Computing Intersections of Horn Theories for Reasoning with Models*

Thomas Eiter

Institut für Informatik
Universität Gießen
Arndtstraße 2, D-35392 Gießen
eiter@informatik.uni-giessen.de

Toshihide Ibaraki

Dept. Applied Mathematics & Physics
Graduate School of Engineering,
Kyoto University, Kyoto 606
ibaraki@kuamp.kyoto-u.ac.jp

Kazuhisa Makino

Dept. Systems & Human Science
Graduate School of Engineering Science,
Osaka University, Toyonaka, Osaka 560
makino@sys.es.osaka-u.ac.jp

Abstract

We consider computational issues when combining logical knowledge bases represented by their characteristic models; in particular, we study taking their logical intersection. We present efficient algorithms or prove intractability for the major computation problems for Horn knowledge bases. We also consider an extension to Horn theories, for which negative results are obtained. They indicate that generalizing the positive results beyond Horn theories is not immediate.

Introduction

More recently, model-based reasoning has been proposed as an alternative to the traditional approach of representing and accessing a logical knowledge base through formulas, cf. (Dechter & Pearl 1992; Kautz, Kearns, & Selman 1993; 1995; Kavvadias, Papadimitriou, & Sideri 1993; Khardon & Roth 1996; 1997). In this approach, a logical knowledge base KB is represented by a subset S of its models, which are commonly called *characteristic models*, rather than by a set of formulas. Reasoning from KB becomes then as easy as to test, given a query α, whether α is true in all models of S; for suitable α, this can be decided efficiently. Note that it has also be shown that abduction from a KB represented by its characteristic models is polynomial (Kautz, Kearns, & Selman 1993; Khardon & Roth 1996), while this problem is intractable under formula representation (Selman & Levesque 1990; Eiter & Gottlob 1995).

This time speed up comes at the price of space; indeed, the formula-based and the model-based approach are orthogonal, in the sense that while a KB may have small representation in the one formalism, it has an exponentially sized larger in the other. The intertranslatability of the two approaches, in particular for Horn theories, has been addressed in (Kautz, Kearns, & Selman 1993; 1995; Kavvadias, Papadimitriou, & Sideri 1993; Khardon 1995; Khardon & Roth 1996).

* We gratefully acknowledge the partial support of the Scientific Grant in Aid by the Ministry of Education, Science and Culture of Japan (Grant 06044112). The major part of this research was conducted while the first author visited Kyoto University in 1995 and 1998.

A number of techniques for efficient model-based representation of various fragments of propositional logic have been devised, cf. (Kautz, Kearns, & Selman 1995; Khardon & Roth 1996; 1997). However, little attention has been paid so far on the important issue of how under this representation different knowledge bases $KB_1, \ldots, KB_n$ can be combined into a single KB.

The semantical issue of combining knowledge bases has been studied in the recent literature, see e.g. (Baral, Kraus, & Minker 1991; Subrahmanian 1994; Liberatore & Schaerf 1998), and we do not pick up this issue here; rather, we are interested in tools for operations at the technical level.

In this context, a principal operation is taking the logical intersection of $KB_1, \ldots, KB_n$, i.e., the resulting knowledge base KB should have the models which are common to all KB_i's. While this operation is easily accomplished under formula-based representation (just take $KB := \bigcup_i KB_i$), this task appears to be much more complicated under model-based representation. In fact, it is a priory not clear, how from the characteristic models of the single KB_i's the characteristic models of KB should be efficiently constructed, and which computational cost is intrinsic to this problem; yet, even an efficient algorithm for simply deciding the consistency of KB is unclear.

In this paper, we address this issue and study the problems of computing characteristic as well as arbitrary models of the logical intersection $\Sigma = \Sigma_1 \cap \cdots \cap \Sigma_n$ of propositional theories Σ_i. We focus on Σ_i's which are Horn theories; such theories are frequently encountered in the context of knowledge representation, and received the major attention in (Dechter & Pearl 1992; Kautz, Kearns, & Selman 1993; 1995; Khardon 1995; Khardon & Roth 1996). In particular, we consider the following main problems.

Problem MODEL

Input: Sets of characteristic models $M_i \subseteq \{0,1\}^n$, representing Horn theories Σ_i, $i = 1, 2, \ldots, l$.

Output: Model v in $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$ if $\Sigma \neq \emptyset$; else, "No".

Problem CMODEL

Input: Sets of characteristic models $M_i \subseteq \{0,1\}^n$, representing Horn theories Σ_i, $i = 1, 2, \ldots, l$.

Output: A characteristic model v in $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$ if $\Sigma \neq \emptyset$; otherwise, "No".

Problem ALL-MODELS

Input: Sets of characteristic models $M_i \subseteq \{0,1\}^n$, representing Horn theories Σ_i, $i = 1, 2, \ldots, l$.

Output: All models v in $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$.

Problem ALL-CMODELS

Input: Sets of characteristic models $M_i \subseteq \{0,1\}^n$, representing Horn theories Σ_i, $i = 1, 2, \ldots, l$.

Output: All characteristic models v in $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$.

Notice that problem MODEL contains the consistency problem of Σ as a subproblem; if we have an efficient algorithm for MODEL, then we can use it for an efficient check whether Σ is consistent, i.e., $\Sigma \neq \emptyset$. Such a test is another principal operation. Note that by the results of (Dowling & Gallier 1984), problem MODEL and the consistency check can be done in linear time under formula representation.

Obviously, problem MODEL is not harder than problem CMODEL, since any procedure for the latter can be used for the former. However, it remains to see whether the computation of an arbitrary model can be done more efficiently.

Problem ALL-MODELS generalizes the first problem, and is of interest for the issue of producing all models of Σ; ideally, this is done one at a time, so that we can stop any time when no further models are desired. Such a procedure is valuable, if we want to select some model in Σ after seeing part of it; e.g., it can be helpful in case-based reasoning, if one tries to find a "model" of the reality which fits a given description.

Problem ALL-CMODELS is the counterpart of the previous problem for characteristic models; here, we are interested in the complete output, as it is the requested representation of Σ in terms of its characteristic models.

From the results in (Kautz, Kearns, & Selman 1993), it easily follows that the output size of problem ALL-MODELS may be exponential in the input size, even if $l = 1$. However, it was unknown whether a similar result holds for ALL-CMODELS. In Section , we show this by example where the output of ALL-CMODELS has 2^n models, while $l = 2$ and $|M_1| = |M_2| = 2n$.

Since ALL-MODELS and ALL-CMODELS may have exponential output, they are clearly not solvable in polynomial time. Observe that our latter result improves on (Gogic, Papadimitriou, & Sideri 1998, Theorem 6), which states that for $l = 2$, ALL-CMODELS is not polynomial unless P = NP.

However, this does not rule out the possibility of an algorithm which enumerates the models with *polynomial delay* (Johnson, Yannakakis, & Papadimitriou 1988), i.e., the next model is always output in time polynomial in the input size, and the algorithm stops in polynomial time after the last output. Any such algorithm runs in *polynomial total time* (Johnson, Yannakakis, & Papadimitriou 1988), i.e., polynomial in the *combined* size of input and output; if no polynomial total time algorithm exists, then a problem is considered to be hard. As ALL-MODELS outputs more models than ALL-CMODELS, there appear more chances of having a polynomial total time algorithm for ALL-MODELS.

Detailed Proofs of all results are given in the full paper, which contains more results (Eiter, Makino, & Ibaraki 1998).

Preliminaries

We assume a standard propositional language with atoms $x_1, x_2, \ldots, x_n$, where each x_i takes either value 1 (true) or 0 (false). Negated atoms are denoted by $\overline{x}_i$.

A *model* v is a vector in $\{0,1\}^n$, whose i-th component is denoted by v_i. For models v, w, we denote by $v \leq w$ the usual componentwise ordering, i.e., $v_i \leq w_i$ for all $i = 1, 2, \ldots, n$, where $0 \leq 1$; $v < w$ means $v \neq w$ and $v \leq w$. As usual, $v \geq w$ is the reverse ordering. For $B \subseteq \{1, \ldots, n\}$, we denote by x^B the model v such that $v_i = 1$, if $i \in B$ and $v_i = 0$, if $i \notin B$, for all $i = 1, \ldots, n$.

A *theory* is any set $\Sigma \subseteq \{0,1\}^n$ of models; its cardinality is denoted by $|\Sigma|$. By $\min(\Sigma)$ and $\max(\Sigma)$ we denote the sets of minimal and maximal models in Σ under $<$, respectively, where $v \in \Sigma$ is *maximal* (resp., *minimal*) model in Σ, if there is no $w \in \Sigma$ such that $w > v$ (resp., $w < v$).

A propositional clause $C = \ell_1 \vee \cdots \vee \ell_k$ is *Horn*, if at most one literal ℓ_i is positive, and a CNF is *Horn*, if it contains only Horn clauses. A theory Σ is *Horn*, if there exists a Horn CNF representing it.

Horn theories Σ have a well-known model-theoretic characterization. Denote by $v \bigwedge w$ componentwise AND of vectors $v, w \in \{0,1\}^n$, and by $Cl_\wedge(S)$ the closure of $S \subseteq \{0,1\}^n$ under $\bigwedge$. Then, Σ is Horn, if and only if $\Sigma = Cl_\wedge(\Sigma)$. Note that as a consequence, any Horn theory Σ has the *least* (i.e., *unique minimal*) *model* $v = \bigwedge_{w \in \Sigma} w$, i.e., $\min(\Sigma) = \{v\}$.

E.g., consider $\Sigma_1 = \{(0101), (1001), (1000)\}$ and $\Sigma_2 = \{(0101), (1001), (1000), (0001), (0000)\}$. Then, for $v = (0101)$, $w = (1000)$, we have $w, v \in \Sigma_1$, while $v \bigwedge w = (0000) \notin \Sigma_1$; hence Σ_1 is not Horn. On the other hand, $Cl_\wedge(\Sigma_2) = \Sigma_2$, thus Σ_2 is Horn.

For any Horn theory Σ, a model $v \in \Sigma$ is called *characteristic* (Kautz, Kearns, & Selman 1993) (or *extreme* (Dechter & Pearl 1992)), if $v \notin Cl_\wedge(\Sigma \setminus \{v\})$. The set of all characteristic models of Σ, the *characteristic set of* Σ, is denoted by $C^*(\Sigma)$. Note that every Horn theory Σ has a unique characteristic set $C^*(\Sigma)$ and that $\max(\Sigma) \subseteq C^*(\Sigma)$. E.g., $(0101) \in C^*(\Sigma_2)$, while $(0000) \notin C^*(\Sigma_2)$; it holds that $C^*(\Sigma_2) = \Sigma_1$.

Finding Some Model

We start with the following lemma, which is useful for solving problem MODEL.

Lemma 1 *Let $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, 2, \ldots, l$, be Horn theories, and let $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$. Then any $v \in \Sigma$ satisfies*

$$v \;\geq\; \bigvee_{i=1}^{l} \Big(\bigwedge_{w \in C^*(\Sigma_i)} w \Big). \qquad (1)$$

Proof. First note that $v = \bigwedge_{w \in Q_1} w = \bigwedge_{w \in Q_2} w = \ldots = \bigwedge_{w \in Q_l} w$ holds for some $Q_i \subseteq C^*(\Sigma_i)$, $i = 1, 2, \ldots, l$, by the definitions of v and $C^*(\Sigma_i)$. Then we have $v \geq \bigwedge_{w \in C^*(\Sigma_i)} w$ for all i, and hence (1). □

Based on the lemma, a model of Σ is found as follows.

Clearly, Σ has no model, if some Σ_i is empty; if not, then consider the least models $v_1, \ldots, v_l$ of $\Sigma_1, \ldots, \Sigma_l$. If they all coincide, then $v = v_1$ is a model of Σ, which is the output. Otherwise, exploiting Lemma 1, we look at the least upper bound of $v_1, \ldots, v_l$ as a new candidate u for a model; in fact, any $v \in \Sigma$ must satisfy $u \leq v$. Since v must be generated from characteristic models in each Σ_i, we can discard all characteristic models which for sure do not contribute in that. Since the resulting theories are Horn, we can iterate and build a chain $C : u^{(1)} < u^{(2)} < \cdots < u^{(k)}$ such that either $u^{(k)}$ is found to be a model of Σ, or $\Sigma = \emptyset$ is detected.

The formal description of this algorithm is as follows.

Algorithm MODEL

Input: Characteristic sets $M_i = C^*(\Sigma_i)$, representing Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, \ldots, l$.

Output: Model $v \in \Sigma = \bigcap_{i=1}^{l} \Sigma_i$ if $\Sigma \neq \emptyset$; else, "No".

Step 0. **for** each $i = 1, 2, \ldots, l$ **do** $Q_i := M_i$;

Step 1. **if** $Q_i = \emptyset$ for some i **then** output "No" and halt;

Step 2. **if** $\bigwedge_{w \in Q_1} w = \bigwedge_{w \in Q_2} w = \ldots = \bigwedge_{w \in Q_l} w$
then output $v = \bigwedge_{w \in Q_1} w$ and halt;

Step 3. $u := \bigvee_{i=1}^{l} (\bigwedge_{w \in Q_i} w)$;
for each $i = 1, \ldots, l$ **do** $Q_i := \{w \in Q_i \mid w \geq u\}$;
goto Step 1. □

Example 1 Let $M_1 = C^*(\Sigma_1) = \{(0110), (0011), (1010)\}$ and $M_2 = C^*(\Sigma_2) = \{(1110), (0111), (0011)\}$.

In step 2, we have $\bigwedge_{w \in Q_1} = (0010)$ and $\bigwedge_{w \in Q_2} = (0010)$; hence, $v = (0010)$ is output. Note that $\Sigma = \{(0110), (0010), (0011)\}$; thus, output of $v = (0010)$ is correct. □

An analysis of its run time gives the following result.

Theorem 1 *Problem* MODEL *can be solved in $O(n^2 \sum_{i=1}^{l} |M_i|)$ time.* □

In fact, algorithm MODEL finds a distinguished model of Σ; it is not hard to see from its working that the output is the *least model* of Σ. Thus,

Corollary 1 *Algorithm* MODEL *finds the least model v of $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$ in $O(n^2 \sum_{i=1}^{l} |M_i|)$ time if $\Sigma \neq \emptyset$, and outputs "No" if $\Sigma = \emptyset$.* □

Algorithm MODEL has run time about size of the input times the number of propositional atoms, and is thus almost quadratic in the worst case.

In the full paper, we describe an improved version MODEL+ which runs in $O(n \sum_{i=1}^{l} |M_i|)$ time, i.e., in linear time. This is achieved by using appropriate data structures, including cross-reference lists and counters which help in avoiding that the same bit of the input is examined more than a constant number of times. We only note the result.

Proposition 1 *Given the characteristic sets $C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, 2, \ldots, l$, deciding consistency and computing the least model of $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$ is possible in $O(n \sum_{i=1}^{l} |M_i|)$ time, i.e., in linear time.* □

Finding Some Characteristic Model

Also problem CMODEL can be solved in polynomial time.

Basically, we can proceed as follows. We construct the least model u of $\Sigma = \bigcap_i \Sigma_i$ as a candidate in $C^*(\Sigma)$; this is possible using algorithm MODEL. Then, two cases arise:

(i) $u \in C^*(\Sigma)$; in this case, we can output u and stop.

(ii) $u \notin C^*(\Sigma)$; here, u is replaced by a new larger candidate model $u' > u$, $u' \in \Sigma$, and the process is continued.

Since any chain of models $u = u^{(1)} < u^{(2)} < \ldots < u^{(k)}$ is bounded, the algorithm eventually finds some characteristic model (as any maximal model is characteristic) and halts. The problem is recognizing which case applies, and to select in (ii) a proper u'. It can be seen that $u \in C^*(\Sigma)$ holds, if (but not only if) the following condition holds. Let $Q_i = \{v > u \mid v \in M_i\}$ and $P_{ij} = \{w \in Q_i \mid w_j = 1\}$.

$$\forall j : u_j = 0 \implies \bigcap_{i=1}^{l} Cl_\wedge(P_{ij}) = \emptyset \qquad (2)$$

On the other hand, if for some j, (2) is violated, then any model $v \in Cl_\wedge(P_{ij})$ is a model of Σ with $v > u$; since some characteristic model $\geq v$ exists, we can safely select $u' = v$ and replace each M_i by the set $\{w \geq u' \mid w \in P_{ij}\}$.

Example 2 Let again $M_1 = C^*(\Sigma_1) = \{(0110), (0011), (1010)\}$ and $M_2 = C^*(\Sigma_2) = \{(1110), (0111), (0011)\}$.

The least model of $\Sigma = \Sigma_1 \cap \Sigma_2$ is $u = u^{(1)} = (0010)$. Thus, we have $Q_1^{(1)} = M_1$ and $Q_2^{(1)} = M_2$. For $j = 2$, we have $P_{12}^{(1)} = \{(0110)\}$ and $P_{22}^{(1)} = \{(1110), (0111)\}$; hence, $(0110) \in Cl_\wedge(P_{12}^{(1)}) \cap Cl_\wedge(P_{22}^{(1)})$ violates (2). Thus, we set $u^{(2)} = (0110)$ and continue; we set $M_1^{(2)} := \{(0110)$ and $M_2^{(2)} := \{(1110), (0111)\}$. Then, we obtain $Q_1^{(2)} = \emptyset$ and $Q_2^{(2)} = \{(1110), (0111)\}$. Consequently, for each j, $P_{1j}^{(2)}$ is empty, which means that condition (2) is true; hence, $v = u^{(2)}$ is output. Note that $C^*(\Sigma) = \{(0110), (0011)\}$; thus, output of $v = (0110)$ is correct. □

Improving the above method, we can save on time by exploiting the observation that if some j with $u_j = 0$ satisfies (2), then we need not check if (2) holds for this j later again. Indeed, this means that there exists no $w' \in \Sigma$ such that $w' \geq u$ and $w'_j = 1$.

Formally, our algorithm can be written as follows.

Algorithm CMODEL

Input: Characteristic sets $M_i = C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, 2, \dots, l$.

Output: A model $v \in C^*(\Sigma)$, where $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$, if $\Sigma \neq \emptyset$; otherwise, "No".

Step 1. find the least model u in Σ;
  if no such u exists **then** output "No"
  else for each $i = 1, 2, \dots, l$ **do**
    $Q_i := \{w \in M_i \mid w \geq u\}$;

Step 2. **for** each $j = 1, 2, \dots, n$ **do**
  if $u_j = 0$ **then begin**
    for each $i = 1, 2, \dots, l$ **do**
      $P_{ij} := \{w \in Q_i \mid w_j = 1\}$;
    if $\bigcap_{i=1}^{l} Cl_\wedge(P_{ij}) \neq \emptyset$ **then begin**
      find a model w' in $\bigcap_{i=1}^{l} Cl_\wedge(P_{ij})$;
      $u := w'$;
      for each $i = 1, 2, \dots, l$ **do**
        $Q_i := \{w \in P_{ij} \mid w \geq u\}$;
    end;
  end;

Step 3. output the model $v := u$. □

An analysis of the running time of this algorithm yields the following result.

Theorem 2 *Problem* CMODEL *can be solved in* $O(n^2 \sum_{i=1}^{l} |M_i|)$ *time.* □

From the working of this algorithm, we see that it outputs some particular characteristic model, namely a maximal model of Σ. We thus obtain the following result.

Corollary 2 *Algorithm* CMODEL *finds a maximal model* v *in* $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$ *in* $O(n^2 \sum_{i=1}^{l} |M_i|)$ *time if* $\Sigma \neq \emptyset$*, and outputs "No" if* $\Sigma = \emptyset$. □

The fact that we can *compute* some characteristic model fast does not automatically mean that we can *recognize* any characteristic model fast; nonetheless, this task can be solved in polynomial time. The key for this result is the following lemma.

Lemma 2 *Let* Σ *be a Horn theory and* v *be a model in* Σ. *Then* $v \notin C^*(\Sigma)$ *holds if and only if* $v = \bigwedge_{w \in \min(\Sigma_v)} w$, *where* $\Sigma_v = \{w \in \Sigma \mid w > v\}$. □

Exploiting this Lemma, we construct the following algorithm for characteristic model checking.

Algorithm CHECK-CMODEL

Input: Characteristic sets $M_i = C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, \dots, l$, and a model $v \in \Sigma = \bigcap_{i=1}^{l} \Sigma_i$.

Output: "Yes", if $v \in C^*(\Sigma)$, otherwise, "No".

Step 0. $S := \emptyset$.

Step 1. **for** each j with $v_j = 0$ **do begin**
  for each $i = 1, 2, \dots, l$ **do**
    $Q_i^{(j)} := \{w \in M_i \mid w \geq v, w_j = 1\}$;
  if $\bigcap_{i=1}^{l} Cl_\wedge(Q_i^{(j)}) \neq \emptyset$ **then begin**
    $w^{(j)}$:= least model in $\bigcap_{i=1}^{l} Cl_\wedge(Q_i^{(j)})$;
    $S := S \cup \{w^{(j)}\}$;
  end;
end;

Step 2. **if** $v = \bigwedge_{w^{(j)} \in S} w^{(j)}$ **then** output "No"
  else output "Yes". □

Example 3 Let as above $M_1 = C^*(\Sigma_1) = \{(0110), (0011), (1010)\}$ and $M_2 = C^*(\Sigma_2) = \{(1110), (0111), (0011)\}$, and suppose $v = (0110)$.

Then, in step 1 of CHECK-CMODEL, $S := \emptyset$; in step 2, j takes values 1 and 4. For $j = 1$, we obtain $Q_1^{(1)} := \emptyset$ and $Q_2^{(1)} := \{(1110)\}$, hence $Cl_\wedge(Q_1^{(1)}) \cap Cl_\wedge(Q_2^{(1)}) = \emptyset$, and S is unchanged. For $j = 4$, we have $Q_1^{(4)} = \emptyset$ again and $Q_2^{(4)} = \{(0111)\}$; hence $S = \emptyset$ is not changed. In step 3, the check $v = \bigwedge_{w^{(j)} \in S} w^{(j)}$ yields false (for empty S, $\bigwedge_{w \in S} w = (11...1)$); hence the output is "Yes". Note that $v = (0110)$ is indeed a characteristic model of Σ. □

An analysis of the running time of the algorithm yields the following result.

Theorem 3 *Given the characteristic sets* $C^*(\Sigma_i)$ *of Horn theories* $\Sigma_i \subseteq \{0,1\}^n$, $i = 1 \dots, l$, *and a model* $v \in \Sigma = \bigcap_{i=1}^{l} \Sigma_i$, *checking if* $v \in C^*(\Sigma)$ *is possible in* $O(n^2 \sum_{i=1}^{l} |M_i|)$ *time.* □

ALL-CMODELS and ALL-MODELS

It is known that for a Horn theory Σ, the number $|\Sigma|$ of its models may be exponential in $|C^*(\Sigma)|$. Thus the output size of ALL-MODELS may be exponential in the input size. For ALL-CMODELS, we derive an analogous result.

Claim 1 *For every* $n \geq 1$*, there exist Horn theories* Σ_1 *and* Σ_2 *such that* $|C^*(\Sigma_1)| = |C^*(\Sigma_2)| = 2n$ *and* $|C^*(\Sigma)| = 2^n$*, where* $\Sigma = \Sigma_1 \cap \Sigma_2$.

Proof. (Sketch) Fix n, and define $S_1, S_2, \subseteq \{0,1\}^{4n}$ as follows. Let $V_i = \{i * n + j \mid j = 1, \dots n\}$, for $i = 0, \dots, 3$ and $V = \bigcup_{i=0}^{3} V_i$ = { 1,...,4n}. Then,

$$S_1 = \{x^{V \setminus (V_2 \cup \{j, 3n+j\})}, x^{V \setminus (V_2 \cup \{n+j, 3n+j\})} \mid 1 \leq j \leq n\},$$
$$S_2 = \{x^{V \setminus (V_3 \cup \{j, 2n+j\})}, x^{V \setminus (V_3 \cup \{n+j, 2n+j\})} \mid 1 \leq j \leq n\}.$$

E.g., for $n = 2$, we have

$$S_1 = \{(01110001), (11010001), (10110010), (11100010)\},$$
$$S_2 = \{(01110100), (11010100), (10111000), (11101000)\}.$$

Observe that $|S_1| = |S_2| = 2n$. Since $S_1 = \max(S_1)$ and $S_2 = \max(S_2)$, there are Horn theories Σ_1 and Σ_2 such that $C^*(\Sigma_1) = S_1$ and $C^*(\Sigma_2) = S_2$. Define

$$S = \{x^B \mid B \subseteq V_0 \cup V_1, j \in B \equiv n + j \notin B, 1 \leq j \leq n\}.$$

For $n = 2$, we have

$S = \{(00110000), (10010000), (01100000), (11000000)$. Observe that $|S| = 2^n$. It can be shown that $S = C^*(\Sigma)$. Since $|S| = 2^n$, the claim is verified. □

Hence, a polynomial time algorithm in the input size for ALL-CMODELS is impossible, which improves (Gogic, Papadimitriou, & Sideri 1998, Theorem 6).

However, even the remaining hope for a polynomial total time algorithm is unlikely to come true, since the following related problem is intractable.

Lemma 3 *The problem* ADD-CMODEL*: Given characteristic sets $C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, \ldots, l$ and $S \subseteq C^*(\Sigma)$, where $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$, decide if some $v \in C^*(\Sigma) \setminus S$ exists; is* NP*-hard, even for $l = 2$.*

Proof. (Sketch) We prove NP-hardness by a reduction from the satisfiability problem (SAT) (Garey & Johnson 1979); we define for a given CNF formula $\Phi = \bigwedge_{i=1}^{m} C_i$ on n atoms polynomially computable sets M_1, M_2, and S of vectors in $\{0,1\}^{n+2m}$, such that $M_1 = C^*(\Sigma_1)$, $M_2 = C^*(\Sigma_2)$ and $S \subseteq C^*(\Sigma_1 \cap \Sigma_2)$. Moreover, $S = C^*(\Sigma_1 \cap \Sigma_2)$ holds iff Φ is unsatisfiable. □

Theorem 4 *There is no polynomial total time algorithm for problem* ALL-CMODELS*, unless* P=NP.

Proof. Assume there is an algorithm A for ALL-CMODELS with polynomial running time $p(I, O)$, where I is the input length and O the output length. We solve ADD-CMODEL using A: Execute A until either (i) it halts or (ii) time $p(I, |S|)$ is reached. In case (i), output "Yes" if A outputs some vector in $C^*(\Sigma) \setminus S$; otherwise, "No". In case (ii), output "Yes", since it implies $C^*(\Sigma) \setminus S \neq \emptyset$. Hence, ADD-CMODEL is solvable in time polynomial in I and $|S|$, which contradicts Lemma 3 unless P=NP. □

In the full paper, we also show that *approximating* $C^*(\Sigma)$ is hard; unless P = NP, there are no polynomial total time algorithms for computing a polynomially larger superset or a polynomial fraction of $C^*(\Sigma)$, respectively.

Contrary to ALL-CMODELS, problem ALL-MODELS has a polynomial total time algorithm. As we show that it is possible to check whether a *partial* vector $v \in \{0, 1, ?\}^n$, where '?' represents unknown, can be completed to a model $w \in \Sigma$ in polynomial time, we can apply the method of dynamic lexicographic ordering (Dechter & Itai 1992) to enumerate all models with polynomial delay. The algorithm uses a bookkeeping vector $mark \in \{0,1\}^n$ and a subroutine PART-MODEL, which has the following specification:

Procedure PART-MODEL

Input: Characteristic sets M_i of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, 2, \ldots, l$, and a list $b_1, b_2, \ldots, b_r$ of values $b_i \in \{0,1\}$, $1 \leq i \leq r \leq n$.

Output: Model $w \in \Sigma = \bigcap_{i=1}^{l} \Sigma_i$ such that $w_i = b_i$ holds for all $i = 1, 2 \ldots, r$, if one exists; "No," otherwise. □

This algorithm can be implemented to run in $O(n|\sum_{i=1}^{l} |M_i|)$ time. The main algorithm is then as follows.

Algorithm ALL-MODELS

Input: Characteristic sets $M_i = C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, 2, \ldots, l$.

Output: All models $v \in \Sigma = \bigcap_{i=1}^{l} \Sigma_i$, if $\Sigma \neq \emptyset$; otherwise, "No".

Step 1. call MODEL to find some model $v \in \Sigma$;
    if the answer is "No", **then** output "No" and halt
    else begin output v;
        $mark := (00\ldots0)$; $i := n$
    end;

Step 2. **if** $mark_i = 0$ **then begin**
        call PART-MODEL($\Sigma_1, .., \Sigma_l, v_1, .., v_{i-1}, 1-v_i$);
        if a model w is returned **then**
        begin output w;
            set $v := w$; $mark_i := 1$;
            for $j = i+1$ **to** n **do** $mark_j := 0$;
            $i := n+1$
        end
    end;

Step 3. **if** $i = 1$ **then** halt
    else begin $i := i-1$; **goto** Step 2 **end**. □

Example 4 Let again $M_1 = C^*(\Sigma_1) = \{(0110), (0011), (1010)\}$ and $M_2 = C^*(\Sigma_2) = \{(1110), (0111), (0011)\}$.

In Step 1, the call to MODEL returns the least model of Σ, which is $v = (0010)$; this model is output and $mark$ is initialized to (0000) and $i := 4$.

In Step 2, PART-MODEL is called for the list $0, 0, 1, 1$ of b_i values (we omit $\Sigma_1, \ldots, \Sigma_l$, which may be accessed as global variables). The model (0011) is returned, which is output and assigned to v; $mark$ is updated to (0001) and i is set to 5 and decreased to 4 in Step 3, where the computation returns to Step 2.

In Step 3, i is decreased to 3, and in next iteration of Step 2, PART-MODEL is called for the b_i values 0,0,0. The answer is "No", and hence i is decreased to 2 in Step 3. Subsequently, in Step 2 PART-MODEL is called for the b_i values 0,1. The model $w = (0110)$ is returned, which is output; $v := (0110)$, $mark := (0100)$, and $i := 5$.

In the next 2 iterations, PART-MODEL is called for b_i values 0,1,1,1 and 0,1,0, respectively, for which "No" is returned; after decreasing i to 1, PART-MODEL is called again for B_i value 1, which also returns "No". Hence, in Step 3 $i = 1$ is true, and the algorithm stops.

Thus, the models output are: (0010), (0110), and (0011); these are precisely the models in Σ. □

The analysis of the time complexity of ALL-MODELS, gives us the next result.

Theorem 5 *Algorithm* ALL-MODELS *is a polynomial delay algorithm for* ALL-MODELS*, where the delay is $O(n^2 \sum_{i=1}^{l} |M_i|)$.* □

Further Results and Conclusion

As shown in (Kautz, Kearns, & Selman 1993), a striking adavantage of characteristic models is that both deduction $\Sigma \models \alpha$ of a CNF formula α and abduction of a query letter q from the characteristic models of a Horn theory is possible in polynomial time. In the full paper, we show the following results:

Theorem 6 • *Deductive Inference $\Sigma \models \alpha$ of a CNF formula α from $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$, given α and the characteristic sets $M_i = C^*(\Sigma_i)$ of the Horn theories $\Sigma_i \subseteq \{0,1\}^n$ for input, is feasible in $O(nm \sum_{i=1}^{l} M_i|)$ time, where m is the number of clauses in α;*

• *abduction of a letter q from Σ, given in the same way, and a given set A of assumptions is* NP*-complete.* □

The results show that deduction scales up gracefully from a single Horn KB to the intersection of multiple Horn KBs, while this is not the case for abduction; it indicates that the tractability result in (Kautz, Kearns, & Selman 1993) is not very robust.

Characteristic models have been generalized to Non-Horn theories by making use of *monotone theory* (Bshouty 1995) in (Khardon & Roth 1996). This approach is promising, since many advantages of Horn theories carry over to Non-Horn theories. In this direction, we investigate in the full paper the class $\mathcal{C}_{EH}$ of *extended Horn* theories, which contains Horn and *reverse* Horn theories, i.e., theories which become Horn by negating all propositions. For this class, we establish the following result.

Theorem 7 *Problem* MODEL *for class $\mathcal{C}_{EH}$ is NP-hard, even if $l = 2$.* □

We prove this result by a reduction from the EXACT-HITTING-SET problem (Garey & Johnson 1979).

Corollary 3 *For class $\mathcal{C}_{EH}$, problem* CMODEL *is* NP*-hard, and there exist no polynomial total time algorithms for* ALL-MODELS *and* ALL-CMODELS*, unless* P=NP. □

Moreover, also deduction and abduction for $\mathcal{C}_{EH}$ are intractable. This indicates that a generalization of characteristic models is not immediately computationally feasible for combining knowledge bases. Moreover, an investigation of relevant classes besides Horn theories which are benign for combination remains to be done.

Further operations in combining theories Σ_i may be needed; e.g., taking the union $\Sigma' = \bigcup_i \Sigma_i$. Notice that Σ' is not necessarily Horn, if all Σ_i are Horn; in that case, Σ' may be approximated by Horn theories (Kautz, Kearns, & Selman 1995; Kavvadias, Papadimitriou, & Sideri 1993).

Our further and future work addresses these and other issues; e.g., we investigate into conditions for Horn Σ' which may serve as a basis for suitable algorithms.

Acknowledgements. We thank the referees for their helpful comments on this paper.

References

Baral, C.; Kraus, S.; and Minker, J. 1991. Combining Multiple Knowledge Bases. *IEEE Transactions on Knowledge and Data Engineering* 3(2):208–220.

Bshouty, N. H. 1995. Exact Learning Boolean Functions via the Monotone Theory. *Information and Computation* 123:146–153.

Dechter, R., and Itai, A. 1992. Finding All Solutions if You can Find One. Technical Report ICS-TR-92-61, University of California at Riverside.

Dechter, R., and Pearl, J. 1992. Structure Identification in Relational Data. *Art. Intelligence* 58:237–270.

Dowling, W., and Gallier, J. H. 1984. Linear-time Algorithms for Testing the Satisfiability of Propositional Horn Theories. *J. Logic Programming* 3:267–284.

Eiter, T., and Gottlob, G. 1995. The Complexity of Logic-Based Abduction. *JACM* 42(1):3–42.

Eiter, T.; Makino, K.; and Ibaraki, T. 1998. Computing Intersections of Horn Theories for Reasoning with Models. TR 9803, Institut für Informatik, Universität Gießen, Germany.

Garey, M., and Johnson, D. S. 1979. *Computers and Intractability*. New York: W. H. Freeman.

Gogic, G.; Papadimitriou, C.; and Sideri, M. 1998. Incremental Recompilation of Knowledge. *J. Artificial Intelligence Research* 8:23–37.

Johnson, D. S.; Yannakakis, M.; and Papadimitriou, C. H. 1988. On Generating All Maximal Independent Sets. *Information Processing Letters* 27:119–123.

Kautz, H.; Kearns, M.; and Selman, B. 1993. Reasoning With Characteristic Models. In *Proc. AAAI-93*.

Kautz, H.; Kearns, M.; and Selman, B. 1995. Horn Approximations of Empirical Data. *Artificial Intelligence* 74:129–245.

Kavvadias, D.; Papadimitriou, C.; and Sideri, M. 1993. On Horn Envelopes and Hypergraph Transversals. In *Proc. ISAAC-93*, LNCS 762, 399–405.

Khardon, R., and Roth, D. 1996. Reasoning with Models. *Artificial Intelligence* 87(1/2):187–213.

Khardon, R., and Roth, D. 1997. Defaults and Relevance in Model-Based Reasoning. *Artificial Intelligence* 97(1/2):169–193.

Khardon, R. 1995. Translating between Horn Representations and their Characteristic Models. *J. Artificial Intelligence Research* 3:349–372.

Liberatore, P., and Schaerf, M. 1998. Arbitration (or How to Merge Knowledge Bases). *IEEE Transactions on Knowledge and Data Engineering* 10(1), 1998.

Selman, B., and Levesque, H. J. 1990. Abductive and Default Reasoning: A Computational Core. In *Proc. AAAI-90*, 343–348.

Subrahmanian, V. 1994. Amalgamating Knowledge Bases. *ACM Trans. on Database Syst.* 19(2):291–331.