the database. Such procedures were frequently used in pairs, each pair corresponding to a single inference rule. For example, for an inference rule "If A then infer B", one would write the two change-triggered procedures "IF-ADDED A then ADD B", and "IF-REMOVED A then REMOVE B."

It is apparent, though, that a given statement can often have multiple derivations. The possibility of multiple and even circular derivations leads to many difficulties when using change-triggered procedures to update databases, for these procedures depend only on the occurrence of changes and not on the form of the derivation. For example, the rule "IF-REMOVED A then REMOVE B" might not be valid if B is derivable from another true assertion C. **Data-dependencies** provide a solution to these problems. Data-dependencies are explicit records of inferences or computations. These records are examined to determine the set of valid derivations, and hence the current set of beliefs (that is, those statements with valid arguments). In some cases, they are erased along with the beliefs they support when changes lead to removing a belief and its consequences from the database [Fikes 75]. In other systems, the dependencies are kept permanently. In this latter case, **dependency-based revision** techniques can use a uniform procedure, sometimes called truth maintenance [Doyle 79a], to mark each database statement as believed or not believed, depending on whether the recorded

derivations currently provide a valid argument for the statement. One might view this sort of dependency analysis as analogous to the mark/sweep garbage collection procedures of list-processing systems.

The data-dependency technique handles consequential updates, but updating to accomodate the basic action effects is another problem altogether. This problem is usually discussed in terms of backtracking which is invoked to remove an apparent inconsistency caused by new information or by an action effect conflicting with a belief for which the program desires to maintain validity (e.g., because it is a goal in a problem solving domain [London 78b]). By **non-chronological backtracking** we mean the general technique of searching only for the set of possible assumptions to change among those beliefs known to be logically supporting the beliefs involved in the inconsistency.[2] This is to be contrasted with the previously popular **chronological backtracking**, in which the change was made at the latest choicepoint. Dependency-directed backtracking techniques usually involve tracing through the recorded data-dependencies of the conflicting beliefs. In some systems, this technique involves analyzing the simultaneous goals for which no solution can be found to see which goals share variables, and which goals are independent of one another.[Pereira 79a, McDermott 77] Sometimes dependency-directed backtracking systems record the source of the inconsistency (the set of conflicting premises or assumptions) to prevent its future reoccurrence [Stallman 77]. This can be viewed as a simple form of learning from mistakes (see <7.3> and <7.5>).

The techniques described above represent a progression toward more and more sophisticated techniques for maintaining model consistency. The use of **non-monotonic inference rules** forms a dimension somewhat orthogonal to the above techniques. Non-monotonic inference rules allow the system to make tentative conclusions based on partial evidence. The typical use of such inferences is in default reasoning, where some condition is assumed to hold until contrary evidence is found. The techniques involved began with the THNOT primitive of PLANNER, which would be satisfied as a goal if its argument goal failed, that is, if no derivation for its argument formula could be found. Another form of non-monotonic inference inference is that of **partial matching**, in which

---

[2] This is sometimes called dependency-directed backtracking [Stallman 77].

any assumptions necessary are made to allow the retrieval of some database object which matches a supplied pattern for the object [Joshi 75a, Hayes-Roth 78]. Non-monotonic data-dependencies record such non-monotonic inferences, and make the truth maintenance algorithms for updating data-dependencies considerably more complex than in the monotonic case.

**Theory-oriented concepts**

The theory-oriented concepts of belief revision encompass the formal study of beliefs and belief systems as objects of interest in their own right. A primary concept in this regard is that of the **situational calculus**, in which all predicates and objects are indexed by situations or "worlds." McCarthy made early proposals along these lines [McCarthy 68], and logicians have studied situational or indicial logics for some time in connection with analyticity, temporal logics, counterfactual conditionals, and the semantics of belief statements and other propositional attitudes. In the situational calculus, each situation variable refers to the world at a given instant. This calculus is itself amenable to the implementation techniques mentioned previously, and is occasionally referenced as an implementational technique.

The study of **theory evolution** is the mathematical analysis of the temporal evolution of sets of formulas, such as the successive sets of statements representing the beliefs of an agent. Whereas the situational calculus approach views the world as an object to which statements refer, the theory evolution approach views worlds as collections of statements. It provides a framework for phrasing questions about inference rules viewed as theory changes, whether monotonic or non-monotonic. One basic such study is that of joint consistency or extension theorems, which study the relation between a consistent theory and its inconsistent extensions in order to characterize or localize the exact nature of the inconsistency.

The **theory of incomplete databases** can be viewed as a restriction of the theory evolution area in which one studies an incomplete database of atomic statements or relations. Typical concerns involve judging questions of possible or necessary truth of a statement in any extension or completion of the incomplete database, and whether the database can be made more succinct by means of the *closed world assumption*, which assumes that any atomic relation not explicitly listed in the database is false.

**Non-monotonic logics** attempt to give a precise semantics to systems of non-monotonic inference rules. This is a tricky problem, for as Sandewall pointed out [Sandewall 72], several distinct (and incompatible) sets of "theorems" may result from a single set of non-monotonic axioms and inference rules. These logics have been investigated both as modal logics and as ordinary predicate logic with non-monotonic inference rules. (See also **conclusion theory** <5.3>, for a related idea.)

**Proof theory** studies proofs as mathematical objects. A typical concern of much of the literature is whether proofs in one inferential system can be rewritten (possibly uniformly) into proofs in another (possibly simpler) inferential system. For example, long arguments frequently must be summarized, and this can be viewed as the rewriting of a long proof using "small" inference steps to a shorter proof using "larger" derived inference steps.

**Meta-theory** is a broad generalization of the formal approach, incorporating proof theory, model theory, and language theory to study systems of languages, proofs and their models as mathematical objects. This approach has the power to formalize all of the approaches mentioned above.

## Applications of Belief Revision

Belief revision techniques enjoy an important position in the technology available to implement or formulate other AI techniques and concepts. We have mentioned above the application of belief revision techniques to **execution monitoring** when modelling the world while executing plans of actions. In addition, data-dependencies lend themselves to **explanation** of program beliefs and actions, by presenting valid arguments as explanations of beliefs. These same explanations can guide human experts in **transfer of expertise**, the task of correcting and extending databases of program expertise. In this task, the explanations of program beliefs help the expert to see what mistaken beliefs or incorrect rules are used by the program. Closely related to this, the explanations also help a human (or computer) programmer in the **debugging of procedures** by tracing the records of the computation to find the procedures responsible for some unhappy result. These explanations might also help in **causal analysis** of beliefs and actions, although this is yet largely unexplored. The issue is to identify causal relationships among actions and beliefs held by the system. The revision

techniques allow crude forms of **hypothetical reasoning** or **perturbation analysis** to determine the effects on the set of beliefs of a change in some particular belief. A final application is in the **control of reasoning**, where belief revision and explanation techniques are used to revise the set of goals when some are achieved (e.g., to abandon the subgoals of an achieved goal), or when difficulties in execution force **replanning from failures** (e.g., to abandon the portions of the plan made irrelevant by the failure).

## Related Topics and Fields

Finally, we mention some related topics and fields which touch on belief revision but which we do not consider in any detail. For example, the psychological study of **human belief** and change of belief postulates mechanisms not captured by the computer techniques discussed above. **Epistemology** is an enormous field concerned in part with the justification of belief and the evaluation of the grounds for belief. **Inductive inference** treats the revision of hypotheses in reaction to a training sequence of partial examples and non-examples. **Hypothetical reasoning** and **counterfactual conditionals** are closely related to the belief revision problem, and have their own large literature. **Decision theory** typically applies probabilistic inference techniques (especially Bayesian methods) to update the utility measures for each option for action when given new information. **Conclusion theory** is an extension of decision theory suggested by Tukey to overcome the impermanence of decision theoretic judgements in the face of new evidence. It is intended to capture the human technique of drawing a conclusion and sticking by it unless overwhelming evidence to the contrary forces the abandonment of the conclusion. Finally, **multi-valued logics** and **fuzzy logics** each try to treat the problem of uncertainty of predictions by extending the range of truth values assigned to formulas.

## 2. Descriptor Categories

1. **Global issues**
    1. The frame problem - how to revise beliefs to account for actions
    2. Choosing between alternate competing belief revisions
    3. Change of Belief or Mind
    4. How to revise beliefs to account for new information
2. **Representations for revising beliefs**
    1. Manual updates
    2. Operator descriptions
    3. Context-layered databases
    4. Change-triggered procedures
    5. Logical data-dependencies
    6. Causal dependencies
    7. Dependency-like representations
3. **Procedures for revising beliefs**
    1. Chronological backtracking
    2. Operator application
    3. Dependency-based revision
    4. Non-chronological backtracking
4. **Non-monotonic inference techniques**
    1. Default reasoning
    2. Non-monotonic justifications
    3. Negation in databases
    4. Partial matching

5. Minimal model circumscription

**5. Inexact inferential techniques**

1. Probabilistic inference

2. Decision theory

3. Conclusion theory

4. Multi-valued logics

5. Fuzzy logics

**6. Theoretical issues**

1. Logics of rational belief

2. Temporal logics

3. Situational calculus

4. Theory evolution theory

5. Proof theory

6. Non-monotonic logics

7. Meta theory

8. Theory of incomplete databases

9. Modal logics

10. Reasoning about uncertainty

**7. Applications**

1. Explanation

2. Modelling and execution

3. Replanning from failures

4. Transfer of expertise

5. Learning

6. Debugging of computer programs and mechanisms

7. Evolutionary computer systems

8. Causal analysis

9. Perturbation analysis

10. Control of reasoning

**8. Related issues**

1. Human belief and memory

2. Epistemology

3. Inductive inference

4. Hypothetical reasoning and counterfactual conditionals

5. Scientific theory change

## 3. Indexed Citations

### 1. Global issues

**1. The frame problem - how to revise beliefs to account for actions**

[Doyle 80], [Havel 78], [Hayes 70], [Hayes 71], [Hayes 73], [Hayes 74], [Hendrix 73], [Joshi 75b], [London 78a], [London 78b], [McCarthy 69], [McCarthy 77], [McDermott 78c], [McDermott 79a], [McDermott 79b], [Moore 75], [Raphael 71], [Raphael 76], [Reiter 78a], [Reiter 80], [Sacerdoti 79], [Sandewall 72], [Schwind 78], [Shrobe 79b], [Sridharan 76], [Sridharan 77], [Srinivasan 76], [Stepankova 78], [Stepankova 76a], [Stepankova 76b], [Thompson 79], [Waldinger 77].

**2. Choosing between alternate competing belief revisions**

[deKleer 79a], [deKleer 79b], [Doyle 79a], [Doyle 80], [Fahlman 74], [Feldman 77], [Fox 77], [Good 68], [Good 77b], [Goodman 73], [Hajek 77], [Hajek 78b], [Harman 73], [Hart 72], [Hayes-Roth 78], [Joshi 78a], [Latombe 79], [McDermott 74a], [Quine 70], [Quine 78], [Schmidt 77], [Schmidt 78], [Simon 66], [Sridharan 76], [Sridharan 77], [Srinivasan 76].

**3. Change of Belief or Mind**

[Dennett 78], [Harman 73], [Minsky 79], [Suppes 77].

#### 4. How to revise beliefs to account for new information

[Charniak 78], [Dacey 78], [Doyle 78a], [Doyle 78b], [Doyle 79a], [Doyle 80], [Dummett 73], [Friedman 79], [Good 52], [Good 77a], [GumbM 78], [GumbR 78], [Harper 76], [Hayes 73], [Joshi 78b], [Kowalski 78], [May 76], [McDermott 74a], [Quine 53], [Quine 78], [Rescher 64], [Reiter 78a], [Reiter 80], [Rosenberg 78], [Rosenberg 79], [Rubin 75], [Shafer 76], [Skyrms 66], [Sridharan 78], [Teller 76], [Waterman 75], [Winston 75].

## 2. Representations for revising beliefs

### 1. Manual updates

[Fahlman 74].

### 2. Operator descriptions

[Crocker 77], [Daniel 77], [Fikes 72a], [Fikes 72b], [Fikes 71], [Havel 78], [Hayes 71], [Hayes 73], [Hewitt 72], [London 78b], [McDermott 77], [McDermott 78a], [Raphael 71], [Raphael 76], [Rieger 76a], [Sacerdoti 74], [Sacerdoti 77], [Sacerdoti 79], [Sandewall 72], [Schmidt 77], [Schmidt 78], [Shrobe 79a], [Shrobe 79b], [Sridharan 75], [Sridharan 76], [Sridharan 77], [Sridharan 78], [Stepankova 78], [Stepankova 76b], [Tate 76], [Tate 77].

### 3. Context-layered databases

[Charniak 79], [Fahlman 73], [Fahlman 74], [Fikes 75], [Hewitt 72], [Hewitt 75a], [Kornfeld 79a], [Kornfeld 79b], [McDermott 74a], [McDermott 75], [McDermott 77], [McDermott 78a], [McDermott 74b], [Minsky 75], [Moore 75], [Raphael 71], [Rubin 75], [Rulifson 72], [Sacerdoti 77], [Sussman 75a], [Sussman 72], [Sussman 71], [Tate 76], [Tate 77], [Thompson 79], [Waldinger 77].

### 4. Change-triggered procedures

[Bobrow 77], [Charniak 78], [Charniak 79], [deKleer 78a], [Hayes 73], [Hayes-Roth 78], [Hewitt 72], [Hewitt 75a], [Joshi 78a], [Joshi 75a], [Kornfeld 79a], [Kornfeld 79b], [McDermott 77], [McDermott 78a], [McDermott 74b], [Minsky 75], [Moore 75], [Rieger 77a], [Rieger 77b], [Roberts 77], [Rosenberg 78], [Rosenberg 79], [Rulifson 72], [Rychener 79], [Rychener 78], [Steele 78], [Sussman 72], [Sussman 71], [Swartout 77], [Waterman 75].

### 5. Logical data-dependencies

[Abelson 69], [Carbonell 79], [Charniak 78], [Charniak 79], [Cohen 80], [Colby 73], [Cox 77], [Cox 78], [Daniel 77], [DavisR 76], [deKleer 76], [deKleer 77], [deKleer 78a], [deKleer 78b], [deKleer 79c], [Doyle 76], [Doyle 78a], [Doyle 78b], [Doyle 78c], [Doyle 79a], [Doyle 80], [Doyle 79b], [Fikes 75], [Fikes 80], [Friedman 79], [Geiser 75], [Goguen 80], [GumbM 78], [GumbR 78], [Hayes 73], [Hayes 74], [Hayes 75], [Kahn 77], [Katz 76], [Latombe 76], [Latombe 79], [London 78a], [London 78b], [McAllester 78], [McAllester 79a], [McAllester 79b], [McDermott 75], [McDermott 77], [McDermott 78a], [McDermott 78b], [McDermott 78c], [McDermott 79a], [McDermott 79b], [Moriconi 77], [Nevins 74], [Reiter 80], [Rich 79], [Schmidt 77], [Schmidt 78], [Shrobe 79a], [Shrobe 79b], [Shrobe 79c], [Sridharan 76], [Stallman 77], [Stansfield 78], [Steele 78], [Sussman 75b], [Tate 76], [Tate 77], [Thompson 79], [Weiner 79], [Weyhrauch 78].

### 6. Causal dependencies

[Charniak 78], [Hart 72], [London 78a], [London 78b], [Miller 79], [Schmidt 77], [Schmidt 78], [Sridharan 77], [Thompson 79].

### 7. Dependency-like representations

[Abelson 73], [Crocker 77], [DavisR 80], [Geiser 75], [Goguen 80], [Marcus 79], [McDermott 74a], [McDermott 77], [McDermott 78a], [Miller 79], [Minsky 79], [Rieger 76a], [Rieger 76b], [Rieger 77c], [Sridharan 75], [Sridharan 77], [Sridharan 78], [Srinivasan 76], [Weiner 79], [Winston 78].

## 3. Procedures for revising beliefs

### 1. Chronological backtracking

[Berliner 74], [Charniak 79], [Sussman 72], [Sussman 71].

### 2. Operator application

[Crocker 77], [Daniel 77], [Fikes 72b], [Fikes 71], [Hayes 73], [Hewitt 72], [London 78b], [McDermott 77], [McDermott 78a], [Rieger 76b], [Rieger 77c], [Sacerdoti 74], [Shrobe 79a], [Sridharan 77], [Tate 77], [Thompson 79], [Waldinger 77].

### 3. Dependency-based revision

[Carbonell 79], [Charniak 79], [Cohen 80], [Cox 77], [Cox 78],

[Daniel 77], [deKleer 77], [deKleer 78a], [deKleer 78b], [deKleer 79c], [Doyle 76], [Doyle 78a], [Doyle 78b], [Doyle 78c], [Doyle 79a], [Doyle 80], [Fikes 75], [Fikes 80], [Friedman 79], [Hayes 73], [Hayes 75], [Kahn 77], [London 78a], [London 78b], [McAllester 78], [McAllester 79a], [McAllester 79b], [McDermott 75], [McDermott 77], [McDermott 78a], [McDermott 78b], [McDermott 78c], [McDermott 79a], [McDermott 79b], [Moriconi 77], [Nevins 74], [Reiter 80], [Shrobe 79a], [Shrobe 79b], [Shrobe 79c], [Stallman 77], [Stansfield 78], [Steele 78], [Sussman 77], [Thompson 79].

#### 4. Non-chronological backtracking

[Carbonell 79], [Charniak 79], [Cox 77], [Cox 78], [deKleer 77], [deKleer 78a], [deKleer 78b], [deKleer 79c], [Doyle 76], [Doyle 78a], [Doyle 78b], [Doyle 78c], [Doyle 79a], [Fahlman 73], [Fahlman 74], [Freuder 76], [Gaschnig 74], [Gaschnig 77], [Gaschnig 78], [Gaschnig 79], [GumbM 78], [Haralick 79], [Hewitt 75a], [Katz 76], [Latombe 76], [Latombe 79], [London 78a], [London 78b], [McAllester 78], [McAllester 79b], [McDermott 74a], [Nevins 74], [Pereira 79a], [Pereira 79b], [Rieger 77b], [Rubin 75], [Sacerdoti 79], [Schmidt 77], [Shrobe 79b], [Sridharan 76], [Sridharan 77], [Srinivasan 76], [Stallman 77], [Steele 78], [Sussman 77], [Tate 77].

### 4. Non-monotonic inference techniques

#### 1. Default reasoning

[Bobrow 77], [Carbonell 79], [Charniak 79], [Clark 78], [Collins 78], [deKleer 79c], [Harman 73], [Hayes 73], [Hayes 77], [Hewitt 72], [Joshi 78b], [Kornfeld 79a], [Kornfeld 79b], [Kramosil 75], [Lehrer 69], [Levesque 79], [McAllester 78], [McAllester 79a], [McAllester 79b], [McCarthy 69], [McDermott 74a], [McDermott 77], [McDermott 78a], [McDermott 78c], [McDermott 79a], [McDermott 79b], [Miller 79], [Minsky 75], [Reiter 78a], [Reiter 80], [Roberts 77], [Rosenberg 78], [Rosenberg 79], [Rubin 75], [Sandewall 72], [Schubert 79], [Scriven 59], [Stallman 77], [Sussman 71].

#### 2. Non-monotonic justifications

[Charniak 79], [deKleer 77], [deKleer 79c], [Doyle 76], [Doyle 78a], [Doyle 78b], [Doyle 78c], [Doyle 79a], [McDermott 78c], [McDermott 79a], [McDermott 79b], [Shrobe 79a], [Shrobe 79b], [Shrobe 79c].

#### 3. Negation in databases

[Davies 74], [Reiter 78a], [Reiter 78b].

### 4. Partial matching

[Abelson 69], [Hayes-Roth 78], [Joshi 78a], [Joshi 75a], [Joshi 75b], [Rosenschein 76], [Rosenschein 75], [Sridharan 75].

### 5. Minimal model circumscription

[McCarthy 77], [McCarthy 80].

## 5. Inexact inferential techniques

### 1. Probabilistic inference

[Colby 69], [Colby 73], [Collins 78], [Dacey 78], [DavisR 76], [DavisR 80], [Duda 76], [Giles 76], [Good 52], [Good 61], [Good 68], [Good 77a], [Good 77b], [Hajek 77], [Hajek 78b], [Harper 76], [Hart 72], [Havel 78], [Hayes 74], [May 76], [Minsky 61], [Scriven 59], [Shafer 76], [Skyrms 66], [Stansfield 78], [Swartout 77], [Teller 76].

### 2. Decision theory

[Feldman 74], [Feldman 77], [Good 52], [Good 77a], [Havel 78], [Lehrer 71], [Lehrer 75], [Tukey 60].

### 3. Conclusion theory

[Dacey 78], [Tukey 60].

### 4. Multi-valued logics

[Belnap 76], [Fikes 75], [Gaines 79], [Haack 78], [Hayes 70], [Srinivasan 76], [Stickel 77].

### 5. Fuzzy logics

[Friedman 79], [Gaines 76], [Gaines 79], [Haack 78], [Hayes 74], [Zadeh 75], [Zadeh 79].

## 6. Theoretical issues

### 1. Logics of rational belief

[Belnap 76], [Doyle 78c], [Doyle 80], [Dummett 73], [Giles 76], [Good 52], [Harper 76], [Hayes 73], [Hintikka 62], [Lehrer 71], [Lehrer 75], [Lehrer 69], [McCarthy 69], [Moore 77], [Moore 79a], [Quine 78],

[Rescher 64], [Rosenschein 79], [Shafer 76], [Shrobe 79a], [Skyrms 66].

**2. Temporal logics**

[Haack 78], [Hayes 70], [Hayes 71], [Hendrix 73], [Kahn 77], [McCarthy 68], [McCarthy 69], [McDermott 77], [McDermott 78a], [Rosenschein 79], [Schwind 78].

**3. Situational calculus**

[Crocker 77], [Havel 78], [Hayes 71], [Hayes 73], [Hewitt 75a], [Kornfeld 79a], [Kornfeld 79b], [McCarthy 68], [McCarthy 69], [McCarthy 77], [McDermott 77], [McDermott 78a], [Moore 75], [Moore 77], [Moore 79a], [Raphael 71], [Raphael 76], [Sandewall 72], [Shrobe 79a], [Stepankova 78], [Stepankova 76a].

**4. Theory evolution theory**

[GumbM 78], [GumbR 78], [GumbR 79], [Havel 78], [Lipski 76], [Lipski 78], [Quine 53], [Rescher 64], [Rosenschein 79], [Stepankova 76a], [Stepankova 76b].

**5. Proof theory**

[Boolos 79], [Clark 78], [Cohen 80], [DavisR 77], [DavisR 78], [DavisR 80], [Geiser 75], [Kreisel 77], [McDermott 78c], [McDermott 79a], [McDermott 79b], [Statman 74], [Stickel 77].

**6. Non-monotonic logics**

[DavisM 80], [Kramosil 75], [Levesque 79], [McCarthy 77], [McDermott 78b], [McDermott 78c], [McDermott 79a], [McDermott 79b], [Reiter 78b], [Reiter 80].

**7. Meta theory**

[Bobrow 77], [BrownF 79], [Collins 78], [Doyle 76], [Doyle 78c], [Doyle 80], [Minsky 65], [Weyhrauch 78].

**8. Theory of incomplete databases**

[Gaines 79], [Hajek 77], [Hajek 78a], [Hajek 78b], [Jaegermann 78], [Lipski 76], [Lipski 77], [Lipski 78], [Reiter 78b].

**9. Modal logics**

[Boolos 79], [Gaines 79], [Haack 78], [Hayes 70], [Hintikka 62],

[McCarthy 68], [McCarthy 69], [Moore 79b], [Schwind 78].

### 10. Reasoning about uncertainty

[Hewitt 75a], [Moore 77], [Moore 79a], [Stepankova 78].

## 7. Applications

### 1. Explanation

[BrownA 76], [Collins 78], [DavisR 76], [DavisR 77], [DavisR 78], [DavisR 80], [deKleer 76], [deKleer 79a], [deKleer 79b], [deKleer 78b], [Doyle 79b], [Geiser 75], [Goguen 80], [Hart 72], [Hewitt 75b], [Lipski 78], [McDermott 74a], [McDermott 77], [McDermott 78a], [Rosenschein 75], [Schmidt 77], [Schmidt 78], [Scriven 59], [Shrobe 79b], [Sridharan 75], [Sridharan 78], [Stallman 77], [Steele 78], [Sussman 75b], [Swartout 77], [Weiner 79].

### 2. Modelling and execution

[Daniel 77], [Doyle 80], [Fahlman 74], [Fikes 72a], [Hayes 75], [Hendrix 73], [London 78a], [London 78b], [McDermott 77], [McDermott 78a], [Sacerdoti 77], [Sacerdoti 79], [Sandewall 72], [Shrobe 79a], [Thompson 79].

### 3. Replanning from failures

[Daniel 77], [Doyle 80], [Fahlman 74], [Fikes 72a], [Hayes 75], [Latombe 76], [Latombe 79], [London 78a], [London 78b], [McDermott 77], [McDermott 78a], [Rieger 77b], [Sacerdoti 77], [Sacerdoti 79], [Shrobe 79b], [Srinivas 78], [Swartout 77], [Thompson 79].

### 4. Transfer of expertise

[DavisR 76], [DavisR 77], [DavisR 78], [DavisR 80], [Hart 72], [Joshi 78b], [McDermott 74a], [Rychener 78], [Rychener 79].

### 5. Learning

[Berliner 74], [Carbonell 79], [Cohen 80], [Fikes 72b], [Friedman 79], [Hart 72], [Lenat 77], [Minsky 62], [Minsky 61], [Minsky 65], [Minsky 79], [Stallman 77], [Sussman 75a], [Waterman 75], [Winston 75], [Winston 78].

### 6. Debugging of computer programs and mechanisms

[BrownA 76], [Doyle 80], [Doyle 79b], [Goldstein 74], [Goldstein 75],

[Katz 76], [McDermott 77], [McDermott 78a], [Miller 79], [Moriconi 77], [Rich 79], [Rieger 76b], [Rieger 77c], [Sacerdoti 79], [Shrobe 79a], [Srinivas 78], [Sussman 75a], [Sussman 77].

**7. Evolutionary computer systems**

[Doyle 80], [Hewitt 75b], [Moriconi 77], [Rich 79], [Rychener 78], [Sandewall 79], [Shrobe 79a], [Sussman 75a], [Sussman 79].

**8. Causal analysis**

[Berliner 74], [BrownA 76], [deKleer 76], [deKleer 79a], [deKleer 79b], [Fikes 72b], [Goldstein 74], [Goldstein 75], [Good 61], [Hart 72], [Latombe 79], [Minsky 62], [Minsky 61], [Rieger 76b], [Rieger 77c], [Rosenschein 75], [Shrobe 79b], [Sridharan 75], [Srinivas 78], [Stallman 77], [Sussman 75a], [Waldinger 77].

**9. Perturbation analysis**

[Berliner 74], [deKleer 78b], [Rich 79], [Shrobe 79a], [Stallman 77], [Steele 78].

**10. Control of reasoning**

[Abelson 73], [Cox 77], [deKleer 77], [deKleer 78b], [Doyle 76], [Doyle 78a], [Doyle 78b], [Doyle 78c], [Doyle 79a], [Doyle 80], [Hayes 74], [Hewitt 75a], [McDermott 77], [McDermott 78a], [Minsky 65], [Minsky 79], [Rosenberg 78], [Rosenberg 79], [Rosenschein 76], [Rosenschein 75], [Sacerdoti 79], [Shrobe 79a], [Shrobe 79b], [Shrobe 79c].

## 8. Related issues

**1. Human belief and memory**

[Abelson 69], [Abelson 73], [Colby 69], [Colby 73], [Collins 78], [Dennett 78], [Gaines 76], [Goguen 80], [Good 52], [Harman 73], [Minsky 65], [Minsky 75], [Minsky 79], [Moore 79b], [Quine 78], [Suppes 77], [Tukey 60], [Weiner 79].

**2. Epistemology**

[Belnap 76], [Chisolm 66], [Dummett 73], [Gaines 76], [Gaines 79], [Griffiths 67], [Harman 73], [Harper 76], [Hayes 70], [Hayes 74], [Hintikka 62], [Lehrer 71], [Lehrer 75], [Lehrer 69], [McCarthy 68], [McCarthy 69], [McCarthy 77], [Minsky 65], [Quine 53], [Quine 78],

[Scriven 59], [Shafer 76], [Skyrms 66], [Sosa 75], [Tukey 60], [Turner 78].

**3. Inductive inference**

[Friedman 79], [Goodman 73], [Hajek 77], [Hajek 78a], [Harman 73], [Harper 76], [Hayes 74], [Hayes-Roth 78], [Lehrer 71], [Lehrer 75], [Minsky 62], [Minsky 61], [Skyrms 66], [Suppes 77], [Tukey 60].

**4. Hypothetical reasoning and counterfactual conditionals**

[Goodman 73], [Harper 76], [Hayes 73], [Lewis 73], [Rescher 64], [Simon 66], [Skyrms 66], [Sosa 75], [Turner 78].

**5. Scientific theory change**

[Goodman 73], [Hempel 66], [Kuhn 62], [Lakatos 76], [Skyrms 66].

## 4. References

[Abelson 69] Abelson, R. P., and Reich, C. M., "Implicational molecules: a method for extracting meaning from input sentences," in *IJCAI-69*, pp. 641-645, 1969.
<2.5>, <4.4>, <8.1>.

[Abelson 73] Abelson, R. P., " The structure of belief systems," in R. Schank and K. M. Colby (ed.), *Computer Models of Thought and Language*, W. H. Freeman, San Franciso, 1973.
<2.7>, <7.10>, <8.1>.

[Belnap 76] Belnap, N. D., "How a computer should think," in G. Ryle (ed.), *Contemporary Aspects of Philosophy*, Oriel Press, Stocksfield, 1976.
<5.4>, <6.1>, <8.2>.

[Berliner 74] Berliner, H. J., *Chess as problem solving: the development of a tactics analyzer*, Ph.D. thesis, CMU, 1974.
<3.1>, <7.5>, <7.8>, <7.9>.

[Bobrow 77] Bobrow, D. G., and Winograd, T., "An overview of KRL, a knowledge representation language," *Cognitive Science* 1, (1), 1977.
<2.4>, <4.1>, <6.7>.

[Boolos 79] Boolos, G., *The Unprovability of Consistency: An essay in modal logic*, Cambridge University Press, 1979.
<6.5>, <6.9>.

[BrownA 76] Brown, A. L., *Qualitative knowledge, causal reasoning, and the localization of failures*, MIT, Technical Report TR-362, 1976.
<7.1>, <7.6>, <7.8>.

[BrownF 79] Brown, F. M., "A theorem prover for meta theory," in *Fourth Workshop on Automated Deduction*, Austin, Texas, 1979.
<6.7>.

[Carbonell 79] Carbonell, J. G., "Towards a self-extending parser," in *ACL-79*, 1979.
<2.5>, <3.3>, <3.4>, <4.1>, <7.5>.

[Charniak 78] Charniak, E., "On the use of framed knowledge in language comprehension," *Artificial Intelligence* 11, 1978, 225-265.
<1.4>, <2.4>, <2.5>, <2.6>.

[Charniak 79] Charniak, E., Riesbeck, C., and McDermott, D., *Artificial Intelligence Programming*, L. E. Eribaum, Baltimore, 1979.
<2.3>, <2.4>, <2.5>, <3.1>, <3.3>, <3.4>, <4.1>, <4.2>.

[Chisolm 66] Chisolm, R. M., *Theory of Knowledge*, Prentice Hall, Englewood Cliffs, NJ, 1966.
<8.2>.

[Clark 78] Clark, K. L., "Negation as failure," in H. Gallaire and J. Minker (ed.), *Logic and Data Bases*, pp. 293-322, Plenum Press, New York, 1978.
<4.1>, <6.5>.

[Cohen 80] Cohen, D., *A program that learns to prove theorems*, Ph.D. thesis, Department of Computer Science, Carnegie Mellon University, 1980.
<2.5>, <3.3>, <6.5>, <7.5>.

[Colby 69] Colby, K. M., and Smith, D. C., "Dialogues between humans and an artificial belief system," in *IJCAI 69*, pp. 319-324, 1969.
<5.1>, <8.1>.

[Colby 73] Colby, K. M., "Simulations of belief systems," in R. C. Schank and K. M. Colby (ed.), *Computer Models of Thought and Language*, pp. 251-286, W. H. Freeman, San Francisco, 1973.
<2.5>, <5.1>, <8.1>.

[Collins 78] Collins, A., "Fragments of a theory of human plausible reasoning," in *TINLAP-78*, pp. 194-201, 1978.
<4.1>, <5.1>, <6.7>, <7.1>, <8.1>.

[Cox 77] Cox, P. T., *Deduction plans: a graphical proof procedure for the first-order predicate calculus*, Department of Computer Science, University of Waterloo, Technical Report Report CS-77-28, 1977.
<2.5>, <3.3>, <3.4>, <7.10>.

[Cox 78] Cox, P. T., "Locating the source of unification failure," in *CSCSI-78*, pp. 20-29, 1978.
<2.5>, <3.3>, <3.4>.

[Crocker 77] Crocker, S. D., *State deltas: a formalism for representing segments of computation*, University of Southern California, Information Sciences Institute, Technical Report RR-77-61, 1977.
<2.2>, <2.7>, <3.2>, <6.3>.

[Dacey 78] Dacey, R., "A theory of conclusions," *Philosophy of Science* 45, 1978, 563-574.
<1.4>, <5.1>, <5.3>.

[Daniel 77] Daniel, L., *Using artificial intelligence techniques to generate project networks*, Department of AI, University of Edinburgh, Technical Report, 1977.
<2.2>, <2.5>, <3.2>, <3.3>, <7.2>, <7.3>.

[Davies 74] Davies, D. J. M., "Representing negation in a planner system," *AISB-74*, 1974, 20-36.
<4.3>.

[DavisM 80] Davis, M., "Notes on the mathematics of non-monotonic logic," *Artificial Intelligence*, 1980. (To appear)
<6.6>.

[DavisR 76] Davis, R., *Applications of meta level knowledge to the construction, maintainance and use of large knowledge bases*, Stanford AI Lab, Technical Report Memo AIM-283, 1976.
<2.5>, <5.1>, <7.1>, <7.4>.

[DavisR 77] Davis, R., "Interactive transfer of expertise: acquisition of new inference rules," in *IJCAI-77*, pp. 321-328, 1977.
<6.5>, <7.1>, <7.4>.

[DavisR 78] Davis, R., "Knowledge acquisition in rule-based systems -- knowledge about representations as a basis for system construction and maintenance," in D. Waterman and F. Hayes-Roth (ed.), *Pattern-Directed Inference Systems*, pp. 99-134, Academic Press, New York, 1978.
<6.5>, <7.1>, <7.4>.

[DavisR 80] Davis, R., and Lenat, D. B., *Knowlege Based Systems in Artificial Intelligence: Two case studies*, McGraw-Hill, New York, 1980.
<2.7>, <5.1>, <6.5>, <7.1>, <7.4>.

[deKleer 76] de Kleer, J., *Local methods for localization of failures in electronic circuits*, MIT, Technical Report Memo 394, 1976.
<2.5>, <7.1>, <7.8>.

[deKleer 77] de Kleer, J., Doyle, J., Steele, G. L. Jr., and Sussman, G. J., "Explicit control of reasoning," in *ACM Symp. on Artificial Intelligence and Programming Languages*, Rochester, New York, 1977. also MIT, Memo 427, also *Artificial Intelligence: A MIT Perspective* (P. H. Winston and R. H. Brown, ed.), MIT Press.
<2.5>, <3.3>, <3.4>, <4.2>, <7.10>.

[deKleer 78a] de Kleer, J., Doyle, J., Rich, C., Steele, G. L. Jr., and Sussman, G. J., *AMORD: a deductive procedure system*, MIT, Technical Report Memo 435, 1978.
<2.4>, <2.5>, <3.3>, <3.4>.

[deKleer 78b] de Kleer, J. and Sussman, G. J., *Propagation of constraints applied to circuit synthesis*, MIT, Technical Report Memo 485, 1978.
<2.5>, <3.3>, <3.4>, <7.1>, <7.9>, <7.10>.

[deKleer 79a] de Kleer, J., *Causal and teleological reasoning in circuit recognition*, Ph.D. thesis, MIT Department of Electrical Engineering and Computer Science, 1979.

<1.2>, <7.1>, <7.8>.

[deKleer 79b] de Kleer, J., "The origin and resolution of ambiguities in causal arguments," in *IJCAI-79*, pp. 197-203, 1979.
<1.2>, <7.1>, <7.8>.

[deKleer 79c] de Kleer, J. and Harris, G., *Truth Maintenance Systems in Problem Solving*, Xerox PARC, Technical Report draft, 1979.
<2.5>, <3.3>, <3.4>, <4.1>, <4.2>.

[Dennett 78] Dennett, D. C., "How to change your mind," in *Brainstorms*, pp. 300-309, Bradford Books, Montgomery, Vermont, 1978.
<1.3>, <8.1>.

[Doyle 76] Doyle, J., *The use of dependency relationships in the control of reasoning*, MIT, Technical Report Working Paper 133, 1976.
<2.5>, <3.3>, <3.4>, <4.2>, <6.7>, <7.10>.

[Doyle 78a] Doyle, J., *Truth maintenance systems for problem solving*, MIT, Technical Report TR-419, 1978.
<1.4>, <2.5>, <3.3>, <3.4>, <4.2>, <7.10>.

[Doyle 78b] Doyle, J., *A glimpse of truth maintenance*, MIT, Technical Report Memo 461, 1978. also Fourth Workshop on Automated Deduction, 90-96, also IJCAI-79, 232-237, also *Artificial Intelligence: an MIT Perspective* (P. H. Winston and R. H. Brown, ed.), MIT Press.
<1.4>, <2.5>, <3.3>, <3.4>, <4.2>, <7.10>.

[Doyle 78c] Doyle, J., *Reflexive interpreters*, MIT Department of Electrical Engineering and Computer Science, Technical Report Unpublished Ph.D. proposal, 1978.
<2.5>, <3.3>, <3.4>, <4.2>, <6.1>, <6.7>, <7.10>.

[Doyle 79a] Doyle, J., "A truth maintenance system," *Artificial Intelligence* 12, (3), 1979.
<1.2>, <1.4>, <2.5>, <3.3>, <3.4>, <4.2>, <7.10>.

[Doyle 79b] Doyle, J., "Historical Annotations and Humble Databases," *submitted to ACM SIGMOD Record*, 1979.
<2.5>, <7.1>, <7.6>.

[Doyle 80] Doyle, J., *Reflexive reasoner: a model for deliberation, action, and introspection*, Ph.D. thesis, MIT, 1980. (forthcoming).
<1.1>, <1.2>, <1.4>, <2.5>, <3.3>, <6.1>, <6.7>, <7.2>, <7.3>, <7.6>, <7.7>, <7.10>.

[Duda 76] Duda, R. O., Hart, P. E., and Nilsson, N. J., "Subjective bayesian methods for rule-based inference systems," in *AFIPS Conference Proceedings, Vol. 45*, pp. 1075-1082, National Computer Conference, 1976.
<5.1>.

[Dummett 73] Dummett, M. A. E., "The justification of deduction," in *Proc. British Academy, Vol. LIX.*, 1973.
<1.4>, <6.1>, <8.2>.

[Fahlman 73] Fahlman, S. E., *A hypothesis-frame system for recognition problems*, MIT, Technical Report Working Paper 73, 1973. summarized in [Minsky 75].
<2.3>, <3.4>.

[Fahlman 74] Fahlman, S. E., "A planning system for robot construction tasks," *Artificial Intelligence* 5, 1974, 1-49.
<1.2>, <2.1>, <2.3>, <3.4>, <7.2>, <7.3>.

[Feldman 74] Feldman, J. A., and Yakimovsky, Y., "Decision theory and artificial intelligence I: a semantics-based region analyzer," *Artificial Intelligence* 5, 1974, 349-371.
<5.2>.

[Feldman 77] Feldman, J. A., and Sproull, R. F., "Decision theory and artificial intelligence II: the hungry monkey," *Cognitive Science* 1, 1977, 158-192.
<1.2>, <5.2>.

[Fikes 71] Fikes, R. E., and Nilsson, N. J., "STRIPS: a new approach to the application of theorem proving to problem solving," *Artificial Intelligence* 2, 1971, 189-208.
<2.2>, <3.2>.

[Fikes 72a] Fikes, R. E., "Monitored execution of robot plans produced by STRIPS," in *IFIP 1971*, pp. 189-194, North-Holland, Amsterdam, 1972.
<2.2>, <7.2>, <7.3>.

[Fikes 72b] Fikes, R. E., Hart, P. E., and Nilsson, N. J., "Learning and executing generalized robot plans," *Artificial Intelligence* 3, 1972, 251-288.
<2.2>, <3.2>, <7.5>, <7.8>.

[Fikes 75] Fikes, R. E., "Deductive retrieval mechanisms for state description models," in *IJCAI-75*, pp. 99-106, 1975.
<2.3>, <2.5>, <3.3>, <5.4>.

[Fikes 80] Fikes, R. E., *A knowledge-based personal assistance project*, SRI International, Technical Report, 1980. in preparation.
<2.5>, <3.3>.

[Fox 77] Fox, M. S., and Mostow, D. J., "Maximal consistent interpretations of errorful data in hierarchically modeled domains," in *IJCAI-77*, pp. 165-171, 1977.
<1.2>.

[Freuder 76] Freuder, E., *Synthesizing constraint expressions*, MIT Artificial Intelligence Lab, Technical Report Memo 360, 1976.
<3.4>.

[Friedman 79] Friedman, L., *Plausible inference: a multi-valued logic for problem solving*, Jet Propulsion Laboratory, Technical Report 79-11, 1979.
<1.4>, <2.5>, <3.3>, <5.5>, <7.5>, <8.3>.

[Gaines 76] Gaines, B. R., "Foundations of fuzzy reasoning," *International Journal of Man-Machine Studies* 8, 1976, 623-668.
<5.5>, <8.1>, <8.2>.

[Gaines 79] Gaines, B. R., "Logical foundations for database systems," *International Journal for Man-Machine Studies* 11, 1979, 481-500.
<5.4>, <5.5>, <6.8>, <6.9>, <8.2>.

[Gaschnig 74] Gaschnig, J., "A constraint satisfaction method for inference making," in *Proc. 12th Allerton Conf. on Circuit and System Theory*, pp. 866-874, University of Illinois, Urbana-Champaign, 1974.
<3.4>.

[Gaschnig 77] Gaschnig, J., "A general backtrack algorithm that eliminates most redundant tests," in *IJCAI-77*, pp. 457, 1977.
<3.4>.

[Gaschnig 78] Gaschnig, J., "Experimental case studies of backtrack vs. Waltz-type vs. new algorithms for satisficing assignment problems," in *CSCSI-78*, pp. 268-277, 1978.
<3.4>.

[Gaschnig 79] Gaschnig, J., *Performance measurement and analysis of certain search algorithms*, CMU, Technical Report CMU-CSD-79-124, 1979.
<3.4>.

[Geiser 75] Geiser, J. R., "Commenting Proofs," in *IJCAI-75*, pp. 34-39, 1975.
<2.5>, <2.7>, <6.5>, <7.1>.

[Giles 76] Giles, R., "A logic for subjective belief," in Harper and Hooker (ed.), *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, pp. 41-72, Reidel, Dordrecht, 1976.
<5.1>, <6.1>.

[Goguen 80] Goguen, J. A., Linde, C., and Weiner, J. L., Reasoning and natural explanations, 1980.
<2.5>, <2.7>, <7.1>, <8.1>.

[Goldstein 74] Goldstein, I. P., *Understanding simple picture programs*, MIT, Technical Report TR-294, 1974.
<7.6>, <7.8>.

[Goldstein 75] Goldstein, I. P., "Summary of MYCROFT: a system for understanding simple picture programs," *Artificial Intelligence* 6, 1975, 249-288.
<7.6>, <7.8>.

[Good 52] Good, I. J., "Rational decisions," *Journal of the Royal Statistical Society B 14*, 1952, 107-114. <1.4>, <5.1>, <5.2>, <6.1>, <8.1>.

[Good 61] Good, I. J., " A causal calculus," *British Journal for the Philosophy of Science* 11, 12, 13, 1961, 305-318, 43-51, 88.
<5.1>, <7.8>.

[Good 68] Good, I. J., "Corroboration, explanation, evolving probability, simplicity and a sharpened razor," *British Journal for the Philosophy of Science* 19, 25, 1968, 123-143, 289.
<1.2>, <5.1>.

[Good 77a] Good, I. J., "Dynamic probability, computer chess, and the measurement of knowledge," in E. W. Elcock and D. Michie (ed.), *Machine Intelligence 8*, pp. 139-150, Wylie, New York, 1977.
<1.4>, <5.1>, <5.2>.

[Good 77b] Good, I. J., "Explicativity: a mathematical theory of explanation with statistical applications," *Proc. Royal Society of London A 354*, 1977, 303-330.
<1.2>, <5.1>.

[Goodman 73] Goodman, N., "The problem of counterfactual conditionals," in *Fact, Fiction, and Forecast (third edition)*, pp. 3-27, Bobbs-Merrill, New York, 1973.
<1.2>, <8.3>, <8.4>, <8.5>.

[Griffiths 67] Griffiths, A. P., *Knowledge and Belief*, Oxford University Press, London, 1967.
<8.2>.

[GumbM 78] Gumb, M. A. and Gumb, R. D., "Logical techniques for pinpointing inconsistencies in the knowledge base," in *41st Annual Conf. American Society of Information Sciences*, 1978.
<1.4>, <2.5>, <3.4>, <6.4>.

[GumbR 78] Gumb, R. D., "Computational aspects of evolving theories," *ACM SIGART Newsletter*, (67), 1978, 13.
<1.4>, <2.5>, <6.4>.

[GumbR 79] Gumb, R. D., *Evolving Theories*, Haven, Flushing, NY, 1979.
<6.4>.

[Haack 78] Haack, S., *Philosophy of Logics*, Cambridge University Press, Cambridge, 1978.
<5.4>, <5.5>, <6.2>, <6.9>.

[Hajek 77] Hajek, P., and Havranek, T., "On generation of inductive hypotheses," *International Journal of Man-Machine Studies* 9, 1977, 415-438.
<1.2>, <5.1>, <6.8>, <8.3>.

[Hajek 78a] Hajek, P., and Havranek, T., "The GUHA method - its aims and techniques," *International Journal of Man-Machine Studies* 10, 1978, 3-22.
<6.8>, <8.3>.

[Hajek 78b] Hajek, P., and Havranek, T., *Mechanizing Hypothesis Formation - Mathematical Foundations of a General Theory*, Springer-Verlag, Heidelberg, 1978.
<1.2>, <5.1>, <6.8>.

[Haralick 79] Haralick, R. M., and Elliot, G. L., "Increasing tree search efficiency for constraint satisfaction problems," in *IJCAI-79*, pp. 356-364, 1979.
<3.4>.

[Harman 73] Harman, G., *Thought*, Princeton University Press, 1973.
<1.2>, <1.3>, <4.1>, <8.1>, <8.2>, <8.3>.

[Harper 76] Harper, W. L., "Rational belief change, Popper functions, and counterfactuals," in Harper and Hooker (ed.), *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, pp. 73-115, Reidel, Dordrecht, 1976.
<1.4>, <5.1>, <6.1>, <8.2>, <8.3>, <8.4>.

[Hart 72] Hart, P. E., Nilsson, N. J., and Robinson, A. E., *A causality representation for enriched robot task domains*, SRI, Technical Report TN-62, 1972.
<1.2>, <2.6>, <5.1>, <7.1>, <7.4>, <7.5>, <7.8>.

[Havel 78] Havel, I. M., and Kramosil, I., "A stochastic approach to robot plan formation," *Kybernetika* 14, 1978, 143-173.
<1.1>, <2.2>, <5.1>, <5.2>, <6.3>, <6.4>.

[Hayes-Roth 78] Hayes-Roth, F., "The role of partial and best matches in knowledge systems," in D. A. Waterman and F. Hayes-Roth (ed.), *Pattern-Directed Inference Systems*, pp. 557-574, Academic Press, New York, 1978.
<1.2>, <2.4>, <4.4>, <8.3>.

[Hayes 70] Hayes, P. J., "Robotologic," in B. Meltzer and D. Michie (ed.), *Machine Intelligence 5*, pp. 533-554, American Elsevier, New York, 1970.
<1.1>, <5.4>, <6.2>, <6.9>, <8.2>.

[Hayes 71] Hayes, P. J., "A logic of actions," in B. Meltzer and D. Michie (ed.), *Machine Intelligence 6*, pp. 495-520, American Elsevier, New York, 1971.
<1.1>, <2.2>, <6.2>, <6.3>.

[Hayes 73] Hayes, P. J., "The frame problem and related problems in artificial intelligence," in A. Elithorn and D. Jones (ed.), *Artificial and Human Thinking*, Josey-Bass, San Francisco, 1973.
<1.1>, <1.4>, <2.2>, <2.4>, <2.5>, <3.2>, <3.3>, <4.1>, <6.1>, <6.3>, <8.4>.

[Hayes 74] Hayes, P. J., "Some problems and non-problems in representation theory," *AISB-74*, 1974, 63-79.
<1.1>, <2.5>, <5.1>, <5.5>, <7.10>, <8.2>, <8.3>.

[Hayes 75] Hayes, P. J., "A representation for robot plans," in *IJCAI-75*, pp. 181-188, 1975.
<2.5>, <3.3>, <7.2>, <7.3>.

[Hayes 77] Hayes, P. J., *The logic of frames*, Department of Computer Science, University of Essex, Technical Report, 1977.
<4.1>.

[Hempel 66] Hempel, C., *Philosophy of Natural Science*, Prentice Hall, Englewood Cliffs, NJ, 1966.
<8.5>.

[Hendrix 73] Hendrix, G. G., "Modeling simultaneous actions and continuous processes," *Artificial Intelligence* 4, 1973, 145-180.
<1.1>, <6.2>, <7.2>.

[Hewitt 72] Hewitt, C. E., *Description and theoretical analysis (using schemata) of PLANNER: a language for proving theorems and manipulating models in a robot*, MIT, Technical Report TR-258, 1972.
<2.2>, <2.3>, <2.4>, <3.2>, <4.1>.

[Hewitt 75a] Hewitt, C. E., "How to use what you know," in *IJCAI-75*, pp. 189-198, 1975.
<2.3>, <2.4>, <3.4>, <6.3>, <6.10>, <7.10>.

[Hewitt 75b] Hewitt, C. E., and Smith, B., "Towards a programming apprentice," *IEEE Transactions on Software Engineering* SE-1, 1975, 26-45.
<7.1>, <7.7>.

[Hintikka 62] Hintikka, J., *Knowledge and Belief*, Cornell University Press, Ithica, 1962.
<6.1>, <6.9>, <8.2>.

[Jaegermann 78] Jaegermann, M., "Information storage and retrieval systems with incomplete information I," *Fundamenta Informaticae* II, 1978, 17-41.
<6.8>.

[Joshi 75a] Joshi, A. K. and Rosenschein, S. J., "Some problems of inferencing: Relation of inferencing to decomposistion of predicates," in *5th Int. Conf. on Computational Linguistics*, pp. 47-70, 1975.
<2.4>, <4.4>.

[Joshi 75b] Joshi, A. K., and Rosenschein, S. J., "A formalism for relating lexical and pragmatic information: its relevance to recognition and generation," in *TINLAP-1*, pp. 79-83, 1975.
<1.1>, <4.4>.

[Joshi 78a] Joshi, A. K., "Some extensions of a system for inferencing on partial information," in F. Hayes-Roth and D. Waterman (ed.), *Pattern Directed Inference Systems*, pp. 241-257, Academic Press, New York, 1978.
<1.2>, <2.4>, <4.4>.

[Joshi 78b] Joshi, A. K., "A note on partial match of descriptions: can one simultaneously question (retrieve) and inform (update)?," in *TINLAP-2*, 1978.
<1.4>, <4.1>, <7.4>.

[Kahn 77] Kahn, K. M., and Gorry, G. A., "Mechanizing temporal knowledge," *Artificial Intelligence* 9, 1977, 87-108.
<2.5>, <3.3>, <6.2>.

[Katz 76] Katz, S., and Manna, Z., "Logical analysis of programs," *CACM* 19, 1976, 188-206.
<2.5>, <3.4>, <7.6>.

[Kornfeld 79a] Kornfeld, W. A., "ETHER - a parallel problem solving system," in *IJCAI-79*, pp. 490-492, 1979.
<2.3>, <2.4>, <4.1>, <6.3>.

[Kornfeld 79b] Kornfeld, W. A., Using parallel processing for problem solving, Master's thesis, MIT, 1979.
<2.3>, <2.4>, <4.1>, <6.3>.

[Kowalski 78] Kowalski, R., "Logic for data description," in H. Gallaire and J. Minker (ed.), *Logic and Data Bases*, pp. 77-103, Plenum Press, New York, 1978.
<1.4>.

[Kramosil 75] Kramosil, I., "A note on deduction rules with negative premises," in *IJCAI-75*, pp. 53-56, 1975.
<4.1>, <6.6>.

[Kreisel 77] Kreisel, G., "On the kind of data needed for a theory of proofs," in R. Gandy and M. Hyland (ed.), *Logic Colloquium 76*, pp. 111-128, North-Holland, Amsterdam, 1977.
<6.5>.

[Kuhn 62] Kuhn, T., *The Structure of Scientific Revolutions*, University of Chicago Press, 1962.
<8.5>.

[Lakatos 76] Lakatos, I., *Proofs and Refutations: the logic of mathematical discovery*, Cambridge University Press, Cambridge, 1976.
<8.5>.

[Latombe 76] Latombe, J.-C., *Artificial intelligence in computer-aided design: the TROPIC system*, SRI, Technical Report TN-125, 1976.
<2.5>, <3.4>, <7.3>.

[Latombe 79] Latombe, J.-C., "Failure processing in a system for designing complex assemblies," in *IJCAI-79*, pp. 508-515, 1979.
<1.2>, <2.5>, <3.4>, <7.3>, <7.8>.

[Lehrer 69] Lehrer, K., and Paxson, T. Jr., "Knowledge: undefeated justified true belief," *Journal of Philosophy* LXVI, 1969, 225-237.
<4.1>, <6.1>, <8.2>.

[Lehrer 71] Lehrer, K., "Induction and conceptual change," *Synthese* 23, 1971, 206-225.
<5.2>, <6.1>, <8.2>, <8.3>.

[Lehrer 75] Lehrer, K., "Induction, rational acceptance, and minimally inconsistent sets," in G. Maxwell and R. M. Anderson, Jr. (ed.), *Induction, Probability, and Confirmation*, pp. 295-323, University of Minnesota Press, Minneapolis, 1975.
<5.2>, <6.1>, <8.2>, <8.3>.

[Lenat 77] Lenat, D. B., "The ubiquity of discovery," *Artificial Intelligence* 9, 1977, 257-285.
<7.5>.

[Levesque 79] Levesque, H. J., *The representation of incomplete knowledge as applied to a logic of sentences*, Department of Computer Science, University of Toronto, Technical Report draft, 1979.
<4.1>, <6.6>.

[Lewis 73] Lewis, D., *Counterfactuals*, Harvard University Press, Cambridge, 1973.
<8.4>.

[Lipski 76] Lipski, W. Jr., "Informational systems with incomplete information," in S. Michaelson and R. Milner (ed.), *Proc. 3rd Int. Symp. on Automata, Languages and Programming*, pp. 120-130, Edinburgh University Press, Edinburgh, 1976.
<6.4>, <6.8>.

[Lipski 77] Lipski, W. Jr., "On the logic of incomplete information," in J. Gruska (ed.), *Proc. Symp. on Mathematical Foundations of Computer Science*, pp. 374-381, Springer-Verlag, Berlin, 1977.
<6.8>.

[Lipski 78] Lipski, W. Jr., *On semantic issues connected with incomplete information data bases*, Institute of Computer Science, Polish Academy of Sciences, Technical Report 325, 1978.
<6.4>, <6.8>, <7.1>.

[London 78a] London, P. E., "A dependency-based modelling mechanism for problem solving," in *AFIPS Conference Proceedings, Vol. 47*, pp. 263-274, National Computer Conference, 1978.
<1.1>, <2.5>, <2.6>, <3.3>, <3.4>, <7.2>, <7.3>.

[London 78b] London, P. E., *Dependency networks as a representation for modelling in general problem solvers*, Department of Computer Science, University of Maryland, Technical Report TR-698, 1978.
<1.1>, <2.2>, <2.5>, <2.6>, <3.2>, <3.3>, <3.4>, <7.2>, <7.3>

[Marcus 79] Marcus, L., *State deltas that remember*, USC-ISI, Technical Report Microver Note #22, 1979.
<2.7>.

[May 76] May, S., and Harper, W., "Toward an optimization procedure for applying minimum change principles in probability kinematics," in Harper and Hooker (ed.), *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, pp. 137-166, Reidel, Dordrecht, 1976.
<1.4>, <5.1>.

[McAllester 78] McAllester, D. A., *A three-valued truth maintenance system*, MIT, Technical Report Memo 473, 1978.
<2.5>, <3.3>, <3.4>, <4.1>.

[McAllester 79a] McAllester, D. A., Canonical naming under equality, Master's thesis, MIT, EECS, 1979.
<2.5>, <3.3>, <4.1>.

[McAllester 79b] McAllester, D. A., *An outlook on propositional reasoning in deduction*, MIT, Technical Report draft, 1979.
<2.5>, <3.3>, <3.4>, <4.1>.

[McCarthy 68] McCarthy, J., "Programs with common sense," in M. Minsky (ed.), *Semantic Information Processing*, pp. 403-418, MIT Press, Cambridge, 1968.
<6.2>, <6.3>, <6.9>, <8.2>.

[McCarthy 69] McCarthy, J., and Hayes, P. J., "Some philosophical problems from the standpoint of artificial intelligence," in B. Meltzer and D. Michie (ed.), *Machine Intelligence 4*, pp. 463-502, American Elsevier, New York, 1969.
<1.1>, <4.1>, <6.1>, <6.2>, <6.3>, <6.9>, <8.2>.

[McCarthy 77] McCarthy, J., "Epistemological problems of artificial intelligence," in *IJCAI-77*, pp. 1038-1044, 1977.
<1.1>, <4.5>, <6.3>, <6.6>, <8.2>.

[McCarthy 80] McCarthy, J., "Circumscription: a form of non-monotonic reasoning," *Artificial Intelligence* 13, 1980, to appear.
<4.5>.

[McDermott 74a] McDermott, D., *Assimilation of new information by a natural language-understanding system*, MIT, Technical Report AI-TR-291, 1974.
<1.2>, <1.4>, <2.3>, <2.7>, <3.4>, <4.1>, <7.1>, <7.4>.

[McDermott 74b] McDermott, D., and Sussman, G. J., *The CONNIVER reference manual*, MIT, Technical Report Memo 259a, 1974.
<2.3>, <2.4>.

[McDermott 75] McDermott, D., *Very large PLANNER-type data bases*, MIT, Technical Report Memo 339, 1975.
<2.3>, <2.5>, <3.3>.

[McDermott 77] McDermott, D., *Flexibility and efficiency in a computer program for designing circuits*, MIT, Technical Report TR-402, 1977.
<2.2>, <2.3>, <2.4>, <2.5>, <2.7>, <3.2>, <3.3>, <4.1>, <6.2>, <6.3>, <7.1>, <7.2>, <7.3>, <7.6>, <7.10>.

[McDermott 78a] McDermott, D., "Planning and acting," *Cognitive Science* 2, 1978, 71-109.
<2.2>, <2.3>, <2.4>, <2.5>, <2.7>, <3.2>, <3.3>, <4.1>, <6.2>, <6.3>, <7.1>, <7.2>, <7.3>, <7.6>, <7.10>.

[McDermott 78b] McDermott, D., "The last survey of knowledge representation," *AISB-78*, 1978, 206-221.
<2.5>, <3.3>, <6.6>.

[McDermott 78c] McDermott, D., and Doyle, J., *Non-monotonic logic I*, MIT, Technical Report Memo 486, 1978. also, *Artificial Intelligence 13*, to appear.
<1.1>, <2.5>, <3.3>, <4.1>, <4.2>, <6.5>, <6.6>.

[McDermott 79a] McDermott, D., and Doyle, J., "Non-monotonic logic I (extended abstract)," in *Fourth Workshop on Automated Deduction*, Austin, Texas, 1979.
<1.1>, <2.5>, <3.3>, <4.1>, <4.2>, <6.5>, <6.6>.

[McDermott 79b] McDermott, D., and Doyle, J., "An introduction to non-monotonic logic," in *IJCAI-79*, pp. 562-567, 1979.
<1.1>, <2.5>, <3.3>, <4.1>, <4.2>, <6.5>, <6.6>.

[Miller 79] Miller, M. L., *Planning and debugging in elementary programming*, Ph.D. thesis, MIT EECS, 1979.
<2.6>, <2.7>, <4.1>, <7.6>.

[Minsky 61] Minsky, M., "Steps towards artificial intelligence," in E. A. Feigenbaum and J. Feldman (ed.), *Computers and Thought*, pp. 406-450, McGraw-Hill, New York, 1961.
<5.1>, <7.5>, <7.8>, <8.3>.

[Minsky 62] Minsky, M., "Problems of formulation for artificial intelligence," in *Proc. Symp. on Mathematical Problems in Biology*, pp. 35-46, American Mathematical Society Providence, 1962.
<7.5>, <7.8>, <8.3>.

[Minsky 65] Minsky, M., "Matter, mind, and models," in M. Minsky (ed.), *Semantic Information Processing*, pp. 425-432, MIT Press, Cambridge, 1965.
<6.7>, <7.5>, <7.10>, <8.1>, <8.2>.

[Minsky 75] Minsky, M., "A framework for representing knowledge," in P. Winston (ed.), *The Psychology of Computer Vision*, McGraw-Hill, New York, 1975.
<2.3>, <2.4>, <4.1>, <8.1>.

[Minsky 79] Minsky, M., *K-lines: a theory of memory*, MIT, Technical Report Memo 516, 1979.
<1.3>, <2.7>, <7.5>, <7.10>, <8.1>.

[Moore 75] Moore, R. C., *Reasoning from incomplete knowledge in a procedural deduction system*, MIT, Technical Report AI-TR-347, 1975.
<1.1>, <2.3>, <2.4>, <6.3>.

[Moore 77] Moore, R. C., "Reasoning about knowledge and action," in *IJCAI-77*, pp. 223-227, 1977.
<6.1>, <6.3>, <6.10>.

[Moore 79a] Moore, R. C., *Reasoning about knowledge and action*, Ph.D. thesis, MIT Department of Electrical Engineering and Computer Science, 1979.
<6.1>, <6.3>, <6.10>.

[Moore 79b] Moore, R. C., and Hendrix, G. G., *Computational models of beliefs and the semantics of belief sentences*, SRI International, Technical Report TN-187, 1979.
<6.9>, <8.1>.

[Moriconi 77] Moriconi, M., *A system for incrementally designing and verifying programs*, USC-ISI, Technical Report RR-77-65, 1977.
<2.5>, <3.3>, <7.6>, <7.7>.

[Nevins 74] Nevins, A. J., "A human-oriented logic for automatic theorem proving," *JACM* 21, 1974, 606-621.
<2.5>, <3.3>, <3.4>.

[Pereira 79a] Pereira, L. M., *Backtracking intelligently in AND/OR trees*, Centro de Informatica das Universidades de Lisboa, Technical Report, 1979.
<3.4>.

[Pereira 79b] Pereira, L. M., and Porto, A., *Intelligent backtracking and sidetracking in Horn clause programs - the theory*, Universidade de Nova Lisboa, Technical Report CIUNL report #2/79, 1979.
<3.4>.

[Quine 53] Quine, W. V., "Two dogmas of empiricism," in *From a Logical Point of View*, Harvard University Press, Cambridge, 1953.
<1.4>, <6.4>, <8.2>.

[Quine 70] Quine, W. V., *Philosophy of Logic*, Prentice-Hall, Englewood Cliffs, 1970.
<1.2>.

[Quine 78] Quine, W. V., and Ullian, J. S., *The Web of Belief (second edition)*, Random House, New York, 1978.
<1.2>, <1.4>, <6.1>, <8.1>, <8.2>.

[Raphael 71] Raphael, B., "The frame problem in problem-solving systems," in N. V. Findler and B. Meltzer (ed.), *Artificial Intelligence and Heuristic Programming*, pp. 159-169, Edinburgh University Press, Edinburgh, 1971.
<1.1>, <2.2>, <2.3>, <6.3>.

[Raphael 76] Raphael, B., *The Thinking Computer: Mind Inside Matter*, W. H. Freeman, San Fransisco, 1976.
<1.1>, <2.2>, <6.3>.

[Reiter 78a] Reiter, R., "On reasoning by default," in *TINLAP-2*, 1978.
<1.1>, <1.4>, <4.1>, <4.3>.

[Reiter 78b] Reiter, R., "On closed world data bases," in H. Gallaire and J. Minker (ed.), *Logic and Data Bases*, Plenum Press, New York, 1978.
<4.3>, <6.6>, <6.8>.

[Reiter 80] Reiter, R., "A logic for default reasoning," *Artificial Intelligence* 13, 1980, (to appear).
<1.1>, <1.4>, <2.5>, <3.3>, <4.1>, <6.6>.

[Rescher 64] Rescher, N., *Hypothetical Reasoning*, North Holland, Amsterdam, 1964.
<1.4>, <6.1>, <6.4>, <8.4>.

[Rich 79] Rich, C., Shrobe, H. E., and Waters, R. C., *Computer aided evolutionary design for software engineering*, MIT, Technical Report Memo 506, 1979.
<2.5>, <7.6>, <7.7>, <7.9>.

[Rieger 76a] Rieger, C., "One system for two tasks: A commonsense algorithm memory that solves problems and comprehends language," *Artificial Intelligence*, January 1976.
<2.2>, <2.7>.

[Rieger 76b] Rieger, C. and Grinberg, M. R., *The causal representation and simulation of physical mechanisms*, Department of Computer Science, University of Maryland, Technical Report 495, November 1976.
<2.7>, <7.6>, <7.8>.

[Rieger 77a] Rieger, C., "Spontaneous computation in cognitive models," *Cognitive Science* 1, (3), 1977.
<2.4>.

[Rieger 77b] Rieger, C., and London, P., "Subgoal protection and unravelling during plan synthesis," in *IJCAI-77*, pp. 487-493, 1977.
<2.4>, <3.2>, <3.4>, <7.3>.

[Rieger 77c] Rieger, C. and Grinberg, M. R., "The declarative representation and procedural simulation of causality in physical mechanisms," in *IJCAI-77*, 1977.
<2.7>, <7.6>, <7.8>.

[Roberts 77] Roberts, R. B., and Goldstein, I. P., *The FRL manual*, MIT, Technical Report Memo 409, 1977.
<2.4>, <4.1>.

[Rosenberg 78] Rosenberg, S., "SLEUTH: an intelligent noticer," in *CSCSI-78*, pp. 1-10, 1978.
<1.4>, <2.4>, <4.1>, <7.10>.

[Rosenberg 79] Rosenberg, S., "Reasoning in incomplete domains," in *IJCAI-79*, pp. 735-737, 1979.
<1.4>, <2.4>, <4.1>, <7.10>.

[Rosenschein 75] Rosenschein, S. J., *Structuring a pattern space, with applications to lexical information and event interpretation*, Department of Computer and Information Science, University of Pennsylvania, Technical Report, 1975.
<4.4>, <7.1>, <7.8>, <7.10>.

[Rosenschein 76] Rosenschein, S. J., "How does a system know when to stop inferencing," *American Journal of Computational Linguistics*, (Microfiche 36), 1976.

<4.4>, <7.10>.

[Rosenschein 79] Rosenschein, S. J., "Abstract theories of discourse and the formal specification of programs that converse," in *Proc. of the Sloan Workshop on Computational Aspects of Linguistic Structure and Discourse Setting*, Cambridge University Press, Cambridge, 1979.
<6.1>, <6.2>, <6.4>.

[Rubin 75] Rubin, A. D., *Hypothesis formation and evaluation in medical diagnosis*, MIT, Technical Report TR-316, 1975.
<1.4>, <2.3>, <3.4>, <4.1>.

[Rulifson 72] Rulifson, J. F., Waldinger, R. J., and Derksen, J. A., "A language for writing problem-solving programs," in *IFIP 71*, pp. 201-205, North-Holland, Amsterdam, 1972.
<2.3>, <2.4>.

[Rychener 78] Rychener, M. D., and Newell, A., "An instructible production system: basic design issues," in D. Waterman and F. Hayes-Roth (ed.), *Pattern-Directed Inference Systems*, pp. 135-153, Academic Press, New York, 1978.
<2.4>, <7.4>, <7.7>.

[Rychener 79] Rychener, M. D., "A semantic network of production rules in a system for describing computer structures," in *IJCAI-79*, pp. 738-743, 1979.
<2.4>, <7.4>.

[Sacerdoti 74] Sacerdoti, E. D., "Planning in a hierarchy of abstraction spaces," *Artificial Intelligence* 5, 1974, 115-135.
<2.2>, <3.2>.

[Sacerdoti 77] Sacerdoti, E. D., *A Structure for Plans and Behavior*, American Elsevier, New York, 1977.
<2.2>, <2.3>, <7.2>, <7.3>.

[Sacerdoti 79] Sacerdoti, E. D., "Problem solving tactics," in *IJCAI-79*, pp. 1077-1085, 1979.
<1.1>, <2.2>, <3.4>, <7.2>, <7.3>, <7.6>, <7.10>.

[Sandewall 72] Sandewall, E., "An approach to the frame problem, and its implementation," in B. Meltzer and D. Michie (ed.), *Machine Intelligence 7*, pp. 195-204, Wiley, New York, 1972.
<1.1>, <2.2>, <4.1>, <6.3>, <7.2>.

[Sandewall 79] Sandewall, E., "Biological software," in *IJCAI-79*, pp. 744-747, 1979.
<7.7>.

[Schmidt 77] Schmidt, C. F., and Sridharan, N. S., "Plan recognition using a hypothesize and revise paradigm: an example," in *IJCAI-77*, pp. 480-486, 1977.
<1.2>, <2.2>, <2.5>, <2.6>, <3.4>, <7.1>.

[Schmidt 78] Schmidt, C. F., Sridharan, N. S., and Goodson, J. L., "The plan recognition problem: an intersection of psychology and artificial intelligence," *Artificial Intelligence* 11, 1978, 45-83.
<1.2>, <2.2>, <2.5>, <2.6>, <7.1>.

[Schubert 79] Schubert, L. K., "Problems with Parts," in *IJCAI-79*, pp. 778-784, 1979.
<4.1>.

[Schwind 78] Schwind, C. B., "Representing actions by state logic," *AISB-78*, 1978, 304-308.
<1.1>, <6.2>, <6.9>.

[Scriven 59] Scriven, M., "Truisms as the grounds for historical explanations," in P. Gardiner (ed.), *Theories of History*, pp. 443-475, Free Press, Glencoe, Ill., 1959.

<4.1>, <5.1>, <7.1>, <8.2>.

[Shafer 76] Shafer, R., *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, 1976.
<1.4>, <5.1>, <6.1>, <8.2>.

[Shrobe 79a] Shrobe, H. E., *Dependency directed reasoning for complex program understanding*, MIT, Technical Report TR-503, 1979.
<2.2>, <2.5>, <3.2>, <3.3>, <4.2>, <6.1>, <6.3>, <7.2>, <7.6>, <7.7>, <7.9>, <7.10>.

[Shrobe 79b] Shrobe, H. E., "Dependency directed reasoning in the analysis of programs which modify complex data structures," in *IJCAI-79*, pp. 829-835, 1979.
<1.1>, <2.2>, <2.5>, <3.3>, <3.4>, <4.2>, <7.1>, <7.3>, <7.8>, <7.10>.

[Shrobe 79c] Shrobe, H. E., "Explicit control of reasoning in the programmer's apprentice," in *Fourth Workshop on Automated Deduction*, pp. 97-102, Austin, Texas, 1979.
<2.5>, <3.3>, <4.2>, <7.10>.

[Simon 66] Simon, H. A., and Rescher, N., "Cause and counterfactual," *Philosophy of Science* 33, 1966, 323-340.
<1.2>, <8.4>.

[Skyrms 66] Skyrms, B., *Choice and Chance: An Introduction to Inductive Logic*, Dickenson Publishing Co., Belmont, CA, 1966.
<1.4>, <5.1>, <6.1>, <8.2>, <8.3>, <8.4>, <8.5>.

[Sosa 75] Sosa, E., *Causation and Conditionals*, Oxford University Press, London, 1975.
<8.2>, <8.4>.

[Sridharan 75] Sridharan, N. S., *The architecture of BELIEVER - a system for interpreting human actions*, Department of Computer Science, Rutgers University, Technical Report Memo CBM-TR-46, 1975.
<2.2>, <2.7>, <4.4>, <7.1>, <7.8>.

[Sridharan 76] Sridharan, N. S., "The frame and focus problems in AI: discussion in relation to the BELIEVER system," *AISB-76*, 1976, 322-333.
<1.1>, <1.2>, <2.2>, <2.5>, <3.4>.

[Sridharan 77] Sridharan, N. S., and Hawrusik, F., "Representation of actions that have side-effects," in *IJCAI-77*, pp. 265-266, 1977.
<1.1>, <1.2>, <2.2>, <2.6>, <2.7>, <3.2>, <3.4>.

[Sridharan 78] Sridharan, N. S., and Schmidt, C. F., "Knowledge-directed inference in BELIEVER," in D. Waterman and F. Hayes-Roth (ed.), *Pattern-Directed Inference Systems*, Academic Press, New York, 1978.
<1.4>, <2.2>, <2.7>, <7.1>.

[Srinivas 78] Srinivas, S., "Error recovery in robots through failure reason analysis," in *AFIPS Conference Proceedings, Vol. 47*, pp. 275-282, National Computer Conference, 1978.
<7.3>, <7.6>, <7.8>.

[Srinivasan 76] Srinivasan, C. V., "The architecture of coherent information system: a general problem solving system," *IEEE Transactions on Computers* C-25, 1976, 390-402.
<1.1>, <1.2>, <2.7>, <3.4>, <5.4>.

[Stallman 77] Stallman, R. M., and Sussman, G. J., "Forward reasoning and dependency-directed backtracking in a system for computer-aided circuit analysis," *Artificial Intelligence* 9, 1977, 135-196.
<2.5>, <3.3>, <3.4>, <4.1>, <7.1>, <7.5>, <7.8>, <7.9>.

[Stansfield 78] Stansfield, J. L., *Integrating truth maintenance systems with propagation of strength of belief as a means of addressing the fusion problem*, MIT, Technical Report draft, 1978.
<2.5>, <3.3>, <5.1>.

[Statman 74] Statman, R., *Structural complexity of proofs*, Ph.D. thesis, Stanford University, Department of Mathematics, 1974.
<6.5>.

[Steele 78] Steele, G. L. Jr., and Sussman, G. J., *Constraints*, MIT, Technical Report Memo 502, 1978.
<2.4>, <2.5>, <3.3>, <3.4>, <7.1>, <7.9>.

[Stepankova 76a] Stepankova, O., and Havel, I. M., "A logical theory of robot problem solving," *Artificial Intelligence* 7, 1976, 129-161.
<1.1>, <6.3>, <6.4>.

[Stepankova 76b] Stepankova, O., and Havel, I. M., "Incidental and state-dependent phenomena in problem solving," *AISB-78*, 1976, 266-278. and *Kybernetika 13*, pp. 421-438.
<1.1>, <2.2>, <6.4>.

[Stepankova 78] Stepankova, O., "Planning in uncertain environments through situation calculus," *AISB-76*, 1978, 330-339.
<1.1>, <2.2>, <6.3>, <6.10>.

[Stickel 77] Stickel, M. E., *Mechanical theorem proving and artificial intelligence languages*, CMU, Technical Report, 1977.
<5.4>, <6.5>.

[Suppes 77] Suppes, P., "A survey of contemporary learning theories," in R. E. Butts and J. Hintikka (ed.), *Foundational Problems in the Special Sciences*, pp. 217-239, Reidel, Dordrecht, 1977.
<1.3>, <8.1>, <8.3>.

[Sussman 71] Sussman, G. J., Winograd, T., and Charniak, E., *MICRO-PLANNER reference manual*, MIT, Technical Report Memo 203a, 1971.
<2.3>, <2.4>, <3.1>, <4.1>.

[Sussman 72] Sussman, G. J., and McDermott, D., "From PLANNER to CONNIVER - a genetic approach," in *Proc. AFIPS FJCC*, pp. 1171-1179, 1972.
<2.3>, <2.4>, <3.1>.

[Sussman 75a] Sussman, G. J., *A Computer Model of Skill Acquisition*, American Elsevier, New York, 1975.
<2.3>, <7.5>, <7.6>, <7.7>, <7.8>.

[Sussman 75b] Sussman, G. J., and Stallman, R. M., "Heuristic techniques in computer-aided circuit analysis," *IEEE Transactions on Circuits and Systems* CAS-22, 1975, 857-865.
<2.5>, <7.1>.

[Sussman 77] Sussman, G. J., "Electrical design: a problem for artificial intelligence research," in *IJCAI-77*, pp. 894-900, 1977.
<3.3>, <3.4>, <7.6>.

[Sussman 79] Sussman, G. J., Holloway, J., and Knight, T. F. Jr., *Computer aided evolutionary design for digital integrated systems*, MIT, Technical Report Memo 526, 1979.
<7.7>.

[Swartout 77] Swartout, W., *A digitalis therapy advisor with explanations*, MIT LCS, Technical Report TR-176, 1977.
<2.4>, <5.1>, <7.1>, <7.3>.

[Tate 76] Tate, A., *Project planning using a hierarchic non-linear planner*, Department of Artificial Intelligence, University of Edinburgh, Technical Report 25, 1976.
<2.2>, <2.3>, <2.5>.

[Tate 77] Tate, A., "Generating project networks," in *IJCAI-77*, pp. 888-893, 1977.
<2.2>, <2.3>, <2.5>, <3.2>, <3.4>.

[Teller 76] Teller, P., "Conditionalization, observation, and change of preference," in Harper and Hooker (ed.), *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, pp. 205-259, Reidel, Dordrecht, 1976.
<1.4>, <5.1>.

[Thompson 79] Thompson, A., "Network truth maintenance for deduction and modelling," in *IJCAI-79*, pp. 877-879, 1979.
<1.1>, <2.3>, <2.5>, <2.6>, <3.2>, <3.3>, <7.2>, <7.3>.

[Tukey 60] Tukey, J. W., "Conclusions vs decisions," *Technometrics* 2, 1960, 423-433.

<5.2>, <5.3>, <5.8>, <8.1>, <8.2>, <8.3>.

[Turner 78] Turner, R., *Counterfactuals without possible worlds*, University of Essex, Technical Report, 1978.
<8.2>, <8.4>.

[Waldinger 77] Waldinger, R., "Achieving several goals simultaneously," in E. W. Elcock and Donald Michie (ed.), *Machine Intelligence*, Ellis Horwood, Chichester, 1977.
<1.1>, <2.3>, <3.2>, <7.8>.

[Waterman 75] Waterman, D. A., "Adaptive production systems," in *IJCAI-75*, pp. 296-303, 1975.
<1.4>, <2.4>, <7.5>.

[Weiner 79] Weiner, J. L., *The structure of natural explanation: theory and application*, System Development Corp., Santa Monica CA, Technical Report SP-3980, 1979.
<2.5>, <2.7>, <7.1>, <8.1>.

[Weyhrauch 78] Weyhrauch, R. W., *Prolegomena to a theory of mechanized formal reasoning*, SAI, Technical Report AIM-315, 1978.
<2.5>, <6.7>.

[Winston 75] Winston, P. H., "Learning structural descriptions from examples," in *The Psychology of Computer Vision*, pp. 157-209, McGraw-Hill, New York, 1975.
<1.4>, <7.5>.

[Winston 78] Winston, P. H., "Learning by creating and justifying transfer frames," *Artificial Intelligence* 10, 1978, 147-172.
<2.7>, <7.5>.

[Zadeh 75] Zadeh, L., "Fuzzy logic and approximate reasoning," *Synthese* 30, 1975, 407-428.
<5.5>.

[Zadeh 79] Zadeh, L., "Approximate reasoning based on fuzzy logic," in *IJCAI-79*, pp. 1004-1010, 1979.
<5.5>.