

This paper appears in the *Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI-93)*, Washington, D.C., 1993.

Reasoning With Characteristic Models

Henry A. Kautz, Michael J. Kearns, and Bart Selman
AI Principles Research Department
AT&T Bell Laboratories
Murray Hill, NJ 07974
{kautz, mkearns, selman}@research.att.com

Abstract

Formal AI systems traditionally represent knowledge using logical formulas. We will show, however, that for certain kinds of information, a model-based representation is more compact and enables faster reasoning than the corresponding formula-based representation. The central idea behind our work is to represent a large set of models by a subset of *characteristic* models. More specifically, we examine model-based representations of Horn theories, and show that there are large Horn theories that can be exactly represented by an exponentially smaller set of characteristic models.

In addition, we will show that deduction based on a set of characteristic models takes only linear time, thus matching the performance using Horn theories. More surprisingly, *abduction* can be performed in polynomial time using a set of characteristic models, whereas abduction using Horn theories is NP-complete.

1 Introduction

Logical formulas are the traditional means of representing knowledge in formal AI systems [McCarthy and Hayes, 1969]. The information implicit in a set of logical formulas can also be captured by explicating recording the set of models (truth assignments) that satisfy the formulas. Indeed, standard databases are naturally viewed as representations of a single model. However, when dealing with incomplete information, the set of models is generally much too large to be represented

explicitly, because a different model is required for each possible state of affairs. Logical formulas can often provide a compact representation of such incomplete information.

There has, however, been a growing dissatisfaction with the use of logical formulas in actual applications, both because of the difficulty in writing consistent theories, and the tremendous computation problems in reasoning with them. An example of the reaction against the traditional approach is the growing body of research and applications using case-based reasoning (CBR) [Kolodner, 1991]. By identifying the notion of a “case” with that of a “model”, we can view the CBR enterprise as an attempt to bypass (or reduce) the use of logical formulas by storing and directly reasoning with a set of models. While the practical results of CBR are promising, there has been no formal explanation of how model-based representations could be superior to formula-based representations.¹

In this paper, we will prove that for certain kinds of information, a model-based representation is much more compact and enables much faster reasoning than the corresponding formula-based representation. The central idea behind our work is to represent a large set of models by a subset of *characteristic* models, from which all others can be generated efficiently. More specifically, we examine model-based representations of Horn theories, and show that there are large Horn theories that can be exactly represented by exponentially smaller sets of characteristic models.

In addition, we will show that deduction based on a set of characteristic models takes only linear time, thus matching the performance using Horn theories [Dowling and Gallier, 1984]. More surprisingly, *abduction* can be performed in polynomial time using a set of characteristic models, whereas abduction using Horn theories is NP-complete [Selman and Levesque, 1990]. This result is particularly interesting because very few other tractable classes of abduction problems are known [Bylander *et al.*, 1989; Selman, 1990].

2 Horn Theories and Characteristic Models

We assume a standard propositional language, and use $a, b, c, d, p,$ and q to denote propositional variables. A *literal* is either a propositional variable, called a positive literal, or its negation, called a negative literal. A *clause* is a disjunction of literals, and can be represented by the set of literals it contains. A clause C *subsumes* a clause C' iff all the literals in C appear in C' . A set (conjunction) of clauses is called a *clausal theory*, and is represented by the Greek letter Σ . We use n to denote the

¹This is, of course, an oversimplified description of CBR; most CBR systems incorporate both a logical background theory and a set of cases.

length of a theory (*i.e.*, number of literals). A clause is *Horn* if and only if it contains at most one positive literal; a set of such clauses is called a *Horn theory*. (Note that we are not restricting our attention to definite clauses, which contain exactly one positive literal. A Horn clause may be completely negative.)

A *model* is a complete truth assignment for the variables (equivalently, a mapping from the variables to $\{0, 1\}$). We sometimes write a model as a bit vector, *e.g.*, $[010\dots]$, to indicate that variable a is assigned false, b is assigned true, c is assigned false, etc. A model *satisfies* a theory if the theory evaluates to “true” in the model. Another way of saying this is that the theory is *consistent* with the model. When we speak of the “models of a theory Σ ,” we are referring to the set of models that satisfy the theory. This set is denoted by $models(\Sigma)$.

We begin by developing a model-theoretic characterization of Horn theories. The *intersection* of a pair of models is defined as the model that assigns “true” to just those variables that are assigned “true” by both of the pair. The *closure* of a set of models is obtained by repeatedly adding the intersection of the elements of the set to the set until no new models are generated.

Definition: Intersection and Closure

The intersection of models m_1 and m_2 over a set of variables is given by

$$[m_1 \cap m_2](x) \stackrel{def}{=} \begin{cases} 1 & \text{if } m_1(x) = m_2(x) = 1 \\ 0 & \text{otherwise} \end{cases}$$

Where M is a set of models, $closure(M)$ is the smallest set containing M that is closed under \cap .

To illustrate the various definitions given in this section, we will use an example set M_0 of models throughout. Let $M_0 = \{[1110], [0101], [1000]\}$. The closure of this set is given by $M'_0 = M_0 \cup \{[0100], [0000]\}$. See Figure 1.

The notion of closure is particularly relevant in the context of Horn theories:

Theorem 1 [McKinsey 1943]² *A theory Σ is equivalent to a Horn theory if and only if $models(\Sigma)$ is closed under intersection.*

Thus there is a direct correspondence between Horn theories and sets of models that are closed under intersection. For example, consider the closure M'_0 of the

²The proof in McKinsey is for first-order equational theories, and in fact led to the original definition of a Horn clause. A simpler, direct proof for the propositional case appears in [Dechter and Pearl, 1992].




Figure 1: The circled models are M'_0 , which is the closure of the example set of models M_0 .

models in set M_0 defined above. It is not difficult to verify that the models in the closure are exactly the models of the the Horn theory $\Sigma_0 = \{\neg a \vee \neg b \vee c, \neg b \vee \neg c \vee a, \neg a \vee \neg d, b \vee \neg d, b \vee \neg c\}$.

Next, we define the notion of a *characteristic* model. The characteristic models of a closed set M can be thought of as a minimal “basis” for M , that is, a smallest set that can generate all of M by taking intersections. In general, the characteristic models of any finite M can be defined as those elements of M that do not appear in the closure of the rest of M :

Definition: Characteristic Model

Where M is a finite set of models, the set of characteristic models is given by

$$char(M) \stackrel{def}{=} \{m \in M \mid m \notin closure(M - \{m\})\}$$

For example, the characteristic models of M'_0 are [1110], [1000], and [0101]. The other two models in of M'_0 can be obtained from these characteristic models via intersection.

Note that according to this definition the characteristic elements of any set of models are unique and well-defined. Furthermore, we will prove that characteristic models of a set can generate the complete closure of the set. Now, because the set of models of a Horn theory is closed (Theorem 1), it follows that we can identify a Horn theory with just the characteristic elements among its models. (In fact,

henceforth we will simply say “the characteristic models of a Horn theory” to mean the characteristic subset of its models.) In general, this set may be much smaller than the set of all of its models. In summary:

Theorem 2 *Let M be any finite set of models. Then, (1) the closure of the characteristic models of M is equal to the closure of M ; and (2) if M is the models of a Horn theory, then the closure of the characteristic models of M is equal to M .*

Proof: That $\text{closure}(\text{char}(M)) \subseteq \text{closure}(M)$ is obvious. To prove equality, for a given M and distinct $m_1, m_0 \in M$, define $m_1 >_M m_0$ iff there exists $m_2, \dots, m_n \in M$ such that $m_0 = m_1 \cap m_2 \cap \dots \cap m_n$, while $m_0 \neq m_2 \cap \dots \cap m_n$. Define \geq_M as the reflexive and transitive closure of $>_M$. We make the following three claims: (i) The directed graph corresponding to $>_M$ is acyclic, because if $m_1 >_M m_0$ then the number of variables set to “true” in m_1 is greater than the number set to “true” in m_0 . (ii) The elements of M that are maximal under $>_M$ are characteristic. This is so because if $m \in \text{closure}(M - \{m\})$, there must be some $m_1, \dots, m_n \in M$ such that $m = m_1 \cap \dots \cap m_n$. But then (in particular) $m_1 >_M m$, so m is not maximal under $>_M$. (iii) For any $m \in M$, there is a subset M' of elements of M maximal under $>_M$ such that $m = \bigcap M'$. This set is simply defined by $M' = \{m' \mid m' \geq_M m \text{ and } m' \text{ is maximal under } >_M\}$. In graphical terms, M' consists of the sources of the graph obtained by restricting the graph of $>_M$ to nodes that are $\geq_M m$. Therefore, $M \subseteq \text{closure}(\text{char}(M))$, so $\text{closure}(M) = \text{closure}(\text{char}(M))$. Claim (2) then follows from the previous observations together with Theorem 1. ■

As an aside, one should note that notion of a characteristic model is not the same as the standard definition of a maximal model. By definition, any $m \in M$ is a maximal model of M iff there is no $m' \in M$ such that the variables assigned to “true” by m' are a superset of those assigned to “true” by m . It is easy to see that all maximal models of a set (or theory) are characteristic, but the reverse does not hold. For example, the model [1000] in M_0 is an example of a non-maximal characteristic model.

3 Size of Representations

In this section we will examine the most concise way of representing the information inherent in a Horn theory. We have three candidates: a set of Horn clauses; the complete set of models of the theory; and the set of characteristic models of the theory.

We can quickly eliminate the complete set of models from contention. Obviously, it is at least as large as the set of characteristic models, and often much larger.

Furthermore, every Horn theory with K models over n variables can be represented using at most Kn^2 Horn clauses [Dechter and Pearl, 1992]. Thus up to a small polynomial factor, the complete set of models is also always at least as large as the clausal representation.

Neither of the other two representations strictly dominates the other. We first show that in some cases the representation using characteristic models can be exponentially smaller than the *best* representation that uses Horn clauses.

Theorem 3 *There exist Horn theories with $O(n^2)$ characteristic models where the size of the smallest clausal representation is $O(2^n)$.*

Proof: Consider the theory $\Sigma = \{\neg x_1 \vee \neg x_2 \vee \dots \vee \neg x_n \mid x_i \in \{p_i, q_i\}\}$. The size of Σ is $O(2^n)$. Moreover, one can show that there is no shorter clausal form for Σ , by using a proof very similar to the one in [Kautz and Selman, 1992], but the size of its set of characteristic models is polynomial in n . This can be seen as follows. Write a model as a truth assignment to the variables $p_1q_1p_2q_2 \dots p_nq_n$. From the clauses in Σ , it is clear that in each model there must be some pair p_i and q_i where both letters are assigned false (otherwise, there is always some clause eliminating the model). Without loss of generality, let us consider the set of models with p_1 and q_1 are both assigned false. Each of the clauses in Σ is now satisfied, so we can set the other letters to any arbitrary truth assignment. The characteristic models of this set are

[00111111 ... 11]		[00111111 ... 11]
[00011111 ... 11]	[00111111 ... 01]
[00101111 ... 11]		[00111111 ... 10]

The three models in the first column represent all the settings of the second pair of letters. (Note that 00 can be obtained by intersecting the 2nd and the 3rd model.) Each triple handles the possible settings of one of the pairs. From these $3(n - 1)$ models, we can generate via intersections all possible truth assignments to the letters in all pairs other than the first pair. For each pair, we have a similar set of models with that pair set negatively. And, again each set can be generated using $3(n - 1)$ models. So, the total number of characteristic models is at most $O(n^2)$. ■

The following theorem, however, shows that in other cases, the set of characteristic models can be exponentially *larger* than the best equivalent set of Horn clauses.

Theorem 4 *There exist Horn theories of size $O(n)$ with $O(2^{(n/2)})$ characteristic models.*

Proof: Consider the theory Σ given by the clauses $(\neg a \vee \neg b)$, $(\neg c \vee \neg d)$, $(\neg e \vee \neg f)$, etc. The set M of characteristic models of this theory contains all the models where each of the variables in each consecutive pair, such as (a, b) , (c, d) , (e, f) , etc., are assigned opposite truth values (*i.e.*, either [01] or [10]). So, we get the models [010101 . . .], [100101 . . .], [011001 . . .], . . . , [101010 . . .]. There are $2^{(n/2)}$ of such such models, where n is the number of variables. It is easy to see that these are all maximal models of the theory, and as we observed earlier, all such models are characteristic. (One can go on to argue that there are no other characteristic models in this case.) ■

Thus we see that sometimes the characteristic model set representation offers tremendous space-savings over the clausal representation, and vice-versa. This suggests a strategy if one wishes to compactly represent the information in a closed set of models: interleave the generation of both representations, and stop when the smaller one is completed.

The characteristic models in a closed set can be efficiently found by selecting each model which is not equal to the intersection of any two models in the set. This operation takes $O(K^2n)$ time, where K is the total number of models and n the number of variables. The clausal theory can be found using the algorithms described in [Dechter and Pearl, 1992] and [Kautz *et al.*, to appear].

4 Deduction using Characteristic Models

One of the most appealing features of Horn theories is that they allow for fast inference. In the propositional case, queries can be answered in linear-time [Dowling and Gallier, 1984]. However, there is no *a priori* reason why a representation based on characteristic models would also enable fast inference. Nevertheless, in this section, we show that there is indeed a linear-time algorithm for deduction using characteristic models.

We will take a query to be a formula in conjunctive normal form — that is, a conjunction of clauses. It is easy to determine if a query follows from a complete set of models: you simply verify that the query evaluates to “true” on every model. But if the representation is just the set of characteristic models, such a simple approach does not work. For example, let the query α be the formula $a \vee b$, and let the characteristic set of models be $M_0 = \{[1110], [0101], [1000]\}$, as defined earlier. It is easy to see that α evaluates to true in each member of M_0 . However, α does not logically follow from the Horn theory with characteristic model set M_0 ; in other words, α does not hold in every model in the closure of M_0 . For example, the query is false in $[0101] \cap [1000] = [0000]$.

There is, however, a more sophisticated way of evaluating queries on the set of characteristic models, that does yield an efficient sound and complete algorithm. Our approach is based on the idea of a “Horn-strengthening”, which we introduced in [Selman and Kautz, 1991].

Definition: Horn-strengthening

A Horn clause C_H is a Horn-strengthening of a clause C iff C_H is a Horn clause, C_H subsumes C , and there is no other Horn clause that subsumes C and is subsumed by C_H .

Another way of saying this is that a Horn-strengthening of a clause is generated by striking out positive literals from the clause just until a Horn clause is obtained. For example, consider the clause $C = p \vee q \vee \neg r$. The clauses $p \vee \neg r$ and $q \vee \neg r$ are Horn-strengthenings of C . Any Horn clause has just one Horn-strengthening, namely the clause itself.

Suppose the query is a single clause. Then the following theorem shows how to determine if the query follows from a knowledge base represented by a set of characteristic models.

Theorem 5 *Let Σ be a Horn theory and M its set of characteristic models. Further let C be any clause. Then $\Sigma \models C$ iff there exists some Horn-strengthening C_H of C such that C_H evaluates to “true” in every model in M .*

Proof: Suppose $\Sigma \models C$. By Lemma 1 in [Selman and Kautz, 1991], $\Sigma \models C_H$ for some Horn-strengthening C_H of C . So C evaluates to “true” in every model of Σ , and thus in every member of M . On the other hand, suppose that there exists some Horn-strengthening C_H of C such that C_H evaluates to “true” in every model in M . By Theorem 1, because the elements of M are models of a Horn theory C_H , the elements of the closure of M are all models of C_H . But the closure of M is the models of Σ ; thus $\Sigma \models C_H$. Since $C_H \models C$, we have that $\Sigma \models C$. ■

In the previous example, one can determine that $a \vee b$ does not follow from the theory with characteristic models M_0 because neither the Horn-strengthening a nor the Horn-strengthening b hold in all of $\{[1110], [0101], [1000]\}$.

A clause containing k literals has at most k Horn-strengthenings, so one can determine if it follows from a set of characteristic models in k times the cost of evaluating the clause on each characteristic model. In the more general case the query is a conjunction of clauses. Such a query can be replaced by a sequence of queries, one for each conjunct. We therefore obtain the following theorem:

Theorem 6 *Let a Horn theory Σ be represented by its set of characteristic models M , and let α be a formula in conjunctive normal form. It is possible to determine if $\Sigma \models \alpha$ in time $O(|M| \cdot |\alpha|^2)$, where $|M|$ is the total length of the representation of M .*

Finally, using more sophisticated data structures we can bring the complexity down to truly linear time, $O(|M| + |\alpha|)$ [Kautz *et al.*, 1993].

5 Abduction using Characteristic Models

Another central reasoning task for intelligent systems is abduction, or inference to the best explanation [Peirce, 1958]. In an abduction problem, one tries to *explain* an observation by selecting a set of assumptions that, together with other background knowledge, logically entail the observation. This kind of reasoning is central to many systems that perform diagnosis or interpretation, such as the ATMS.

The notion of an explanation can be formally defined as follows [Reiter and de Kleer, 1987]:

Definition: [Explanation] Given a set of clauses Σ , called the background theory, a subset A of the propositional letters, called the assumption set, and a query letter q , an explanation E for q is a minimal subset of unit clauses with letters from among A such that

1. $\Sigma \cup E \models q$, and
2. $\Sigma \cup E$ is consistent.

Note that an explanation E is a set of unit clauses, or equivalently, a single conjunction of literals.

For example, let the background theory be $\Sigma = \{a, \neg a \vee \neg b \vee \neg c \vee d\}$ and let the assumption set $A = \{a, b, c\}$. The conjunction $b \wedge c$ is an explanation for d .

It is obvious that in general abduction is harder than deduction, because the definition involves both a test for entailment and a test for consistency. However, abduction can remain hard even when the background theory is restricted to languages in which both tests can be performed in polynomial time. Selman and Levesque [1989] show that computing such an explanation is NP-complete even when the background theory contains only Horn clauses, despite the fact that the tests take only linear time for such theories. The problem remains hard because all

```

Explain( $M, A, q$ )
  For each  $m$  in  $M$  do
    If  $m \models q$  then
       $E \leftarrow$  all letters in  $A$  that
        are assigned “true” by  $m$ 
      if  $\text{closure}(M) \models (\bigwedge E) \supset q$  then
        Minimize  $E$  by deleting as many
          elements as possible while
          maintaining the condition
          that  $\text{closure}(M) \models (\bigwedge E) \supset q$ .
        return  $E$ 
      endif
    endif
  endfor
  return “false”
end.

```

Figure 2: Polynomial time algorithm for abduction. M is a set of characteristic models, representing a Horn theory; A is the assumption set; and q is the letter to be explained. The procedure returns a subset of A , or “false”, if no explanation exists.

known algorithms have to search through an exponential number of combinations of assumptions to find an explanation that passes both tests.

There are very few restricted clausal forms for which abduction is tractable. One of these is *definite* Horn clauses, which are Horn clauses that contain exactly one positive literal — completely negative clauses are forbidden. However, the expressive power of definite Horn is much more limited than full Horn: In particular, one cannot assert that two assumptions are mutually *incompatible*.

It is therefore interesting to discover that abduction problems can be solved in polynomial time when the background theory is represented by a set of characteristic models. We give the algorithm for this computation in Figure 2.

The abduction algorithm works by searching for a characteristic model in which the query holds. Then it sets E equal to the strongest set of assumptions that are compatible with the model, and tests if this E rules out all models of the background theory in which the query does not hold. This step is performed by the test

$$\text{closure}(M) \models (\bigwedge E) \supset q$$

and can be performed in polynomial time, using the deduction algorithm described

in the previous section. (Note that the formula to be deduced is a single Horn clause.) If the test succeeds, then the assumption set is minimized, by deleting unnecessary assumptions. Otherwise, if no such characteristic model is in the given set, then no explanation for the query exists. Note that the minimization step simply eliminates redundant assumptions, and does not try to find an assumption set of the smallest possible cardinality, so no combinatorial search is necessary.

It is easy to see that if the algorithm does find an explanation it is sound, because the test above verifies that the query follows from the background theory together with the explanation, and the fact that the model m is in M (and thus also in the closure of M) ensures that the background theory and the explanation are mutually consistent. Furthermore, if the algorithm searched through *all* models in the closure of M , rather than just M itself, it would be readily apparent that the algorithm is complete. (The consistency condition requires that the the explanation and the query both hold in at least one model of the background theory.) However, we will argue that it is in fact only necessary to consider the *maximal* models of the background theory; and since, as we observed earlier, the maximal models are a subset of the characteristic models, the algorithm as given is complete.

So suppose m is in $\text{closure}(M)$, and E is a subset of A such that q and all of E hold in m . Let m' be any maximal model of M (and thus, also a maximal model of $\text{closure}(M)$) that subsumes m — at least one such m' must exist. All the variables set to “true” in m are also set to “true” in m' ; and furthermore, q and all of E consist of only *positive* literals. Therefore, q and E both hold in m' as well.

Thus the algorithm is sound and complete. In order to bound its running time, we note that the outer loop executes at most $|M|$ times, the inner (minimizing) loop at most $|A|$ times, and each entailment test requires at most $O(|M| \cdot |A|^2)$ steps. Thus the overall running time is bounded by $O(|M|^2 \cdot |A|^3)$. In summary:

Theorem 7 *Let M be the set of characteristic models of a background Horn theory, let A be an assumption set, and q be a query. Then one can find an abductive explanation of q in time $O(|M|^2 \cdot |A|^3)$.*

Again, using better data structures, we can reduce the complexity to be quadratic in the combined length of the query and knowledge base.

The fact that abduction is hard for clausal Horn theories, but easy when the same background theory is represented by a set of characteristic models, does not, of course, indicate that $P = NP$! It only means that it may be difficult to generate the characteristic models of a given Horn theory: there may be exponentially many characteristic models, or even if there are few, they may be hard to find. None the less, it may be worthwhile to invest the effort to “compile” a useful Horn theory into

its set of characteristic models, just in case the latter representation does indeed turn out to be of reasonable size. This is an example of “knowledge compilation” [Selman and Kautz, 1991], where one is willing to invest a large amount of off-line effort in order to obtain fast run-time inference. Alternatively, one can circumvent the use of a formula-based representation all together by constructing the characteristic models by hand, or by learning them from empirical data.³

6 Conclusions

In this paper, we have demonstrated that, contrary to prevalent wisdom, knowledge-based systems can efficiently use representations based on sets of models rather than logical formulas. Incomplete information does not necessarily make model-based representations unwieldy, because it possible to store only a subset of characteristic models that are equivalent to the entire model set. We showed that for Horn theories neither the formula nor the model-based representation dominates the other in terms of size, and that sometimes one other can offer an exponential savings over the other.

We also showed that the characteristic model representation of Horn theories has very good computational properties, in that *both* deduction and abduction can be performed in polynomial time. On the other hand, all known and foreseeable algorithms for abduction with Horn clauses are of worst-case exponential complexity.

This paper begins to provide a formal framework for understanding the success and limitations of some of the more empirical work in AI that use model-like representations. Earlier proposals to use models in formal inference, such as Levesque’s proposal for “vivid” representations [Levesque, 1986], rely on using a single, database-like model, and thus have difficulty handling incomplete information. As we have seen, our approach is more general, because we represent a *set* of models. We are currently investigating extensions of the notion of a characteristic model to other useful classes of theories.

References

[Bylander *et al.*, 1989] Bylander, Tom; Allemang, Dean; Tanner, Michael C.; and

³As mentioned earlier, if all models are given, finding the characteristic models takes only polynomial time. However, the complexity of learning the characteristic models where the algorithm can only sample from the complete set of models is an interesting open problem. Some preliminary results on the complexity of this problem have recently been obtained by D. Sloan and R. Schapire (personal communication).

- Josephson, John R. 1989. Some results concerning the computational complexity of abduction. In *Proceedings of KR-89*, Toronto, Ontario, Canada. Morgan Kaufmann. 44.
- [Dechter and Pearl, 1992] Dechter, Rina and Pearl, Judea 1992. Structure identification in relational data. *Artificial Intelligence* 58(1–3):237–270.
- [Dowling and Gallier, 1984] Dowling, William F. and Gallier, Jean H. 1984. Linear time algorithms for testing the satisfiability of propositional Horn formula. *Journal of Logic Programming* 3:267–284.
- [Kautz and Selman, 1992] Kautz, Henry and Selman, Bart 1992. Speeding inference by acquiring new concepts. In *Proceedings of AAAI-92*, San Jose, CA.
- [Kautz *et al.*, 1993] Kautz, Henry; Kearns, Michael; and Selman, Bart 1993. Reasoning with characteristic models. Technical report, AT&T Bell Laboratories, Murray Hill, NJ.
- [Kautz *et al.*, to appear] Kautz, Henry A.; Kearns, Michael J.; and Selman, Bart ppear. Horn approximations of empirical data. *Artificial Intelligence*.
- [Kolodner, 1991] Kolodner, Janet L. 1991. Improving human decision making through case-based decision aiding. *AI Magazine* 12(2):52–68.
- [Levesque, 1986] Levesque, Hector 1986. Making believers out of computers. *Artificial Intelligence* 30(1):81.
- [McCarthy and Hayes, 1969] McCarthy, J. and Hayes, P. J. 1969. Some philosophical problems from the standpoint of artificial intelligence. In Michie, D., editor 1969, *Machine Intelligence 4*. Ellis Horwood, Chichester, England. 463ff.
- [McKinsey, 1943] McKinsey, J. C. C. 1943. The decision problem for some classes of sentences without quantifiers. *Journal of Symbolic Logic* 8(3):61.
- [Peirce, 1958] Peirce, Charles S. 1958. *Collected Papers of Charles Sanders Peirce*. Harvard University Press, Cambridge, MA.
- [Reiter and de Kleer, 1987] Reiter, Raymond and Kleer, Johande 1987. Foundations of assumption based truth maintance systems: Preliminary report. In *Proceedings of AAAI-87*. 183.

[Selman and Kautz, 1991] Selman, Bart and Kautz, Henry 1991. Knowledge compilation using Horn approximations. In *Proceedings of AAAI-91*, Anaheim, CA. 904–909.

[Selman and Levesque, 1990] Selman, Bart and Levesque, Hector J. 1990. Abductive and default reasoning: a computational core. In *Proceedings of AAAI-90*, Boston, MA.

[Selman, 1990] Selman, Bart 1990. Tractable default reasoning. Ph.D. Thesis, Department of Computer Science, University of Toronto, Toronto, Ontario.