I F I G

RESEARCH

REPORT

INSTITUT FÜR INFORMATIK

COMPUTING INTERSECTIONS OF HORN THEORIES FOR REASONING WITH MODELS

Thomas Eiter Toshihide Ibaraki Kazuhisa Makino

IFIG RESEARCH REPORT 9803
APRIL 1998

Institut für Informatik
JLU Gießen
Arndtstraße 2
D-35392 Giessen, Germany
Tel: +49-641-99-32141
Fax: +49-641-99-32149
mail@informatik.uni-giessen.de
www.informatik.uni-giessen.de

JUSTUS-LIEBIG-
UNIVERSITÄT
GIESSEN

IFIG RESEARCH REPORT

IFIG RESEARCH REPORT 9803, APRIL 1998

COMPUTING INTERSECTIONS OF HORN THEORIES FOR REASONING WITH MODELS

Thomas Eiter[1] Toshihide Ibaraki[2] Kazuhisa Makino[3]

Abstract. Model-based reasoning has been proposed as an alternative form of representing and accessing logical knowledge bases. In this approach, a knowledge base is represented by a set of characteristic models. In this paper, we consider computational issues when combining logical knowledge bases, which are represented by their characteristic models; in particular, we study taking their logical intersection. We present low-order polynomial time algorithms or prove intractability for the major computation problems in the context of knowledge bases which are Horn theories. In particular, we show that a model of the intersection Σ of Horn theories $\Sigma_1, \ldots, \Sigma_l$, represented by their characteristic models, can be found in linear time, and that some characteristic model of Σ can be found in polynomial time. Moreover, we present an algorithm which enumerates the models of Σ with polynomial delay. The analogous problem for the characteristic models is proved to be intractable, even if the possible exponential size of the output is taken into account. Furthermore, we show that approximate computation of the set of characteristic models is difficult as well. Nonetheless, we show that deduction from Σ is possible for a large class of queries in polynomial time, while abduction turns out to be intractable. We also consider an extension of Horn theories, and prove negative results for the basic questions, indicating that an extension of the positive results beyond Horn theories is not immediate.

[1]Institut für Informatik, Universität Gießen, Arndtstraße 2, D-35392 Gießen, Germany. Email: eiter@informatik.uni-giessen.de

[2]Department of Applied Mathematics and Physics, Graduate School of Informatics, Kyoto University, Kyoto 606, Japan. Email: ibaraki@kuamp.kyoto-u.ac.jp

[3]Department of Systems and Human Science, Graduate School of Engineering Science, Osaka University, Toyonaka, Osaka 560, Japan. Email: makino@sys.es.osaka-u.ac.jp

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

An abstract of this paper containing some of the results will appear in: *Proceedings of the National Conference on Artificial Intelligence (AAAI '98)*, Madison, Wisconsin, July 26–30, 1998.

1 Introduction

Logical languages are widely used as a basis for representing knowledge in advanced knowledge based systems (cf. [17]). The investigation of adequate languages, at the syntactical as well as the semantical level, is an ongoing quest for improving on the capabilities of current systems. In this approach, knowledge has been traditionally represented by means of logical formulas, which are stored in a knowledge base KB; intuitively, such a KB is meant to capture the knowledge about a certain domain and state of affairs, which is often called the "world". The knowledge may be accessed by posing queries to KB, which are typically expressed by logical formulas α. The query α is then answered by deduction or some other inference method from KB; i.e., it is tested whether KB entails the query α ($KB \vdash \alpha$). One of the main disadvantages of this approach is that deciding whether $KB \vdash \alpha$ holds is intractable in already plain settings; e.g., in the propositional context, it is a well-known co-NP-complete problem.

More recently, model-based reasoning has been proposed as an alternative form of representing and accessing a logical knowledge base, cf. [13, 24, 25, 26, 29, 30]. It can be seen as an approach towards Levesque's notion of "vivid" reasoning [31], which for asks for a more straight representations of a knowledge base, from which common-sense reasoning is easier and more suitable than from the traditional one. In model-based reasoning, 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. Moreover, it has also been shown that abduction from a KB represented by its characteristic models is polynomial [24, 25, 29], while this problem is intractable under formula representation [38, 15].

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 one formalism, it has an exponentially larger representation in the other. The intertranslatability of the two approaches, in particular for Horn theories, has been addressed in [24, 25, 26, 27, 29]. A number of techniques for efficient model-based representation of various fragments of propositional logic have been devised, cf. [25, 29, 30]. However, little attention has been paid so far on the important issue of how in this representation different knowledge bases $KB_1, \ldots, KB_n$ can be combined into a single KB.

Main problems studied. The semantical issue of combining knowledge bases, as well as closely related issues, have been studied in the recent literature, see e.g. [2, 1, 18, 41, 23, 36, 39, 7, 33]. We do not intend to discuss the same issue here; rather, we are interested in tools and algorithms for operations at the technical level, which are needed for the implementation of a suitable semantics. In this context, a principal operation is taking the logical intersection of knowledge bases $KB_1, \ldots, KB_n$, 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 priori not clear, how from the characteristic models of the individual KB_i's the characteristic models of KB can be efficiently constructed, and what computational cost is intrinsic to this problem. For example, 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 $\Sigma = \Sigma_1 \cap \cdots \cap \Sigma_l$ of propositional theories Σ_i. We focus on those Σ_i's which are Horn theories; such theories are frequently encountered in the context of knowledge representation, and their study in model-based reasoning received the main attention in [13, 24, 25, 26, 27, 20], and was further discussed in [29]. In particular, we consider the following main problems in the context of model computation. Given the sets of characteristic models $M_1, \ldots, M_l$ representing Horn theories $\Sigma_1, \ldots, \Sigma_l$,

- compute some arbitrary model of the theory $\Sigma = \bigcap_{i=1}^{l}$ (problem MODEL);
- compute some arbitrary characteristic model of Σ (problem CMODEL);
- compute all models of the theory Σ (problem ALL-MODELS); and,
- compute all characteristic models of Σ (problem ALL-CMODELS).

Further problems on models, such as model checking [8, 32], i.e., the recognition of models in Σ and characteristic models, will be considered as well.

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., whether $\Sigma \neq \emptyset$ holds. Note that by the results of [14] (see also [21]), problem MODEL and the consistency check can be solved 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 solving the former problem. However, it remains to see whether the computation of an arbitrary model can be done more efficiently than a characteristic model.

Problem ALL-MODELS generalizes the first problem. Ideally, the generation of models is done one at a time, so that we can stop any time when no further models are desired. Such a procedure is valuable in case-based reasoning, for example, if one tries to find a "model" of the reality which fits a given description, or provides a good approximation for it. More general, such an enumeration procedure can serve as a general purpose method for restricting the search space from the set of all models $\{0, 1\}^n$ to models of a knowledge base Σ, if particular models of Σ are computed.

Problem ALL-CMODELS requests the complete output of Σ in terms of its characteristic models. In ALL-MODELS, we might be satisfied if some models are initially produced fast and then the enumeration slows down; this can be useful if we want to find some "good" model within limited time. On the other hand, in ALL-CMODELS, quick generation of a few characteristic models is less important than a good overall behavior.

From the results in [24], it easily follows that the output size of problem ALL-MODELS may be exponential in the input size (i.e., the number of characteristic models), even if $l = 1$. Hence, a polynomial time algorithm for this problem is impossible, and the notion of efficient computation has to be reconsidered. A proposal in this vein is an algorithm which enumerates the models with *polynomial delay* [22], 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* [22], i.e., polynomial in the *combined* size of input and output; if no polynomial total time algorithm exists, then a problem may be considered as intractable.

As discussed above, the model-based paradigm has been proposed to speed up on-line reasoning. It is therefore important to know, how reasoning from the logical intersection of theories can be accomplished. In the seminal paper [24], deduction and abduction from a Horn theory, represented by its characteristic models, have been considered, and both were shown to be tractable. We thus consider these two modes of reasoning on the intersection Σ of Horn theories $\Sigma_1, \ldots, \Sigma_l$ represented by their characteristic models. The main issues here are whether similar benign results as in [24, 25] can be obtained, and in particular how a suitable reasoning procedure, given the characteristic models of $\Sigma_1, \ldots, \Sigma_l$ and the query, should proceed.

Main results. We have addressed all the problems from above, and found answers to all of them. Some of the results, e.g., that deduction from an intersection can be done fast, and the hardness of computing all characteristic models, are rather unexpected. Briefly, the main results of this paper can be summarized as follows.

• Problems MODEL and CMODEL are both solvable in polynomial time. In fact, we show that the least (i.e., unique minimal) model of Σ is computable in time linear in the input size, and hence problem MODEL is solvable in linear time. As shown in [14], the least model of a Horn theory given by a Horn formula can be found in linear time; hence, we obtain that under both formula- and model-based representation, computing some model of Σ, and in particular the least model of Σ, is possible in linear time. As a consequence, under both representations also the consistency problem, i.e., deciding whether $\Sigma \neq \emptyset$, can be solved in linear time.

• Problem ALL-MODELS can be solved with polynomial delay; we have developed a respective enumeration algorithm which produces one model at a time. Also this result parallels a polynomial time result under formula-based representation. In fact, the models of a Horn theory (and thus of Σ) given by a Horn formula, can be enumerated with polynomial delay; see e.g. [12] for such a procedure. The delay of algorithm is of the same order as the best known in the formula case [12].

• We show that problem ALL-CMODELS has no polynomial time algorithm. We prove this by describing a family of instances I_n to ALL-CMODELS, $n \geq 1$, for which the output of ALL-CMODELS has 2^n models, while $l = 2$ and $|M_1| = |M_2| = 2n$ (Proposition 5.1). Thus, ALL-CMODELS may have exponential output, and is clearly not polynomially solvable. This improves the result [20, Theorem 6], which states that, in our terminology, ALL-CMODELS for $l = 2$ can not be solved in polynomial time *unless* $\mathrm{P} = \mathrm{NP}$.

• Problem ALL-CMODELS has no polynomial *total* time algorithm, unless P= NP. This is a somewhat negative result, since it means that merging Horn knowledge bases under model-based representation is a complex task in general. In fact, we establish this for $l = 2$, i.e., even the intersection of two Horn theories is hard to compute. We derive this result from the following associated decision problem, which is proved NP-complete: Given the characteristic models of $\Sigma_1, \ldots, \Sigma_l$ and a subset S of the characteristic models of $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$, decide whether some characteristic model exists in $\Sigma \setminus S$.

• Since computing the set of characteristic models $C^*(\Sigma)$ is hard, we also consider the issue of efficient approximations. However, we show that also the natural notion of sound and complete approximation of $C^*(\Sigma)$ is hard to compute. More precisely, we prove that any approximation $N \subseteq \{0,1\}^n$ of $C^*(\Sigma)$, which contains at least a polynomial fraction of $C^*(\Sigma)$ and is only polynomially larger than $C^*(\Sigma)$, is hard to compute. This is a rather strong result, since it shows that even if we want only to compute a significantly large part of $C^*(\Sigma)$, and allow (not to much) junk in the output, we face an intractable problem. To our knowledge, such a type of result is novel in the area of model-based reasoning, and our proof technique may be applied to obtain similar results for a wide range if similar problems.

Furthermore, we prove similar results for computing the maximal models of Σ, which constitute a non-polynomial fraction of $C^*(\Sigma)$. This shows that both some natural *quantitative* (in terms of numbers of models) and *qualitative* (semantically described) approximations of $C^*(\Sigma)$ are hard to compute, and reinforces the view that computing $C^*(\Sigma)$ is really difficult.

• Despite the fact that the number of characteristic models of the intersection Σ of Horn theories $\Sigma_1, \ldots, \Sigma_l$ may be exponential, we show that it is possible to answer deductive queries α to Σ in polynomial time. In particular, for any query α given by a CNF formula, deciding whether $\Sigma \models \alpha$ is possible in $O(mn \sum |M_i|)$ time, where m is the number of clauses in α, n is the number of atoms, and $|M_i|$ the number of models in M_i; if α is a single clause or a positive formula, then deciding $\Sigma \models \alpha$ is possible in linear time. These results are promising, since they show that under taking intersections of Horn theories, the fundamental property of model-based reasoning is preserved that any CNF query posed to a Horn theory can be answered in polynomial time [24, 25], while this is intractable under formula-representation.

• On the other hand, abduction from the intersection Σ of Horn theories $\Sigma_1, \ldots, \Sigma_l$ is intractable, even if l is fixed to 2. We prove that deciding whether a query letter q has an explanation from Σ and a set of assumptions A is NP-complete.

This result tells us that not all benign properties of characteristic models are preserved when we consider intersections of theories. In fact, this indicates that the tractability result for abduction from a single Horn theory Σ_1 is not very robust, and that advantage of particular properties is taken in that case, which is no longer possible if two theories Σ_1, Σ_2 are combined (see Section 6.2 for further discussion).

Usage and significance of the results. Our results give a rather complete picture of the computational properties of using the model-based reasoning approach when different Horn knowledge bases $KB_1, \ldots,$ KB_n are combined by taking their logical intersection. Since this is undoubtedly a principal operation, our algorithms and results are significant for any reasoning system which adopts the model-based approach and incorporates this operation, embedded into a sophisticated combination semantics. Our algorithms are described at a detailed level, and can be easily implemented. Moreover, several algorithms run in linear time (and thus of optimal order), and others are of low-polynomial degree; improvements to linear time (if feasible) seem to require much more effort and sophisticated methods.

Furthermore, the algorithms, together with the complexity results, give us more insight into the potential trade-off between off-line compilation and on-line reasoning. For example, by our results, for ad-hoc on-line deductive reasoning from an intersection Σ, using a direct inference method from $\Sigma_1, \ldots, \Sigma_l$ is more advisable than computing first the characteristic models of Σ, and then applying a polynomial algorithm on them (e.g., the one of [25]). Even in case of repetitive queries, a direct method may be more beneficial if Σ has many characteristic models (of course, on the other hand, if Σ is small while the sets of characteristic models of $\Sigma_1, \ldots, \Sigma_l$ are huge, we may be better off with $C^*(\Sigma)$).

Another aspect is dynamic combination of knowledge bases. For example, suppose there is pool of knowledge bases $KB_1, \ldots, KB_n$; for answering a query, at run-time a subcollection of KB_{i_1}, $KB_{i_2}, \ldots, KB_{i_l}$ of relevant knowledge bases is selected which have to be combined. The different relevant subcollections might vary, and if there are many, we would have to store a number of characteristic set. In the worst case, their number may exponential in n. Even for a small pool size n and under the assumption that only a few knowledge bases are relevant for a query, we might need quite some storage. For example, if $n = 10$ and at most three KB_i's are relevant to a query, then we need to store $\binom{10}{2} + \binom{10}{3} = 45 + 120 = 165$ characteristic sets; if the number of relevant $KB_i's$ is increased to four, we need 275 characteristic sets. Thus, in such a scenario, a direct reasoning strategy which employs our deduction algorithms is preferable.

This becomes even more evident, if we take updates and changes to the knowledge bases into account; an update to a single knowledge base KB_i requires to recompile the characteristic subsets of the subcollection to which KB_i belongs; in the above example, their number is $\binom{9}{1} + \binom{9}{2} = 9 + 36 = 45$ (respectively, 129) for subcollections of size at most three (respectively, at most four). Of course, a mixed strategy is viable in which for some subcollections the (small) characteristic sets are prestored and on others direct methods are applied.

For abductive queries, we have a picture similar to deduction yet different. Here, any current method for answering abductive queries requires exponential time; however, while computing the characteristic models requires exponential space in general, abductive queries can be solved in polynomial space. Observe also that an obliterative reasoning approach, in which characteristic models are enumerated and deleted for avoiding space problems, is not profitable, since it is intractable to tell when the last characteristic model has been found.

Extension of this work. Characteristic models have been generalized to Non-Horn theories in [29], by making use of *monotone theory* [6], a characterization of Boolean functions introduced in computational learning. The approach in [29] is promising, since many advantages of Horn theories carry over to Non-Horn theories. In this direction, we further investigate *extended Horn* theories, which contains both Horn and *reverse* Horn theories, i.e., theories which become Horn by negating all elementary propositions.

It appears that for extended Horn theories, both finding some model and finding some characteristic model are intractable. As a consequence, polynomial total time algorithms for finding all models and all characteristic models, respectively, are unlikely to exist. Moreover, this means that both deduction and abduction of atomic queries from an intersection of theories, given by their characteristic models, is intractable in this case. These results indicate that from the computational side, a generalization of the characteristic models approach for intersections of theories is not immediately feasible, in the sense that both the off-line compilation and the on-line reasoning by direct methods are expensive in general.

Structure of the paper. The remainder of this paper is organized as follows. In the next section, we recall some basic concepts and introduce notation. In Section 3, we consider problem MODEL and model checking, i.e., recognition of a model from an intersection. We then address in Section 4 the problem CMODEL, as well as characteristic model checking. After that, we study in Section 5 the problems ALL-MODELS and ALL-CMODELS, where we show that the output of ALL-CMODELS can be exponential in the input. In Section 6, we consider deduction and abduction from the intersection of Horn theories. Section 7 is devoted to address a possible generalization of our results to extended Horn theories. The final Section 8 discusses further aspects and concludes the paper.

In order not to distract from the flow of reading, longer proofs and technical details of proofs have been moved to Appendix A.

2 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 literal ℓ is an atom or its negation.

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 any set $B \subseteq \{1, \ldots, n\}$, 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 a *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. We shall denote by $\widehat{\Sigma}$ the set of clauses from a Horn CNF representing Σ.[1]

Horn theories Σ have a well-known model-theoretic characterization (see e.g. [34], and [13] for a proof in the propositional case). Denote by $v \bigwedge w$ componentwise AND of vectors $v, w \in \{0,1\}^n$, and by

[1]Observe that $\widehat{\Sigma}$ is not uniquely defined; we use this as a conversion of a set of models into an equivalent formula, which is needed in some contexts.

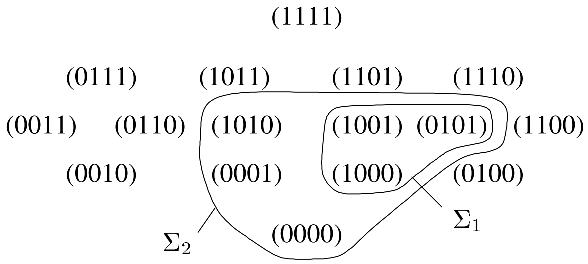

Figure 1: Space of all models for $n = 4$ and theories Σ_1, Σ_2

$Cl_\wedge(S)$ 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 $\Sigma \neq \emptyset$ has the *least* (i.e., *unique minimal*) *model* $v = \bigwedge_{w \in \Sigma} w$, i.e., $\min(\Sigma) = \{v\}$. Here, we use the notation $\bigwedge_{w \in S} w$, where $S \subseteq \{0,1\}$, for the componentwise AND of all vectors in S; in particular, for empty S, by definition $\bigwedge_{w \in S} w = (11...1)$.

E.g., consider $\Sigma_1 = \{(0101), (1001), (1000)\}$ and $\Sigma_2 = \{(0101), (1001), (1000), (0001), (0000)\}$ (see Figure 1). Then, for $v = (0101)$ and $w = (1000)$, we have $v, w \in \Sigma_1$, 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. It can be represented by the Horn CNF $\overline{x}_3 \wedge (\overline{x}_1 \vee \overline{x}_2) \wedge (\overline{x}_2 \vee x_4)$; hence, $\widehat{\Sigma} = \{\overline{x}_3, \overline{x}_1 \vee \overline{x}_2, \overline{x}_2 \vee x_4\}$.

For any Horn theory Σ, a model $v \in \Sigma$ is called *characteristic* [24] (or *extreme* [13]), if $v \notin Cl_\wedge(\Sigma \setminus \{v\})$. The set of all characteristic models of Σ, which we call 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 $C^*(\Sigma_2) = \Sigma_1$. We remark that the characteristic set of Horn theories without negative clauses has been studied in the context of relational databases, where it is known as the generating set [3]; see [28] for a discussion.

Throughout this paper, we suppose that sets of vectors $S \subseteq \{0,1\}^n$ are represented in the standard way, i.e., each model $v \in \{0,1\}^n$ is stored as a sequence $v_1 v_2 \cdots v_n$ of 0's and 1's. However, our algorithms can be adapted for other forms of storage, e.g. a model tree given by a binary decision tree, as well.

3 Finding and Recognizing a Model

In this section, we consider the problem of finding some model of the logical intersection of Horn theories which are represented by their characteristic models. More formally, this problem is specified as follows:

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 $\Sigma \neq \emptyset$; otherwise, "No".

The main result of this section is that such a model, and in fact the least model of Σ, is computable in linear time. Moreover, we obtain that model checking for Σ, i.e., recognizing the members of Σ, is also possible in linear time.

We start with the following lemma, which is useful for our purposes:

Lemma 3.1 *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

$$v \geq \bigvee_{i=1}^{l} \Big(\bigwedge_{w \in C^*(\Sigma_i)} w \Big). \tag{3.1}$$

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 (3.1). $\Box$

Based on the lemma, we can find a model of Σ 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 $\Sigma_1, \ldots, \Sigma_l$, respectively. If they all coincide, then $v = v^{(1)}$ is a model of Σ, which is output. Otherwise, exploiting Lemma 3.1, 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: A model $v \in \Sigma = \bigcap_{i=1}^{l} \Sigma_i$, if $\Sigma \neq \emptyset$; otherwise, "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. $\Box$

Example 3.1 Let $M_1 = C^*(\Sigma_1) = \{(0110), (0011), (1010)\}$ and $M_2 = C^*(\Sigma_2) = \{(1110), (0111), (0011)\}$. The corresponding Horn theories are, under formula-based representation, $\widehat{\Sigma}_1 = \{\overline{x}_1 \vee \overline{x}_2, \overline{x}_1 \vee \overline{x}_4, \overline{x}_2 \vee \overline{x}_4, x_3\}$ and $\widehat{\Sigma}_2 = \{\overline{x}_1 \vee \overline{x}_4, \overline{x}_1 \vee x_2, x_3\}$.

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 $\Sigma = \{(0110), (0010), (0011)\}$ (using formulas, $\widehat{\Sigma} = \{\overline{x}_1, \overline{x}_2 \vee \overline{x}_4, x_3\}$); thus, the output of $v = (0010)$ is correct. $\Box$

An analysis of the run time of the above algorithm gives us the following result.

Theorem 3.1 *Problem* MODEL *can be solved using algorithm* MODEL *in* $O(n^2 \sum_{i=1}^{l} |M_i|)$ *time.*

As an immediate corollary to this result, the consistency of the intersection Σ of Horn theories $\Sigma_1, \dots, \Sigma_l$ is decidable in $O(n^2 \sum_{i=1}^{l} |M_i|)$ time. We do not state this result at this point, since as will show below, the problem can be solved faster.

Recall that since Horn theories are closed under intersection, any Horn theory Σ has the least model $\bigwedge_{w\in\Sigma} w$. In fact, from the working of algorithm MODEL, it is not hard to see that it actually finds this particular model of Σ.

Example 3.2 Let us reconsider Example 3.1. There, MODEL outputs the vector (0010), which is the least model of $\Sigma = \{(0110), (0010), (0011)\}$. □

Corollary 3.1 *Given the characteristic sets $C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, \dots, l$, 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 $\Sigma = \emptyset$.*

Using sophisticated data structures, it is possible to adapt algorithm MODEL such that it runs in time $O(n \sum_{i=1}^{l} |M_i|)$, i.e., in time linear in the input size. Basically, the method is to use lists for cross-references and counters to avoid that the same bit of the input is examined more than a few (constant many) times. We describe this more in detail; the use of similar data structures may be beneficial for speeding up other reasoning algorithms.

The operations we need to perform are

(a) to compute $\bigwedge_{w\in Q_i} w$ (i.e., to compute the set of components k such that $w_k = 1$ holds for all $w \in Q_i$), and $u := \bigvee_{i=1}^{l} (\bigwedge_{w\in Q_i} w)$ and

(b) to update Q_i by removing some models from Q_i.

Recall that the vector u monotonically increases in the execution of MODEL, and observe that the sets Q_i monotonically decrease.

For operation (a), we use counters $\#Q_{i,1}, \#Q_{i,2}, \dots, \#Q_{i,n}$ so that $\#Q_{i,k}$ tells how many models in Q_i have value 1 in component k; i.e., $\#Q_{i,k} = |\{w \in Q_i \mid w_k = 1\}|$. In order to find out the counters with a certain value quickly, we prepare buckets $B_i[0], B_i[1] \dots, B_i[m]$, where $m = |Q_i|$, for each i so that component k (i.e., the counter $\#Q_{i,k}$ via a reference) is in bucket $B_i[\#Q_{i,k}]$. Moreover, we use a counter $\#Q_i$ that tells the number $|Q_i|$ of vectors in Q_i.

For operation (b), we keep lists $L_{i,k}$ of references to all the models $w \in Q_i$ such that $w_k = 0$, and we establish a pointer from each component k with $w_k = 0$ to w in list $L_{i,k}$. Figure 2 shows the data structure for $Q_1 = \{(0101), (0110)\}$.

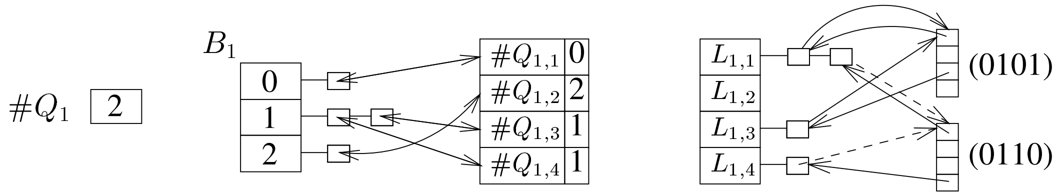

Figure 2: Data structure for set $Q_1 = \{(0101), (0110)\}$ used in MODEL+.

We furthermore prepare a bucket B so that $i \in B$ if $B_i[\#Q_i] \neq \emptyset$.

The algorithm, described in detail below, first scans the input and builds the data structures. After that, it proceeds in a manner similar to MODEL, and processes the sets of models.

Algorithm MODEL+
Input: Characteristic Sets $M_i = C^*(\Sigma_i) \subseteq \{0,1\}^n$ of Horn theories Σ_i, $i = 1, 2, \ldots, l$.
Output: A model $v \in \Sigma$ $(= \bigcap_{i=1}^{l} \Sigma_i)$ if $\Sigma \neq \emptyset$; otherwise, "No".

Step 0. **for** each $i = 1, 2, \ldots, l$ **do** $Q_i := M_i$;
$u := (0, 0, \ldots, 0) \in \{0,1\}^n$.
scan the input to set up the initial counters and buckets described above.

Step 1. **if** $\#Q_i = 0$ holds for some i **then** output "No" and halt;

Step 2. **if** $B = \emptyset$ **then** output $v := u$ and halt
else begin select an arbitrary $i \in B$;
  for each k in $B_i[\#Q_i]$ **do begin**
      (* all models in Q_i have "1" at component k *)
    $u_k := 1$;
    for each $h = 1, 2, \ldots, l$ **do begin**
        (* update the buckets and lists related to Q_h *)
      remove k from $B_h[\#Q_{h,k}]$;
      while there is a model w in $L_{h,k}$ **do begin**
          (* eliminate a $w \in Q_h$ with $w_k = 0$ *)
        $\#Q_h := \#Q_h - 1$;
        for each $j = 1, 2, \ldots, n$ **do**
          if $w_j = 0$ **then** remove w from $L_{h,j}$;
          elsif j is in $B_h[\#Q_{h,j}]$ **then begin**
              (* $w_j = 1$, and update $\#Q_{h,j}$ and $B_h[\cdot]$ *)
            $\#Q_{h,j} := \#Q_{h,j} - 1$;
            move j from $B_h[\#Q_{h,j} + 1]$ to $B_h[\#Q_{h,j}]$
          end{elsif}
      end{while}
    end{for}
  end{for};
  $B := \emptyset$;
  for each $h = 1, 2, \ldots, l$ **do** (* update a bucket B *)
    if $B_h[\#Q_h] \neq \emptyset$ **then** $B := B \cup \{h\}$;
  goto Step 1;
end{if}.

□

Initially, the algorithm sets u to the smallest possible model. If all Q_i are nonempty, but B is empty, i.e., in each Σ_i, at no component k have all the models value 1, then (0,...,0) is a model of each Σ_i, which is output. Otherwise, if some $i \in B$ exists then all models in Q_i (and thus in Σ_i) have value 1 at some component k. If Σ is nonempty, then every model in Σ must have value 1 at this component. Thus, in the candidate model

u the component u_k is set to 1, and all sets Q_j and the data structures are updated accordingly by removing models. If some Q_i becomes empty, then $\Sigma = \emptyset$ is detected; otherwise, the process is continued. To speed up, all selectable components k for Q_i are processed at once.

We omit an example for this algorithm, as it should be clear how it proceeds. The next result establishes that MODEL+ has the desired property.

Theorem 3.2 *Algorithm* MODEL+ *solves problem* MODEL *in* $O(n \sum_{i=1}^{l} |M_i|)$ *time, i.e., in linear time.*

Similarly to algorithm MODEL, we notice the following corollary.

Corollary 3.2 *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 using algorithm MODEL+ in* $O(n \sum_{i=1}^{l} |M_i|)$ *time, i.e., in linear time.*

Yet another corollary is that the problem of model checking for Σ, i.e., recognizing a model from Σ, can be done efficiently.

Corollary 3.3 *Given the characteristic sets* $C^*(\Sigma_i)$ *of Horn theories* $\Sigma_i \subseteq \{0,1\}^n$*,* $i = 1 \ldots, l$*, and some* $v \in \{0,1\}^n$*, deciding whether* $v \in \Sigma = \bigcap_{i=1}^{l} \Sigma_i$ *is possible using algorithm* MODEL+ *in* $O(n \sum_{i=1}^{l} |M_i|)$ *time, i.e., in linear time.*

Proof. Indeed, $v \in \Sigma$ holds if and only if $\bigcap_{i=1}^{l+1} \Sigma_i \neq \emptyset$, where $\Sigma_{l+1} = \{v\}$. Since $C^*(\Sigma_{l+1}) = \{v\}$, we can use algorithm MODEL+ to solve the problem in $O(n \sum_{i=1}^{l+1} |M_i|) = O(n(\sum_{i=1}^{l} |M_i| + 1)) = O(n \sum_{i=1}^{l} |M_i|)$ time (Corollary 3.2). □

4 Finding and Recognizing a Characteristic Model

In this section, we consider the problem of finding some characteristic model of the logical intersection Σ of Horn theories, as well as the problem of recognizing such a model.

The former problem is a first step towards an algorithm for computing all characteristic models; if this problem is hard, then computing all characteristic models is hard as well. The latter problem is relevant to the question of an computational upper bound to the generation of additional characteristic models; if the recognition problem is easy, a sophisticated enumeration procedure may take advantage of this fact and rule out possible candidates for another characteristic model in low-order polynomial time. The main findings are that both computing and recognizing a characteristic model is possible in polynomial time.

4.1 Finding some characteristic model

We first consider computation of some characteristic model, which is the following problem:

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".

For solving this problem, we have to take an approach which is different to finding some arbitrary model. Basically, our method proceeds 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 or its improved version. 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 $C : u = u^{(1)} < u^{(2)} < \ldots < u^{(k)}$ 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'. As we shall prove below, we can exploit the following lemma. Let $Q_i = \{v > u \mid v \in M_i\}$ and $P_{ij} = \{w \in Q_i \mid w_j = 1\}$.

Lemma 4.1 $u \in C^*(\Sigma)$ *holds, if the following condition holds:*

$$\forall j : u_j = 0 \implies \bigcap_{i=1}^{l} Cl_\wedge(P_{ij}) = \emptyset \tag{4.2}$$

(Note that the converse does not hold in general.) On the other hand, if for some j, (4.2) is violated, then any model $v \in Cl_\wedge(P_{ij})$ is a model of Σ with $v > u$; since some characteristic model w such that $w \geq v$ must exist, we can safely select $u' = v$ and replace each M_i by the set $\{w \geq u' \mid w \in P_{ij}\}$. The following example illustrates this algorithm.

Example 4.1 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 (4.2). 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 (4.2) is true; hence, $v = u^{(2)}$ is output.

Note that $C^*(\Sigma) = \{(0110), (0011)\}$; thus, the output of $v = (0110)$ is correct. □

An implementation of this method is straightforward, but rather time consuming. We can save on time by exploiting the following observation: If some j with $u_j = 0$ satisfies (4.2), then we never have to check if (4.2) 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 $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, 2, \ldots, l$.

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 $i = 1, 2, \ldots, l$ **do**
    $Q_i := \{w \in M_i \mid w \geq u\}$;

Step 2. **for** each $j = 1, 2, \ldots, n$ **do**
   if $u_j = 0$ **then begin**

for each $i = 1, 2, \ldots, l$ **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 $i = 1, 2, \ldots, l$ **do**
$Q_i := \{w \in P_{ij} \mid w \geq u\}$;
end;
end;

Step 3. output the model $v := u$. □

Observe that, in this algorithm, the sets P_{ij} are characteristic sets of Horn theories $Cl_\wedge(P_{ij})$. Thus, testing the condition "$\bigcap_{i=1}^{l} Cl_\wedge(P_{ij}) \neq \emptyset$" and finding a model of $\bigcap_{i=1}^{l} Cl_\wedge(P_{ij})$ in Step 3 resorts to an instance of the problem MODEL which we have considered in the previous section, and can be solved in polynomial time.

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

Theorem 4.1 *Problem* CMODEL *can be solved using algorithm CMODEL in $O(n^2 \sum_{i=1}^{l} |M_i|)$ time.*

Similar as in the case of algorithm MODEL, also algorithm CMODEL outputs some particular model of Σ. In fact, from the description of this algorithm, we can easily see that it outputs a maximal model of Σ; recall that $\max(\Sigma) \subseteq C^*(\Sigma)$ holds, while in general not every characteristic model is maximal. We thus obtain the following side result.

Corollary 4.1 *Given the characteristic sets $C^*(\Sigma_i)$ of Horn theories Σ_i, $i = 1, \ldots, l$,* 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 $\Sigma = \emptyset$.* □

Corollaries 3.2 and 4.1 show that the least (i.e., unique smallest) model and some maximal model in Σ can be computed in polynomial time. We come back to the latter result when we will consider abductive reasoning from an intersection.

We remark at this point that finding a maximum model in Σ, i.e., a model which has the largest number of components set to 1, is intractable unless P=NP; this was shown in [20]. For the interested reader, we describe an independently found and different proof in Appendix B. Observe that this problem is also intractable for a theory represented by a Horn formula, while it is trivially solvable in $O(n \sum_{i=1}^{l} |M_i|)$ time, i.e., in linear time from its characteristic set.

4.2 Recognizing a characteristic model

The fact that we can *compute* some characteristic model of the intersection Σ 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 obtaining this result is the following lemma.

Lemma 4.2 *Let Σ be a Horn theory and v be a model in Σ. Then $v \notin C^*(\Sigma)$ holds if and only if*

$$v \neq (11\cdots 1) \text{ and } v \;=\; \bigwedge_{w\in\min(\Sigma_v)} w, \quad \text{where } \Sigma_v = \{w \in \Sigma \mid w > v\}.$$

Proof. The if-part is obvious. For the only-if-part, let $v \notin C^*(\Sigma)$. Then $v = \bigwedge_{w\in\Sigma_v} w$. Clearly, $\bigwedge_{w\in\min(\Sigma_v)} w \geq \bigwedge_{w\in\Sigma_v} w\,(= v)$ as $\min(\Sigma_v) \subseteq \Sigma_v$. If $\bigwedge_{w\in\min(\Sigma_v)} w > \bigwedge_{w\in\Sigma_v} w$, then a component j exists such that $u_j = 0$ for some $u \in \Sigma_v$ and $w_j = 1$ for all $w \in \min(\Sigma_v)$. However, this contradicts that some $w \in \min(\Sigma_v)$ exists with $w \leq u$. □

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 $\Sigma_i \subseteq \{0,1\}^n$, $i = 1,\ldots,l$, and a model $v \in \Sigma = \bigcap_{i=1}^{l} \Sigma_i$.

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

Step 0. **if** $v = (11\ldots1)$ **then** output "Yes" and halt
  else $S := \emptyset$.

Step 1. **for** each j with $v_j = 0$ **do begin**
  for each $i = 1,2,\ldots,l$ **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**
   $w^{(j)}$:= the 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 4.2 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 0 of CHECK-CMODEL, $S := \emptyset$; in Step 1, 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 2, the check $v = \bigwedge_{w^{(j)}\in S} w^{(j)}$ yields false (recall that 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. □

Similar as in algorithm CMODEL, the sets $Q_i^{(j)}$ are the characteristic sets of Horn theories $Cl_\wedge(Q_i^{(j)})$, and thus testing the condition "$\bigcap_{i=1}^{l} Cl_\wedge(P_{ij}) \neq \emptyset$" and finding the least model of $\bigcap_{i=1}^{l} Cl_\wedge(P_{ij})$ in Step 3 can be done in polynomial time.

An analysis of the running time of algorithm CHECK-CMODEL yields then the following result.

Theorem 4.2 *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 using algorithm* CHECK-CMODEL *in $O(n^2 \sum_{i=1}^{l} |M_i|)$ time.*

We conclude this section by remarking that CHECK-CMODEL and CMODEL can be suitably combined into another algorithm for computing a (not necessarily maximal) characteristic model in polynomial time.

5 Computing all Characteristic Models and all Models

We now turn to the issue of generating all models and all characteristic models of a theory Σ, where Σ is the intersection of Horn theories $\Sigma_1, \dots, \Sigma_l$. Let us first consider computing all characteristic models.

5.1 Computing the characteristic set of an intersection

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

Proposition 5.1 *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. Fix n, and define sets of vectors $S_1, S_2, \subseteq \{0,1\}^{4n}$ as follows. Let $V_i = \{i \cdot n + j \mid j = 1, \dots n\}$, for $i = 0, \dots, 3$ and $V = \bigcup_{i=0}^{3} V_i$ = {1,…,4n}; observe that V_0 contains the first n components, V_1 the next n components etc.

Then,

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

Notice that in S_1, every vector has the penultimate block of n bits set to 0. The other blocks are set to 1, and some bit j in the last block as well as the bit at the same relative position j in either the first or second block, is switched to 0. The set S_2 is similar; with respect to S_1, the penultimate and last block are exchanged.

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

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

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 $C^*(\Sigma_2) = S_2$.

Since all models x^B in $C^*(\Sigma_1)$ (resp., $C^*(\Sigma_2)$) satisfy $B \cap V_2 = \emptyset$ (resp., $B \cap V_3 = \emptyset$), all models $x^B \in \Sigma$ satisfy

$$B \subseteq V_0 \cup V_1, \tag{5.3}$$

i.e., the last $2n$ bits of a model in Σ are always 0.

Define

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

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

$$S = \{(00110000),\ (10010000),\ (01100000),\ (11000000)\}.$$

Observe that $|S| = 2^n$. It can be shown that the equation

$$S = C^*(\Sigma) \tag{5.4}$$

holds (see appendix); since $|S| = 2^n$, this proves the result. □

Let us now state the problem of computing the characteristic models of an intersection more formally.

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 $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$.

The previous proposition tells us that the output size of this problem can be exponential in its input size. Therefore, a polynomial time algorithm in the input size is impossible. This improves the result [20, Theorem 6], which states the ALL-CMODELS for $l = 2$ is not solvable in polynomial time unless P = NP; by Proposition 5.1, this is true regardless of whether P = NP holds.

However, we still might hope that ALL-CMODELS has a polynomial total time algorithm. However, this hope does not come true, as the following related problem is intractable.

Problem ADD-CMODEL
Input: Characteristic sets $M_i \subseteq \{0,1\}^n$, of Horn theories Σ_i, $i = 1, 2, \ldots, l$, and a set $S \subseteq C^*(\Sigma)$, where $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$.
Question: $C^*(\Sigma) \setminus S \neq \emptyset$?

Theorem 5.1 *Problem* ADD-CMODEL *is* NP-*complete, and* NP-*hardness holds even if* $l = 2$ *is fixed.*

Proof. Given a candidate model v, we can by Theorem 4.2 check the condition $v \in C^*(\Sigma) \setminus S$ in polynomial time. Thus ADD-CMODEL is in NP.

We prove NP-hardness by a reduction from the satisfiability problem (SAT) [19]; we define for a given CNF formula $\Phi = \bigwedge_{i=1}^{m} C_i$ on n atoms $x_1, \ldots, x_n$ 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 if and only if Φ is unsatisfiable.

Without loss of generality, we make the following assumptions: (i) Every literal in $L = \{x_i, \overline{x}_i \mid 1 \leq i \leq n\}$ appears in Φ, but no literal appears in all clauses; and (ii) Φ does not become a tautology by fixing the truth value of any two atoms x_i and x_j. It is easy to see that these restrictions on Φ do not affect the NP-completeness of SAT.

Define $V = V_L \cup V_1 \cup V_2$, where

$$\begin{aligned} V_L &= \{1, 2, \ldots, n, \overline{1}, \overline{2}, \ldots, \overline{n}\}, \\ V_1 &= \{n+1, n+2, \ldots, n+m\}, \\ V_2 &= \{n+m+1, n+m+2, \ldots, n+2m\}. \end{aligned}$$

Intuitively, the elements in V_L correspond to the literals in L, and the elements j in V_i, $i = 1, 2$, correspond to clauses C_j in Φ. Now we define the instance of our problem as follows:

$$C^*(\Sigma_1) = T_{1,1} \cup T_{1,2}, \tag{5.5}$$

$$
\begin{aligned}
\text{where } T_{1,1} &= \{x^{(V_1\setminus\{n+j\})\cup(V_L\setminus\{q\})} \mid n+j \in V_1, q \in C_j\}, \\
T_{1,2} &= \{x^{(V_L\setminus\{k,\overline{k}\})\cup V_2} \mid k \in V_L\}; \\
C^*(\Sigma_2) &= T_{2,1} \cup T_{2,2}, \qquad (5.6)\\
\text{where } T_{2,1} &= \{x^{(V_2\setminus\{n+m+j\})\cup(V_L\setminus\{q\})} \mid n+m+j \in V_2, q \in C_j\}, \\
T_{2,2} &= \{x^{(V_L\setminus\{k,\overline{k}\})\cup V_1} \mid k \in V_L\}; \\
S &= S_1 \cup S_2, \qquad (5.7)\\
\text{where } S_1 &= \{x^{V_L\setminus\{k,\overline{k}\}} \mid k \in V_L\}, \qquad (5.8)\\
S_2 &= \{x^{V_L\setminus\{k,\overline{k},q\}} \mid k,q \in V_L \text{ with } q \neq k, \overline{k}\}, \qquad (5.9)
\end{aligned}
$$

where $q \in C_j$ denotes that the literal corresponding to q appears in clause C_j (e.g., for a $C_1 = (x_1 \vee \overline{x}_3 \vee x_4)$, we write $1, \overline{3}, 4 \in C_1$), and $\overline{\overline{k}} = k$.

Observe that all vectors in $T_{1,1}$ have value 0 at the components in V_2. Intuitively, every vector in $T_{1,1}$ represents the choice of a literal $q \in C_j$, which is represented by switching in blocks of 1's for V_1 and V_L the components corresponding to q and C_j to 0. By selecting one such vector for every clause, we obtain a collection of literals such that satisfying all these literals makes Φ true. The interpretation of the vectors in $T_{2,1}$ is similar, with the roles of V_1 and V_2 interchanged. An arbitrary selection of literals might include opposite literals q and $\overline{q}$; such illegal selections are trapped by the vectors in $T_{1,2}$ and $T_{2,2}$, which give rise to the characteristic vectors of Σ in S. An additional characteristic model of Σ exists, precisely if there exists a legal choice of literals which satisfies the formula Φ.

The details of the proof can be found in the appendix. □

The result may be intuitively explained by the fact that a characteristic model is a special model, which must satisfy some intersection condition. While it is feasible to check this condition for a given model, it is difficult to find model which satisfies this condition and additional constraints. There is an exponential number of candidates, and we have no efficient method at hand by which this candidate space can be substantially reduced. For problems of a similar characteristics, a polynomial time algorithm is usually obtained by exploiting some benign structural property; in the case of computing an additional characteristic model, no such property is apparent and the result tells us that any such property is hard to find.

Corollary 5.1 *Given the characteristic sets $M_i \subseteq \{0,1\}^n$ of Horn theories Σ_i, $i = 1, 2, \ldots, l$, and a set $S \subseteq C^*(\Sigma)$, where $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$, deciding whether $S = C^*(\Sigma)$ is* co-NP*-complete.*

Exploiting Theorem 5.1, we obtain the next theorem.

Theorem 5.2 *There is no polynomial total time algorithm for problem* ALL-CMODELS*, unless* P=NP.

Proof. Towards a contradiction, assume that there is an algorithm $\mathcal{A}$ for ALL-CMODELS with polynomial running time $p(I, O)$, where I is the input length and O the output length. We then solve ADD-CMODEL using $\mathcal{A}$ as follows. Execute $\mathcal{A}$ until either (i) it halts or (ii) time $p(I, |S|)$ is reached. In case (i), output "Yes" if $\mathcal{A}$ 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 Theorem 5.1 unless P=NP. □

Observe that this result strengthens [20, Theorem 6] in another way, by stating that no polynomial algorithm exists even if we relativize the run time by taking a possible exponential output size into account.

Practically speaking, this means that computing all characteristic models of an intersection is a hard problem.

5.2 Approximation of the characteristic set

In the previous subsection, we have shown that computing the characteristic set of the intersection Σ of Horn theories $\Sigma_1, \ldots, \Sigma_l$ is an intractable problem. As with other hard problems in the context of reasoning (cf. [10, 37]), it is thus natural to ask whether we can compute a suitable approximation of $C^*(\Sigma)$ in polynomial total time.

Towards this goal, we first have to agree on what a suitable approximation of $C^*(\Sigma)$ is. Recall that $C^*(\Sigma)$ is the unique smallest set $S \subseteq \Sigma$ of models such that $\Sigma = Cl_\wedge(S)$ holds. A reasonable requirement is that an approximation M of $C^*(\Sigma)$ should only contain models in Σ, i.e., $M \subseteq \Sigma$ should hold. This assures $Cl_\wedge(M) \subseteq \Sigma$; in a sense, this is *soundness* of the representation. On the other hand, it would be desirable that all models in Σ occur in $Cl_\wedge(M)$; i.e., M is *complete* with respect to Σ.

Let us call any set of models M which is sound and complete with respect to Σ (i.e., $C^*(\Sigma) \subseteq M \subseteq \Sigma$ holds), a *conservative approximation* of $C^*(\Sigma)$. Observe that $M = C^*(\Sigma)$ and $M = \Sigma$ are the best and weakest conservative approximations of $C^*(\Sigma)$, respectively. A conservative approximation might be seen as a non-optimal compact representation of Σ, which is however sound and complete for the purpose of reasoning from Σ.

It is now natural to ask whether finding a reasonably sized conservative approximation M of $C^*(\Sigma)$ is tractable, i.e., possible in output polynomial time. Clearly, an M whose size is exponential in the size of $C^*(\Sigma)$ is not reasonable, and thus we limit attention to those M whose size is polynomial in the size of $C^*(\Sigma)$. The next result, however, tells us that finding any arbitrary such conservative approximation is also an intractable problem.

Theorem 5.3 *Let $p(\cdot)$ be any polynomial.*[2] *Then, given the characteristic sets $M_i = C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, 2 \ldots, l$, there is no polynomial total time algorithm for computing a conservative approximation M for $C^*(\Sigma)$, where $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$, such that $|M| \leq p(|C^*(\Sigma)|)$, unless Pol*=NP. *This holds if $l = 2$ is fixed.*

Proof. Assume such a polynomial total time algorithm $\mathcal{A}$ exists. Then, an output-polynomial total algorithm for ALL-CMODELS exists, since we can first apply $\mathcal{A}$, and then remove from its output M every model v such that $v \notin C^*(\Sigma)$ in polynomial time (Theorem 4.2); observe that the size of the intermediate result M is polynomial in the output $C^*(\Sigma)$. By Theorem 5.2, ALL-CMODELS has no polynomial total time algorithm unless P = NP, from which the result follows. $\square$

This result shows that for gaining tractability, we have to give up on conservative approximations. Thus, either soundness or completeness of the approximation (or both) has to be abandoned. It seems natural, however, to retain soundness, since completeness may be dispensable for answering certain queries to a knowledge base (see Section 6.1 for further discussion).

When giving up completeness, we have to decide which part of $C^*(\Sigma)$ should be omitted, in order to be able to use the result of the approximation. This is not straightforward, however, and depends on the intended use of the knowledge base. We do not embark on this general issue here, but point out some principal limitations to such an approach. Our next result shows that any approximation of $C^*(\Sigma)$, regardless

[2] Here, and in the rest of this paper, we assume as usual that polynomials are monotone increasing.

of being sound or not, which returns a polynomial-size fraction of $C^*(\Sigma)$ and is polynomially bounded in $|C^*(\Sigma)|$, is intractable, i.e., there is no polynomial total algorithm for its computation.

Theorem 5.4 *Let $p(\cdot)$, $q(\cdot)$ be any polynomials. Then, there is no polynomial total time algorithm $\mathcal{A}$ for computing, given the characteristic sets $M_i = C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, 2 \ldots, l$, a set of models $N \subseteq \{0,1\}^n$ such that (i) $|C^*(\Sigma)| \leq q(|N \cap C^*(\Sigma)|)$ and (ii) $|N| \leq p(|C^*(\Sigma)|)$, unless* $\mathrm{P} = \mathrm{NP}$. *This holds if $l = 2$ is fixed.*

As a consequence, there is no polynomial total time algorithm for computing half of the characteristic set, say, or any constant fraction of it. Thus, since a quantitative approximation of $C^*(\Sigma)$ is infeasible, we would have to consider qualitative approximations, i.e., meaningful semantical portions of $C^*(\Sigma)$ which are sufficient for certain purposes.

An example would be the maximal models $\max(\Sigma)$ of an intersection Σ. Recall that $\max(\Sigma)$ is included in $C^*(\Sigma)$, and as easily seen, this set may be exponentially smaller than $C^*(\Sigma)$, and thus the above results do not apply. Moreover, it is easily seen that $\max(\Sigma)$ is sound and complete with respect to answering *negative* deductive queries α to Σ, i.e., deciding whether $\Sigma \models \alpha$, where $\alpha = C_1 \wedge \cdots \wedge C_m$ is a conjunction of negative clauses $C_i = \overline{x}_{i_1} \vee \cdots \vee \overline{x}_{i_k}$ (for reasoning from Σ, see Section 6).

Thus, we might be interested whether for this particular portion of $C^*(\Sigma)$, a polynomial algorithm is feasible. As we know from Corollary 4.1, computing some maximal model of $\Sigma = \bigcap_i \Sigma_i$ is possible in polynomial time, and from the proof of Proposition 5.1, it follows that exponentially many maximal models may exist. Thus, all we can expect is a polynomial total time algorithm.

It turns out that there is no such algorithm, and also approximation of $\max(\Sigma)$ is hard. By a slight adaptation of the proof of Theorem 5.1, we obtain that finding an additional maximal model is NP-hard; moreover, it follows from Corollary 3.3 and Lemma 4.1 (cf. also Lemma 6.1 below) that recognizing a maximal model is polynomial. Thus, by analogous argumentation as in the proofs in the previous subsections, we obtain the following result.

Theorem 5.5 *Given the characteristic sets $M_i = C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, \ldots, l$, (i) it is* co-NP*-complete to decide whether* $\max(\Sigma) = S$*, where $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$ and S is a given set of maximal models of Σ, (ii) there is no polynomial total time algorithm for computing* $\max(\Sigma)$*, unless* $\mathrm{P} = \mathrm{NP}$*, (iii) there is no polynomial total time algorithm for computing a polynomial approximation of* $\max(\Sigma)$*, unless* $\mathrm{P} = \mathrm{NP}$.

Here, "polynomial approximation" in (iii) is understood in the setting of Theorem 5.4.

There is no reason for raising one's hands in desperation about all these negative results. After all, one of the ideas behind characteristic models was off-line compilation for efficient on-line reasoning. For such off-line compilation, we may be willing to pay a high computational price. The results from above just tell us that in the case of intersection of knowledge bases, we indeed have to pay that price. However, this does not mean that we should abandon the search for reasonable and good algorithms for compilation. In the rest of this section, we present an algorithm for enumerating all arbitrary models of Σ with polynomial delay; this algorithm may be used as a basis for an algorithm computing $C^*(\Sigma)$ in some contexts.

5.3 Computing all models of an intersection

Let us now consider the problem of computing all models of an intersection. The formal statement of this problem is as follows:

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 $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$.

It turns out that this problem is easier than the related problem ALL-CMODELS, as we shall present a polynomial delay algorithm for it.

Informally, the reason is that we can output some model and then systematically shrink for the next step the theory $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$ to a subset Σ' of models, such that no model in Σ' has been output so far and finding a model in Σ' is efficiently possible.

The algorithm we present here is based on our results from above and the method of *dynamic lexicographic enumeration* [12]. This method improves on a previous technique in [40], and was used for efficient enumeration of the models of a Horn theory represented by a Horn formula. The idea is to restrict Σ to the subset Σ' of models different from the models $v^{(1)}$, $v^{(2)}$, …, $v^{(k)} = v$ which have been output in the previous steps, and to select from Σ' a model w which has the largest common prefix with v. By clever bookkeeping of which prefixes have been considered, it is possible to find such a model in Σ' (so $\Sigma \neq \emptyset$) quite efficiently.

The bookkeeping is done by maintaining a binary vector $mark \in \{0,1\}^n$, where the value of $mark_i$ indicates whether the search for the models $w \in \Sigma$ with common prefix up to $i-1$ (i.e., $v_j = w_j$ for $1 \leq j < i$ and $v_i \neq w_i$) has already been successfully attempted ($mark_i = 1$) or not ($mark_i = 0$); after the output of the first model $v^{(1)}$, $mark$ is initialized to the zero vector $(00\ldots0)$.

The algorithm, ALL-MODELS, uses 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: A 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 any such model exists; "No", otherwise. □

By means of this procedure, it is possible to check whether a partial vector (given by $b_1, \ldots, b_r$) can be completed to a model in Σ, and such a model is returned in case. Observe that this procedure can be implemented as described in the proof of Lemma 6.1, and such that it returns even the least model among all possible outputs.

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($\Sigma_1, \ldots, \Sigma_l, v_1, \ldots, v_{i-1}, 1 - v_i$);
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. □

(The algorithm can be reformulated to be slightly more efficient; we use this more readable version for the sake of simplicity). We illustrate the algorithm on the following example.

Example 5.1 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 the 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 5.6 *Algorithm* ALL-MODELS *is a polynomial delay algorithm for problem* ALL-MODELS*, where the delay is $O(n^2 \sum_{i=1}^{l} |M_i|)$, i.e., number of atoms times input length.*

By combining the algorithms ALL-MODELS and CHECK-CMODEL, we obtain an algorithm for enumerating all characteristic models of Σ, which is however not a polynomial delay algorithm. Nonetheless, by using ALL-MODELS, we restrict the search space from all vectors in $\{0, 1\}^n$ to models in Σ; if Σ is small, or its size is polynomial in the size of $C^*(\Sigma)$, then this algorithm runs in polynomial total time. The algorithm may by particularly attractive if the size of the input $I = M_1, \ldots, M_l \subseteq \{0, 1\}^n$ is small measured in the number n; observe that in case I is exponential in n, computing all models as well as all characteristic models is possible in time polynomial in the input size by a brute force search.

6 Reasoning from an Intersection

In this section, we turn our attention to reasoning from an intersection Σ of Horn theories $\Sigma_1, \ldots, \Sigma_l$. In particular, we first consider answering of a deductive query α posed to Σ, and then abduction in the setting where for a propositional letter q, an explanation on the basis of a set A of assumptions and the theory Σ should be found.

6.1 Deduction

One of the striking advantages of model-based reasoning is that large classes of queries to a knowledge base can be evaluated efficiently. It has been shown in [24] that deduction of an arbitrary CNF formula α from a Horn theory Σ is polynomial, if Σ is represented by its characteristic models $C^*(\Sigma)$. To evaluate $\Sigma \models \alpha$, it is sufficient to check whether $\Sigma \models C$ for each clause C in α; this problem can be solved by checking whether some Horn strengthening C' of C, i.e., a Horn clause C' obtained from C by removing all but one positive literal, is true in all characteristic models. As shown in [24, 25], $\Sigma \models \alpha$ is decidable in $O(|C^*(\Sigma)| \cdot |\alpha|^2)$ time, where $|\alpha|$ is the length of α.

Following this paradigm, an ad-hoc query α posed to an intersection Σ of Horn theories $\Sigma_1, \ldots, \Sigma_l$ can be answered in the following way:

1. Compute $C^*(\Sigma)$;
2. apply any (fast) algorithm for deciding $\Sigma \models \alpha$ from $C^*(\Sigma)$.

Example 6.1 Reconsider the theories $\widehat{\Sigma}_1 = \{\overline{x}_1 \vee \overline{x}_2, \overline{x}_1 \vee \overline{x}_4, \overline{x}_2 \vee \overline{x}_4, x_3\}$ and $\widehat{\Sigma}_2 = \{\overline{x}_1 \vee \overline{x}_4, \overline{x}_1 \vee x_2, x_3\}$ from Example 3.1, whose characteristic sets are $C^*(\Sigma_1) = \{(0110), (0011), (1010)\}$ and $C^*(\Sigma_2) = \{(1110), (0111), (0011)\}$, respectively. Suppose we want to know whether $\Sigma \models x_1 \vee x_4 \vee \overline{x}_3$, where $\Sigma = \Sigma_1 \cap \Sigma_2$; observe that the query α is not Horn. After computing $C^*(\Sigma) = \{(0110), (0011)\}$, we check whether $C^*(\Sigma) \models x_1 \vee \overline{x}_3$ or $C^*(\Sigma) \models x_4 \vee \overline{x}_3$ holds, since $x_1 \vee \overline{x}_3$ and $x_4 \vee \overline{x}_3$ are the Horn strengthenings of α. However, both clauses evaluate to false on (0110) and hence $\Sigma \not\models \alpha$ is concluded. Indeed, observe that from $\widehat{\Sigma} = \{\overline{x}_1, \overline{x}_2 \vee \overline{x}_4, x_3\}$, the query α is not derivable. On the other hand, $\Sigma \models \overline{x}_1 \vee x_2$, since x_1 is false in all models of $C^*(\Sigma)$. $\Box$

By the results of the previous section, this approach is infeasible, however, since the computation of $C^*(\Sigma)$ may need truly exponential time. Nonetheless, it is possible to evaluate $\Sigma \models \alpha$ efficiently, by a method which bypasses the computation of $C^*(\Sigma)$. The reason is that the test $\Sigma \models C$, where C is a single clause, can be reduced to a consistency test, which is efficiently solvable. This is a consequence of the next lemma. Let, for any formula ϕ, $mod(\phi)$ denote the set of its models.

Lemma 6.1 *Given the characteristic sets $C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, 2, \ldots, l$, and literals $\ell_1, \ldots, \ell_k$, deciding whether $\Delta = mod(\ell_1 \wedge \cdots \wedge \ell_k) \cap \bigcap_{i=1}^{l} \Sigma_i \neq \emptyset$ and finding the least model of Δ (if it exists) is possible in $O(n \sum_{i=1}^{l} |M_i|)$ time, i.e., in linear time.*

Proof. We can obtain an algorithm as desired by a slight adaptation of the algorithm MODEL+, which fixes the values of components of models according to $\ell_1, \ldots, \ell_k$.

Suppose that $\ell_j = x_{i_j}$, $j = 1, \ldots, h$ and $\ell_j = \overline{x}_{i_j}$, for $j = h+1, \ldots, k$, and that no opposite literals are among $\ell_1, \ldots, \ell_k$ (otherwise, $\Delta = \emptyset$). Modify MODEL+ as follows.

Let $w \in \{0,1\}^n$ be the vector which has value 1 at the components i_j, for all $j = 1, \ldots, h$ and value 0 at all others; i.e., w is the least model of $x_{i_1} \wedge \cdots \wedge x_{i_h}$, and set $N := \{i_j \mid h+1 \leq j \leq k\}$. Then,

- replace in Step 0 the assignment "$Q_i := M_i$" by "$Q_i := \{v \in M_i \mid v \geq w\}$," and the assignment "$u := (0,0,\ldots,0)$" by "$u := w$";
- replace in Step 2 the assignment "$u_k := 1$" by the conditional statement "**if** $k \in N$ **then** output "No" and halt **else** $u_k := 1$" (note that the k there is not connected to the k in the statement of the lemma).

Along the argumentation of the proof of Theorem 3.2, it can be shown that the modified algorithm correctly outputs a model (in fact, the least model) of Δ, if one exists, and "No" otherwise; observe that the search through the space of models is restricted from $\{0,1\}^n$ to all models of $x_{i_1} \wedge \cdots \wedge x_{i_h}$, and that the search is stopped as soon it is recognized that the least model in $mod(x_{i_1} \wedge \cdots \wedge x_{i_h}) \cap \bigcap_{i=1}^{l} \Sigma_i$ must have value 1 at some component j such that $\overline{x}_j$ occurs among $\ell_{h+1}, \ldots, \ell_k$.

It is easy to see that by the above modifications, the order of the run time is not affected and remains $O(n \sum_{i=1}^{l} |M_i|)$. This proves the lemma. □

Theorem 6.1 *Given the characteristic sets $C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1,2,\ldots,l$, and a clause $C = \ell_1 \vee \cdots \vee \ell_k$, deciding whether $\Sigma \models C$ holds is possible in $O(n \sum_{i=1}^{l} |M_i|)$ time, i.e., in linear time.*

Proof. Clearly, $\Sigma \models C$ if and only if $\Sigma \cap mod(\neg C) = \emptyset$ holds. Since $\neg C$ is equivalent to $\overline{\ell}_1 \wedge \cdots \wedge \overline{\ell}_k$, where $\overline{\ell}_i$ denotes the opposite to literal ℓ_i, the result immediately follows from Lemma 6.1. □

Corollary 6.1 *Given the characteristic sets $C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1,2,\ldots,l$, and a CNF formula α, deciding whether $\Sigma \models \alpha$ holds is possible in $O(nm \sum_{i=1}^{l} |M_i|)$ time, where m is the number of clauses in α.*

Proof. Since $\Sigma \models \alpha$ iff $\Sigma \models C$ for each clause C in α, this follows from Theorem 6.1. □

For a particular important class of formulas, we obtain the following result. Recall that a formula ϕ (not necessarily in CNF) is *positive*, if each atoms occurs in it under an even number of negations; in particular, every negation-free formula is positive.

Theorem 6.2 *Given the characteristic sets $C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1,2,\ldots,l$, and a positive formula α, deciding whether $\Sigma \models \alpha$ holds is possible in $O(n \sum_{i=1}^{l} |M_i| + |\alpha|)$ time, where $|\alpha|$ denotes the length of α, i.e., in linear time.*

Proof. Since ϕ is positive, it holds for any theory Σ that $\Sigma \models \alpha$ if and only if $v \models \alpha$, for each $v \in \min(\Sigma)$ (see e.g. [29, Section 3]).

Since Σ is Horn, it has a unique minimal model u (provided $\Sigma \neq \emptyset$), which can be constructed in $O(n \sum_{i=1}^{l} |M_i|)$ time (Corollary 3.2). Moreover, checking whether $v \models \alpha$ is possible in time $O(n + |\alpha|)$. Hence, the result follows. □

6.2 Abduction

Abduction [35] is a principal mode of reasoning which is heavily used in our daily life reasoning. Informally, abduction is the task of finding an explanation for certain observations, based on some background theory describing the relationships between causes and effects. There is a growing literature on this subject, which has been recognized as an important principle of common-sense reasoning (see e.g. [5]) but still has many further applications (see e.g. references in [15]).

More formally, abduction can be defined as follows.

Definition 6.1 Let Σ be a theory, A be a subset of the atoms, and q be an atom. Then, a subset E of literals on atoms from A is an *explanation* for q from Σ and A, if (i) $\widehat{\Sigma} \cup E$ is consistent, and $\widehat{\Sigma} \cup E \models q$. [3]

(Recall that $\widehat{\Sigma}$ transforms a Horn theory Σ into an equivalent set of Horn clauses.) Usually, one is interested in *minimal* explanations, i.e., explanations E which do not contain any other explanation properly.

Example 6.2 Consider the theory $\widehat{\Sigma} = \{\overline{x}_1 \vee \overline{x}_4, \overline{x}_4 \vee \overline{x}_3, \overline{x}_1 \vee x_2\}$. Suppose we want to explain $q = x_2$ from $A = \{x_1, x_4\}$. Then, we find that $E = \{x_1\}$ is an explanation. Indeed, $\widehat{\Sigma} \cup \{x_1\}$ is consistent, and $\widehat{\Sigma} \cup \{x_1\} \models x_2$. Moreover, E is minimal. On the other hand, $E' = \{x_1, \overline{x}_4\}$ is an alternative, non-minimal explanation of x_2. $\Box$

One of the main obstacles for an implementation of abduction is its intrinsic computational cost; under formula-based representation, finding an abductive explanation is NP-complete in the Horn case [38], and is Σ_2^p-complete for general propositional theories [15], which is the prototypical complexity of many nonmonotonic reasoning problems.

However, as shown in [24, 25], finding an explanation is polynomial in the Horn case if Σ is represented by its characteristic models. This was a quite an encouraging result, since it shows that both deduction and abduction from a Horn theory can be done in polynomial time. Since by the results of the previous subsection, also deduction from the intersection Σ of Horn theories $\Sigma_1, \ldots, \Sigma_l$ can be done in polynomial time, it would be advantageous if a similar result can be obtained for abduction.

However, it turns out that the desired generalization of the positive result in [25] is not apparent.

Theorem 6.3 *Given the characteristic sets $M_i = C^*(\Sigma_i) \subseteq \{0,1\}^n$ of Horn theories Σ_i, $i = 1, \ldots, l$, an assumption set $A \subseteq \{x_1, \ldots, x_n\}$, and an atom q from $x_1, \ldots, x_n$, deciding whether q has an explanation from $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$ and A is* NP*-complete.*

Proof. The problem is in NP, since we can guess an explanation E and check in polynomial time whether $\widehat{\Sigma} \cup E$ is consistent (Lemma 6.1) and whether $\widehat{\Sigma} \cup E \models q$, by testing the equivalent condition $\Sigma \models E \supset q$ (Theorem 6.1).

The NP-hardness part is shown by a proper modification of the reduction used in the proof of Theorem 5.1. There, we have constructed from a CNF formula Φ the characteristic sets $C^*(\Sigma_1) = T_{1,1} \cup T_{1,2}$ and $C^*(\Sigma_2) = T_{2,1} \cup T_{2,2}$ of Horn theories Σ_1 and Σ_2, respectively, along with a subset S of the characteristic set of $\Sigma = \Sigma_1 \cap \Sigma_2$, such that some characteristic model $v \in C^*(\Sigma) \setminus S$ exists if and only if Φ is satisfiable.

[3] Observe that in some texts, explanations must be sets of positive literals. As with Horn theories, it is known (cf. [29]) that an explanation exists only if an explanation containing merely positive literals exists; in fact, all minimal explanations are of this form.

We modify the construction as follows. Introduce a new component (i.e., atom) "0", and set this component to 0 for all vectors in $T_{1,2}$ and $T_{2,2}$, and to 1 for all vectors in $T_{1,1}$ and $T_{2,1}$; denote the resulting sets by $T'_{i,j}$, for $i, j = 1, 2$, and let $C^*(\Sigma'_i) = T'_{i,1} \cup T'_{i,2}$, for $i = 1, 2$.

Observe that any vector resulting from the intersection of a set of vectors in $T'_{i,1}$ has value 1 at component 0, while any vector, i.e., resulting from an intersection which involves some vector in $T'_{i,2}$, has value 0 at this component. Moreover, since all vectors in $C^*(\Sigma'_i)$ are incomparable, this set is indeed the characteristic set of a Horn theory Σ'_i, for $i = 1, 2$. Following the argumentation in the proof of Theorem 5.1, it can be seen that each model in $\Sigma' = \Sigma'_1 \cap \Sigma'_2$ has the form x^B for some $B \subseteq V_L \cup \{0\}$, and that each model $v^{(k)} = x^{V_L \setminus \{k, \overline{k}\}}$ belongs to Σ', where $k \in V_L$; notice that $v^{(k)}$ has component 0 set to 0.

Let q be the propositional atom corresponding to the newly introduced component 0, and let A be the propositional atoms corresponding to all other components (alternatively, we could also set $A = V_L$). Intuitively, if we want to explain q, then we must find a model v in Σ which has value 1 at the component 0, and such that if we fix the values of the literals in A to those in v, then it is not possible to switch component 0 to 0 and still have a model of Σ; since $v^{(k)}$ has components k, $\overline{k}$, and 0 all set to 0, such a v must correspond to a choice of literals whose satisfaction makes Φ true.

In the appendix, we show that q has an explanation from Σ and A if and only if Φ is satisfiable; the result follows from this. $\square$

This result shows that the tractability result for abduction in [24] is not very robust. The intuitive reason for the positive result in [24] is that if an explanation exists, then some explanation can be easily found from the maximal models of Σ, which are included in $C^*(\Sigma)$. However, in the case where Σ is an intersection of theories, $\max(\Sigma)$ is not explicitly given, and an exponential number of maximal models may exist. While computing some maximal model is tractable, the computation of a maximal model which gives rise to an explanation E is NP-hard.

Thus, in the general case, abduction from an intersection is intractable. It might therefore be suspected that a strategy of computing $C^*(\Sigma)$ and then running the polynomial algorithm from [24] is useful. However, this may not always be the case, since $C^*(\Sigma)$ requires exponential space in general (and thus its computation takes exponential time), while evaluation of an abductive query is always possible in polynomial space and exponential time, and in some cases even in polynomial time. An example is the following special case, which follows immediately from Theorem 6.1 by simple exhaustive search.

Theorem 6.4 *Given the characteristic sets $M_i = C^*(\Sigma_i) \subseteq \{0,1\}^n$ of Horn theories Σ_i, $i = 1, \ldots, l$, an assumption set $A \subseteq \{x_1, \ldots, x_n\}$, and an atom q from $x_1, \ldots, x_n$, finding an explanation for q from $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$ and A is possible in polynomial time, if the size of A is $O(\log n)$. Moreover, it is possible in $O(n \sum_{i=1}^{l} |M_i|)$ time, i.e., in linear time, if $|A| \leq k$ for some constant k.*

The conclusion we can draw is that we have to look into particular query profiles (frequent or not, tractable or not, etc), and that the decision on which strategy should be followed should be based on the results of this inquiry. Observe that even if space is not an issue, answering polynomially solvable abductive queries may take much longer (even exponentially longer) from the off-line compiled $C^*(\Sigma)$ than under on-line evaluation from $\Sigma_1, \ldots, \Sigma_l$.

7 Non-Horn Theories

In this section, we consider a possible generalization of our results to non-Horn theories. In particular, we consider a class of *extended Horn theories*, which includes Horn theories and a close variant thereof. We shall show that for this particular class, the main problems considered in the previous sections are intractable.

7.1 Generalized characteristic models

We first review monotone theory of Boolean functions introduced in [6], and then recall the definition of characteristic models for arbitrary classes $\mathcal{C}$ of Boolean functions.

For any model $b \in \{0,1\}^n$, we define a partial order $\leq_b$ over $\{0,1\}^n$ by that $v \leq_b w$ holds if and only if $v \oplus b \leq w \oplus b$ holds, where $\oplus$ denotes the XOR operation (i.e., componentwise addition modulo 2; e.g., $(1100) \oplus (0110) = (1010)$). $v \leq_b w$ can also be written as $w \geq_b v$, and $v <_b w$ (resp., $v >_b w$) denotes $v \neq w$ and $v \leq_b w$ (resp., $v \geq_b w$). In other words, if $b_i = 0$, then the order on the i-th component is normal, i.e., $0 <_{b_i} 1$; on the other hand, if $b_i = 1$, the order is reversed, i.e., $1 <_{b_i} 0$. The *monotone extension* of a model $z \in \{0,1\}^n$ with respect to b is defined by

$$\mathcal{M}_b(z) = \{v \in \{0,1\}^n \mid v \geq_b z\},$$

and the *monotone extension* of a theory $\Sigma \subseteq \{0,1\}^n$ with respect to b is defined by

$$\mathcal{M}_b(\Sigma) = \bigcup_{z\in\Sigma} M_b(z).$$

The set of *minimal models* of Σ with respect to b is defined by

$$\min_b(\Sigma) = \{z \mid z \in \Sigma \text{ and no } v \in \Sigma \text{ satisfies } v <_b z\}.$$

Observe that $\min(\Sigma) = \min_{(00\ldots0)}(\Sigma)$ and $\max(\Sigma) = \min_{(11\ldots1)}(\Sigma)$, respectively. $\mathcal{M}_b(\Sigma)$ is now rewritten as

$$\mathcal{M}_b(\Sigma) \;=\; \bigcup_{z\in\min_b(\Sigma)} \mathcal{M}_b(z). \tag{7.1}$$

This is because $\mathcal{M}_b(v) \subseteq \mathcal{M}_b(w)$ holds for all pairs of v and w such that $v \geq_b w$.

It is easy to show the following properties:

$$\Sigma \subseteq \mathcal{M}_b(\Sigma) \tag{7.2}$$

$$b \notin \Sigma \iff b \notin \mathcal{M}_b(\Sigma), \tag{7.3}$$

for all $b \in \{0,1\}^n$. Furthermore, $\mathcal{M}_b$ is monotonic in Σ, distributes over unions $\Sigma_1 \cup \Sigma_2$, and satisfies $\mathcal{M}_b(\Sigma_1 \cap \Sigma_2) \subseteq \mathcal{M}_b(\Sigma_1) \cap \mathcal{M}_b(\Sigma_2)$. Hence, by using (7.3) and (7.2), we obtain

$$\bigcap_{b\notin\Sigma} \mathcal{M}_b(\Sigma) \;\subseteq\; \Sigma \;\subseteq\; \bigcap_{b\in\{0,1\}^n} \mathcal{M}_b(\Sigma) \;\subseteq\; \bigcap_{b\notin\Sigma} \mathcal{M}_b(\Sigma).$$

Consequently, Σ is characterized as follows.

Proposition 7.1

$$\Sigma \;=\; \bigcap_{b\in\{0,1\}^n} \mathcal{M}_b(\Sigma) \;=\; \bigcap_{b\notin\Sigma} \mathcal{M}_b(\Sigma). \tag{7.4}$$

In the right hand side of (7.4), not all models $b \notin \Sigma$ may be necessary to represent Σ, i.e., $\Sigma = \bigcap_{b\in B} \mathcal{M}_b(\Sigma)$ may hold for some $B \subseteq \{0,1\}^n \setminus \Sigma$. This leads to the following definition.

Definition 7.1 ([6]) *A set of models B is called a* basis for a theory Σ, *if* $\Sigma = \bigcap_{b\in B} \mathcal{M}_b(\Sigma)$ *holds. Furthermore, B is called a* basis for a class of theories $\mathcal{C}$, *if it is a basis for all the theories in $\mathcal{C}$.* □

Clearly, $\{0,1\}^n$ and $\{0,1\}^n \setminus \Sigma$ are bases for any theory Σ, and $\{0,1\}^n$ is a basis for any class of theories $\mathcal{C}$. It is known that for the class of Horn theories $\mathcal{C}_H$,

$$B_H = \{b \mid \parallel b \parallel \geq n-1\}, \tag{7.5}$$

is a basis [29], where $\parallel x \parallel = \sum_{i=1}^{n} x_i$.

Call a theory Σ *reverse Horn* [29], if by negating all atoms x_i, the resulting theory is Horn; i.e., Σ is reverse Horn, if and only if Σ is closed under union of models (i.e., $v, w \in \Sigma$ implies $v \bigvee w \in \Sigma$, where $\bigvee$ is componentwise OR; e.g., $(1100) \bigvee (0110) = (1110)$). It is easy to see that

$$B_{RH} = \{b \mid \parallel b \parallel \leq 1\} \tag{7.6}$$

is a basis of the class of reverse Horn theories $\mathcal{C}_{RH}$.

Monotone theory and the concept of basis has been used to define characteristic models of arbitrary theories as follows.

Definition 7.2 ([29]) *Let $\mathcal{C}$ be a class of theories, and let B be a basis for $\mathcal{C}$. For a theory $\Sigma \in \mathcal{C}$, we define the* set of characteristic models $\Gamma^B(\Sigma)$ *with respect to B as follows:*

$$\Gamma^B(\Sigma) = \bigcup_{b\in B} \min_b(\Sigma). \tag{7.7}$$

This definition can be regarded as a generalization of that for Horn theories, since

$$C^*(\Sigma) = \Gamma^{B_H}(\Sigma) \tag{7.8}$$

holds for all Horn theories Σ [29]. Note that $\max(\Sigma)$, which is a subset of $C^*(\Sigma)$, can be represented by $\max(\Sigma) = \min_{(11\ldots1)}(\Sigma)$. Any other model v in Σ is minimal with respect to some b with $\parallel b \parallel = n-1$.

7.2 Extended Horn theories

As a generalization of the class of Horn theories (see (7.5)) and reverse Horn theories (see (7.6)), let us define $B_{EH} \subseteq \{0,1\}^n$ by

$$B_{EH} = \{b \in \{0,1\}^n \mid \parallel b \parallel \geq n-1 \text{ or } \parallel b \parallel \leq 1\}. \tag{7.9}$$

A theory $\Sigma \subseteq \{0,1\}^n$ is called *extended Horn* if B_{EH} is a basis for Σ, and let $\mathcal{C}_{EH}$ denote the class of extended Horn theories. Clearly, any Horn theory and reverse Horn theory are always extended Horn.

In the remainder of this section, we consider the problems MODEL, CMODEL, ALL-MODELS, and ALL-CMODELS (see Sections 3–5) for $\mathcal{C}_{EH}$ in place of $\mathcal{C}_H$. That is, the input sets M_i, $i = 1, 2, \ldots, l$ are the sets of characteristic models of extended Horn theories $\Sigma_1, \ldots, \Sigma_l$.

Since the class $\mathcal{C}_{EH}$ is a natural and modest extension to Horn theories, we could expect that the positive results from the previous sections carry over to it. Unfortunately, this is not the case. Already problem MODEL, which is solvable in linear time for $\mathcal{C}_H$, is intractable.

Theorem 7.1 *Problem* MODEL *for class $\mathcal{C}_{EH}$ is NP-hard, even if $l = 2$.*

Proof. We reduce the following NP-complete problem [19] to our problem.

Problem EXACT-HITTING-SET
Input: A collection $\mathcal{S} = \{S_1, S_2, \ldots, S_m\}$ of subsets of a finite set $S = \{1, 2, \ldots, r\}$.
Question: Does $\mathcal{S}$ have an exact hitting set, i.e., a subset $H \subseteq S$ such that $|H \cap S_i| = 1$ for all i ?

Without loss of generality, we may assume that $|S_i| = 3$ holds for all i [19]. Set $n = 2m + r$ and let $V = \{1, \ldots, n\}$. Define $Q_1, Q_2 \subseteq \{0,1\}^n$ by

$$\begin{aligned} Q_1 &= \{x^{V \setminus \{m+k, m+l, m+r+i\}} \mid k, l \in S_i, k \neq l\}, \\ Q_2 &= \{x^{\{i, m+h\}} \mid h \in S_i\}; \end{aligned}$$

the models in Q_1, Q_2 can be illustrated as follows.

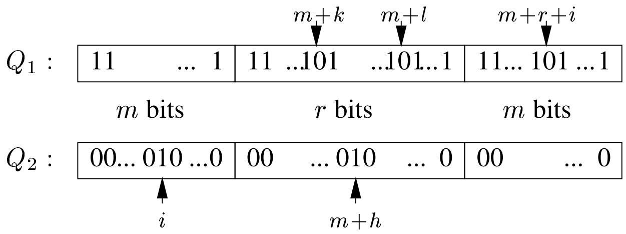

Let $\Sigma_1 = Cl_\wedge(Q_1)$ and $\Sigma_2 = Cl_\vee(Q_2)$, where $Cl_\vee(Q)$ denotes the union closure of Q (dual to the intersection closure). Obviously, $\Sigma_1, \Sigma_2 \in \mathcal{C}_{EH}$, because Σ_1 and Σ_2 are Horn and reverse Horn theories, respectively.

Informally, a model in Q_1 corresponds to the exclusion of the elements k and l from S_i for forming a hitting set H, while a model in Q_2 corresponds to the inclusion of $h \in S_i$ in the hitting set H; they are dual ways of expressing the choice for an element h in S_i. Note that the first m components of the intersection of some models in Q_1 are always 1, and similarly the last m components of the union of some models in Q_2 are always 0. Hence, any model $v \in \Sigma_1 \cap \Sigma_2$ must correspond to the choice of exactly one element from each set S_i, $i = 1, \ldots, m$.

To prove the result, we show (see appendix) that

(i) the set of characteristic models of Σ_i, with respect to class $\mathcal{C}_{EH}$ (i.e., $M_i = \Gamma^{B_{EH}}(\Sigma_i)$) can be obtained from Q_i (and thus, from $\mathcal{S}$) in time polynomial in n and $|Q_i|$, for $i = 1, 2$; and

(ii) $\Sigma_1 \cap \Sigma_2 \neq \emptyset$ if and only if $\mathcal{S}$ has an exact hitting set. □

Corollary 7.1 *For the class $\mathcal{C}_{EH}$, problem* CMODEL *is NP-hard, and there exist no polynomial total time algorithms for the problems* ALL-MODELS *and* ALL-CMODELS, *unless* P=NP.

Proof. NP-hardness of CMODEL is immediate from Theorem 7.1. The latter part can be shown by applying an argument similar to the proof of Theorem 5.2. □

Corollary 7.2 *For the class $\mathcal{C}_{EH}$, both answering a deductive query α and finding an abductive explanation is* co-NP-*hard, even if α is an atom and the set of assumptions A is empty, respectively.* □

8 Conclusion

In this paper, we have considered the problem of taking the intersection $\Sigma = \bigcap_i \Sigma_i$ of Horn theories Σ_i, which are represented by their characteristic models. We found both positive and negative results.

On the positive side, we have shown that deciding consistency and computing some model or characteristic model of Σ are polynomial, and that deductive queries α in CNF to Σ can be answered in polynomial time. More precisely, we presented algorithms which solve model finding, model checking and inference $\Sigma \models C$ of a clause C in $O(n \sum_{i=1}^{l} |M_i|)$ time, i.e., in time linear in the input size. For characteristic model computation, characteristic model checking, and enumerating all models we have described algorithms which work in $O(n^2 \sum_{i=1}^{l} |M_i|)$ time, or in the last case, have this upper bound on the delay between subsequent outputs.

On the negative side, we have shown that computing all characteristic models of Σ is hard, even if the number of models is taken into account. In technical terms, we have shown that there is no polynomial total time algorithm for computing all characteristic models unless $\mathrm{P} = \mathrm{NP}$. The intrinsic difficulty of this problem is further unveiled by our results that also computing an approximation of the set of characteristic models is a hard problem, both for general quantitative notion (a polynomially-sized fraction or superset) and a qualitative notion in terms of the maximal models of a theory. Moreover, we have shown that abductive reasoning from an intersection Σ is intractable; this contrasts with the result in [24], which shows that abductive reasoning from the given characteristic models of Σ is polynomial.

As we have discussed, all these results shed further light on the suitability and computational aspects of the model-based reasoning approach. They tell us that on-line reasoning versus off-line compilation for reasoning from an intersection has to be deliberated, and off-line computation and on-line usage for reasoning may not pay off (e.g., for deductive reasoning). For more insight, we need a study of the typical structure of knowledge bases and query profiles, which we lack to date.

Further issues remain for research. One direction is an extension of our results to other classes of theories. As we have shown, for extended Horn theories, all the main problems which we have considered for Horn theories are intractable. This result indicates that the characteristic models approach is from the computational side not immediately feasible when combining knowledge bases. An investigation which classes of theories besides Horn theories are benign for combination remains to be done.

Another issue concerns a possible combination of the model-based and formula-based approach, in order to have complementary representations of a knowledge base which are suitable for different purposes. It may appear that in such a context, some of the above difficult problems, e.g., computing the characteristic set, is easier. In fact recognizing the characteristic models of Σ is not known to be co-NP-complete, and maybe even polynomial, if the input theories $\Sigma_1, \ldots \Sigma_l$ are represented both by their characteristic models and sets of Horn clauses.

Finally, we comment here that problem MODEL is somewhat related to the *extension problem* for double Horn functions [16], where the extension problem is to establish a Boolean function f that is consistent with a given *partially defined Boolean function* (pdBf) (T, F) (i.e., $f(v) = 1$ (resp. 0) holds for all $v \in T$ (resp., $v \in F$)) [4, 11], and a double Horn extension f is a natural restriction of Horn function. This relationship is by the kind of efficient algorithms for solving the extension problem and problem MODEL, which results from related inherent subproblems. However, no deep semantical relation exists.

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, even if all Σ_i are Horn. Such a theory may be approximated by Horn theories, as described in [25, 26, 9, 20].

A Appendix: Proofs

Theorem 3.1 *Problem* MODEL *can be solved using algorithm* MODEL *in* $O(n^2 \sum_{i=1}^{l} |M_i|)$ *time.*

Proof. We first prove that algorithm MODEL is correct. Let $Q_i^{(j)}$ denote the set Q_i in Step 1 of the j-th iteration, and let $u^{(j)} = \bigvee_{i=1}^{l}(\bigwedge_{w \in Q_i^{(j)}} w)$ denote the model u obtained in Step 3 of the j-th iteration. Consider the first iteration. If $\bigwedge_{w \in Q_1^{(1)}} w = \bigwedge_{w \in Q_2^{(1)}} w = \ldots = \bigwedge_{w \in Q_l^{(1)}} w$, then obviously $v = \bigwedge_{w \in Q_1^{(1)}} w$ is in Σ. Otherwise, we claim that

$$v \in \bigcap_{i=1}^{l} Cl_\wedge(Q_i^{(1)}) \text{ if and only if } v \in \bigcap_{i=1}^{l} Cl_\wedge(Q_i^{(2)}). \tag{1.10}$$

The if-part holds since $Q_i^{(2)} \subseteq Q_i^{(1)}$ holds for all i. For the converse direction, note that any model $v \in \bigcap_{i=1}^{l} Cl_\wedge(Q_i^{(1)})$ satisfies $v \geq u^{(1)}$ by Lemma 3.1. This means that v can be represented by $v = \bigwedge_{w \in Q_1'} w = \bigwedge_{w \in Q_2'} w = \ldots = \bigwedge_{w \in Q_l'} w$ for some $Q_i' \subseteq \{w \in Q_i^{(1)} \mid w \geq u^{(1)}\} = Q_i^{(2)}$. This proves the only-if-part.

Now (1.10) implies that, if $Q_i^{(2)} = \emptyset$ holds for some i, then $\Sigma = \bigcap_{i=1}^{l} Cl_\wedge(Q_i^{(1)})) = \emptyset$; otherwise, in order to find a model $v \in \Sigma$, we only check if there is a model $v \in \bigcap_{i=1}^{l} Cl_\wedge(Q_i^{(2)})$, that is, the problem can be solved by returning to Step 1.

We now iterate the loop of Steps 1-3 for $j = 1, 2, \ldots$. We claim that the iteration finitely terminates. To prove this, we show that $u^{(j)} < u^{(j+1)}$ always holds if algorithm MODEL does not halt in the $(j+1)$-st iteration; as a consequence, it halts after at most $n+1$ iterations.

Since the sets $Q_i^{(j)}$ are monotone nonincreasing with respect to j, $u^{(j)} \leq u^{(j+1)}$ always holds. Let us assume that $u^{(j)} = u^{(j+1)}$ holds for some j. Then, by the definition of $Q_i^{(j+1)}$,

$$u^{(j)} \;\leq\; \bigwedge_{w \in Q_i^{(j+1)}} w \;\leq\; u^{(j+1)} \tag{1.11}$$

holds for all i. Therefore, $u^{(j)} = u^{(j+1)}$ implies $u^{(j)} = \bigwedge_{w \in Q_1^{(j+1)}} w = \bigwedge_{w \in Q_2^{(j+1)}} w = \ldots = \bigwedge_{w \in Q_l^{(j+1)}} w$, and hence MODEL halts in Step 2 of the $(j+1)$-st iteration. This proves our claim.

Finally, since each iteration can be obviously carried out in $O(n \sum_{i=1}^{l} |Q_i^{(j)}|) = O(n \sum_{i=1}^{l} |M_i|)$ time, Algorithm MODEL requires $O(n^2 \sum_{i=1}^{l} |M_i|)$ time in total. □

Corollary 3.1 *Given the characteristic sets $C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, \ldots, l$, 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 $\Sigma = \emptyset$.*

Proof. Define $Q_i^{(j)}$ as in the proof of Theorem 3.1. Let us assume that algorithm MODEL outputs some model v^* in Step 2 of the k-th iteration. Then, by extending (1.10) to $j = 1, 2, \ldots k-1$, we have

$$v \in \bigcap_{i=1}^{l} Cl_\wedge(Q_i^{(1)}) \, (= \Sigma) \iff v \in \bigcap_{i=1}^{l} Cl_\wedge(Q_i^{(2)}) \iff \ldots \iff v \in \bigcap_{i=1}^{l} Cl_\wedge(Q_i^{(k)}). \tag{1.12}$$

Thus $\Sigma = \bigcap_{i=1}^{l} Cl_\wedge(Q_i^{(k)})$ holds. It follows from the definition of v^* that v^* is the unique minimal model in $\bigcap_{i=1}^{l} Cl_\wedge(Q_i^{(k)})$, and thus the least model of Σ. □

Theorem 3.2 *Algorithm* MODEL+ *solves problem* MODEL *in $O(n\sum_{i=1}^{l}|M_i|)$ time, i.e., in linear time.*

Proof. Algorithm MODEL+ is similar to MODEL. Its correctness comes from the following observation. By Lemma 3.1, if $u \in \Sigma$, then $u \geq \bigwedge_{w\in M_i} w$ holds for all $i = 1, 2, \ldots, l$. This implies that if all models w in an M_i satisfy $w_k = 1$ for some k, then $u_k = 1$ must hold. Hence, to compute a model $u \in \Sigma$, we first initialize $u = (00\ldots 0)$, and, for each component k satisfying the above argument, update $u_k := 1$ and remove all models w with $w_k = 0$ from all M_i until either (i) no new k exists or (ii) $M_i = \emptyset$ holds for some i. In case of (i), the current u satisfies $u \in \Sigma$; otherwise, no $u \in \Sigma$ exists. This, combined with the fact that buckets and counters are maintained properly, shows the correctness of MODEL+.

For the time complexity, observe that Step 0 (setting up the data structure) can be done in $O(n\sum_{i=1}^{l}|M_i|)$ time, since each bit of the input can be incorporated into the structures in constant time. The number of iterations of Steps 1 and 2 is at most n, since the numbers of 1 in v strictly increases at each iteration. Thus in total, Step 1 and the maintenance of B in Step 2 require $O(n\sum_{i=1}^{l}|M_i|)$ time, respectively. Furthermore, the n iterations of Step 2 (other than the maintenance of B), can be executed in $O(n\sum_{i=1}^{l}|M_i|)$ time. This is because each component j of any model w is referred only once, each pointer from as well as to a list $L_{h,j}$ is immediately removed after the first reference, and each removal of an entry to $L_{h,j}$ induces only a constant number of counter maintenance steps. Consequently, the overall running time of H-MODEL+ is $O(n\sum_{i=1}^{l}|M_i|)$. □

Theorem 4.1 *Problem* CMODEL *can be solved using algorithm CMODEL in $O(n^2\sum_{i=1}^{l}|M_i|)$ time.*

Proof. To establish the correctness of CMODEL, it remains from the discussion at the beginning of this section to verify Lemma 4.1.
Proof (of Lemma 4.1). We assume that (4.2) holds and $u \notin C^*(\Sigma)$, and derive a contradiction. Then, there exists a model $u' \in \Sigma$ such that $u' > u$ (since $u \notin C^*(\Sigma)$ implies that $u = \bigwedge_{u'\in S} u'$ holds for some $S \subseteq C^*(\Sigma)$, and hence any model u' in S satisfies $u' > u$). Consequently, $u' \in \bigcap_{i=1}^{l} Cl_\wedge(P_{ij})$ must hold for every component j such that $u'_j = 1$ and $u_j = 0$. Since (4.2) is true for u, holds for all j with $u_j = 0$, we then can conclude that there is no such u'; it follows $u \in C^*(\Sigma)$, which is a contradiction. This proves the lemma. ◇

It remains to prove the bound on the time complexity. Step 1 can be done in $O(n\sum_{i=1}^{l}|M_i|)$ time by using algorithm MODEL+ $O(n\sum_{i=1}^{l}|M_i|)$ (Corollary 3.2). In Step 2, for each j, both constructing P_{ij} and updating Q_i for all i can obviously be done in $O(n\sum_{i=1}^{l}|M_i|)$ time. Similarly to Step 1, checking whether $\bigcap_{i=1}^{l} Cl_\wedge(P_{ij}) \neq \emptyset$ and output of some $w' \in \bigcap_{i=1}^{l} Cl_\wedge(P_{ij})$ (if it is not empty) can be done in $O(n\sum_{i=1}^{l}|M_i|)$ time. Thus, the entire Step 2 can be executed in $O(n^2\sum_{i=1}^{l}|M_i|)$ time. In total, $O(n^2\sum_{i=1}^{l}|M_i|)$ time is required. □

Theorem 4.2 *Given the characteristic sets $C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1\ldots, l$, and a model $v \in \Sigma = \bigcap_{i=1}^{l}\Sigma_i$, checking if $v \in C^*(\Sigma)$ is possible using algorithm* CHECK-CMODEL *in $O(n^2\sum_{i=1}^{l}|M_i|)$ time.*

Proof. Note that all models $w^{(j)} \in S$ in algorithm CHECK-CMODEL satisfy $w^{(j)} \geq v$. Thus, by Lemma 4.2, showing

$$S \;\supseteq\; \min(\Sigma_v) \tag{1.13}$$

proves the correctness of CHECK-CMODEL. For every $u \in \min(\Sigma_v)$, there is a component j such that $u_j = 1$ and $v_j = 0$. For such a j, let $Q_i := \{w \in M_i \mid w \geq v, w_j = 1\}$, $i = 1, 2, \ldots, l$. Then $u \in \bigcap_{i=1}^{l} Cl_{\wedge}(Q_i)$ holds. Since $\bigcap_{i=1}^{l} Cl_{\wedge}(Q_i)$ is Horn theory, it has the unique minimal model $w^{(j)}$. However, $u \in \min(\Sigma_v)$ implies $u = w^{(j)}$, and hence (1.13) follows.

For the time complexity of CHECK-CMODEL, Step 0 is possible in constant time. The inner **for** loop in Step 2 is feasible in $O(n|\sum_{i=1}^{l}|M_i|)$ time, and the **if** statement also in $O(n|\sum_{i=1}^{l}|M_i|)$ by virtue of Corollary 3.2. Hence, Step 2 is possible in $O(n^2|\sum_{i=1}^{l}|M_i|)$ time. Step 3 can be done $O(n|\sum_{i=1}^{l}|M_i|)$ time. Hence, algorithm CHECK-CMODEL runs in $O(n^2|\sum_{i=1}^{l}|M_i|)$ time. □

Proposition 5.1 *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. (continued) It remains to show that $S = C^*(\Sigma)$ (5.4) holds.

We first show

$$S \;=\; \max(\Sigma)\ (\subseteq C^*(\Sigma)). \tag{1.14}$$

It is easy to see that $S \subseteq \Sigma$. Assume that there is a model $x^B \in \Sigma$ such that $B \subseteq V_1 \cup V_2$ and $j, n+j \in B$ for some $j \in V_1$. Then, by $j, n+j \in B$ and $x^B \in \Sigma_1$, we have $3n+j \in B$. However, this is a contradiction to (5.3). Hence

$$\{j, n+j\} \not\subseteq B \tag{1.15}$$

holds, which implies the maximality of all models in S, i.e., (1.14). For a non-maximal model $x^B \in \Sigma \setminus S$, we can verify from (5.3) and (1.15) that

$$v \;=\; \bigwedge_{x^B \in S : x^B \geq v} x^B \tag{1.16}$$

holds; i.e., $v \notin C^*(\Sigma)$. This proves our claim (5.4). □

Theorem 5.1 *Problem* ADD-CMODEL *is* NP*-complete, and* NP*-hardness holds even if $l = 2$ is fixed.*

Proof. (continued) Clearly, all models in $C^*(\Sigma_i)$ are maximal; hence, there exist Horn theories Σ_i with the defined characteristic models.

To show that the reduction is appropriate, we will first prove the following containments:

$$\begin{aligned} S &\subseteq \Sigma && (1.17)\\ B &\subseteq V_L \text{ holds for all } x^B \in \Sigma && (1.18)\\ S_1 &\subseteq \max(\Sigma)\ (\subseteq C^*(\Sigma)) && (1.19)\\ S_2 &\subseteq C^*(\Sigma) && (1.20)\\ S &\subseteq C^*(\Sigma). && (1.21) \end{aligned}$$

This shows that (5.5), (5.6) and (5.7) in fact give a legal instance of our problem.

(1.17): Consider $x^B = x^{V_L \setminus \{k, \overline{k}, q\}} (\in S)$, where $q = k$ or $\overline{k}$ is also allowed. By the assumption on Φ, every literal q appears in some clause C_j. Thus $x^{(V_1 \setminus \{n+j\}) \cup (V_L \setminus \{q\})} \in C^*(\Sigma_1)$ holds for some j. This, combined

with $x^{(V_L\setminus\{k,\overline{k}\})\cup V_2} \in C^*(\Sigma_1)$, implies $x^B = x^{(V_1\setminus\{n+j\})\cup(V_L\setminus\{q\})} \bigwedge x^{(V_L\setminus\{k,\overline{k}\})\cup V_2} \in \Sigma_1$. Similarly, we can show $x^B \in \Sigma_2$. Hence (1.17) holds.
(1.18): Since any $x^{B_1} \in \Sigma_1$ satisfies either $V_2 \subseteq B_1$ or $V_2 \cap B_1 = \emptyset$, and no $x^{B_2} \in \Sigma_2$ satisfies $V_2 \subseteq B_2$, we have $V_2 \cap B = \emptyset$ for all $x^B \in \Sigma$. Symmetrically, $V_1 \cap B = \emptyset$ holds for all $x^B \in \Sigma$. Hence (1.18) holds for all $x^B \in \Sigma$.
(1.19): Let $x^B = x^{V_L\setminus\{k,\overline{k}\}} (\in S_1)$. If $x^B \notin \max(\Sigma)$, then, by (1.17) and (1.18), some models in $\{x^{V_L}$, $x^{V_L\setminus\{k\}}$, $x^{V_L\setminus\{\overline{k}\}}\}$ are in Σ. Since no $x^{B_1} \in C^*(\Sigma_1)$ satisfies $B_1 \supseteq V_L$, we have $x^{V_L} \notin \Sigma$. Furthermore, $x^{V_L\setminus\{q\}} \in \Sigma_1$ for $q = k$ or $\overline{k}$ is possible only if

$$x^{V_L\setminus\{q\}} = \bigwedge_{n+j\in V_1} x^{(V_1\setminus\{n+j\})\cup(V_L\setminus\{q\})} \tag{1.22}$$

holds. However, this is impossible by the assumption on Φ that no literal q in L appears in all clauses C_j.
(1.20): For every $v = x^{V_L\setminus\{k,\overline{k},q\}} \in S_2$, there is exactly one $w = x^{V_L\setminus\{k,\overline{k}\}} \in S$ such that $w > v$. Thus, if v can be represented as the intersection of models in $C^*(\Sigma)$, then at least one of the models in

$$\{x^{V_L\setminus\{k,q\}}, x^{V_L\setminus\{\overline{k},q\}}, x^{V_L\setminus\{q\}}\}$$

is contained in $C^*(\Sigma) \setminus S$. However, we will show below (in the proof of (c) $\Rightarrow$ (b)) that, if such a model exists in Σ, then Φ becomes $\top$ by fixing some two atoms in Φ, which contradicts the assumption (ii) on Φ. Therefore, (1.20) holds.
(1.21): Immediate from (1.19) and (1.20). $\diamond$

Clearly $C^*(\Sigma_1)$, $C^*(\Sigma_2)$ and S can be constructed in polynomial time from Φ. Hence, to complete the proof, it remains to show that (a) $C^*(\Sigma) \setminus S \neq \emptyset$ holds if and only if (b) Φ is satisfiable.

It is easy to show that any model u with $u \leq w$ for some $w \in S$ is in $Cl_\wedge(S)$. Thus, (a) is equivalent to the existence of a model $x^B \in \Sigma$ such that $x^B \not\leq w$ holds for all $w \in S$. As a consequence, (a) is also equivalent to (c) the existence of a model $x^B \in \Sigma$ satisfying either $k \in B$ or $\overline{k} \in B$ (or both) for all $k \in V_L$. To prove the equivalence of (a) and (b), we show the equivalence of conditions (b) and (c).
(c) $\Rightarrow$ (b): By $x^B \in \Sigma_1$ and (1.18), x^B can be represented by

$$x^B = \bigwedge_{n+j\in V_1} x^{(V_1\setminus\{n+j\})\cup(V_L\setminus\{q_j\})},$$

where each $q_j \in C_j$ satisfies $q_j \in V_L \setminus B$. Since at least one of $k, \overline{k}$ is contained in B, we can conclude that Φ is satisfiable; a model v such that $\Phi(v) = 1$ can be constructed by fixing $v_k = 1$ if $k \in V_L \setminus B$, 0 if $\overline{k} \in V_L \setminus B$, and 0 or 1 arbitrarily if $k, \overline{k} \notin V_L \setminus B$.
(b) $\Rightarrow$ (c): For a model v with $\Phi(v) = 1$, let

$$V_L \setminus B = \{k \mid v_k = 1\} \cup \{\overline{k} \mid v_k = 0\}.$$

This means that, for each C_j, there is a component $q_j \in C_j \cap (V_L \setminus B)$. Furthermore, since

$$\begin{aligned} x^B &= \textstyle\bigwedge_{n+j\in V_1} \bigwedge_{q_j\in C_j\cap(V_L\setminus B)} x^{(V_1\setminus\{n+j\})\cup(V_L\setminus\{q_j\})} && (\in \Sigma_1) \\ &= \textstyle\bigwedge_{n+m+j\in V_2} \bigwedge_{q_j\in C_j\cap(V_L\setminus B)} x^{(V_2\setminus\{n+m+j\})\cup(V_L\setminus\{q_j\})} && (\in \Sigma_2) \end{aligned}$$

holds, we have a model $x^B \in \Sigma_1 \cap \Sigma_2\ (= \Sigma)$. This completes the proof. $\square$

Theorem 5.4 *Let $p(\cdot)$, $q(\cdot)$ be any polynomials. Then, there is no polynomial total time algorithm $\mathcal{A}$ for computing, given the characteristic sets $M_i = C^*(\Sigma_i)$ of Horn theories $\Sigma_i \subseteq \{0,1\}^n$, $i = 1, 2 \ldots, l$, a set of models $N \subseteq \{0,1\}^n$ such that (i) $|C^*(\Sigma)| \leq q(|N \cap C^*(\Sigma)|)$ and (ii) $|N| \leq p(|C^*(\Sigma)|)$, unless* P = NP. *This holds if $l = 2$ is fixed.*

Proof. We prove this result by an extension to the proof Theorem 5.1 and applying an argument similar as in the proof of Theorem 5.2.

Recall that we have shown in Theorem 5.1 that problem ADD-CMODEL, in the proof, we have described the construction of characteristic sets $C^*(\Sigma_1)$, $C^*(\Sigma_2)$ and a set of models $S \subseteq C^*(\Sigma)$, where $\Sigma = \Sigma_1 \cap \Sigma_2$, from a restricted CNF formula Φ such that $S \neq C^*(\Sigma)$ holds if and and only if Φ is satisfiable. The restrictions on Φ were: (i) Every literal in L appears in Φ, but no literal appears in all clauses; and (ii) Φ does not become a tautology by fixing the truth value of any two atoms x_i and x_j.

Without loss of generality, we may replace (i) by the stronger condition (i'): for each atom x_i, the clause $x_i \vee \overline{x}_i$ occurs in Φ, and require in addition that (iii) if Φ is satisfiable, then it has exponentially many models in the size $|\Phi|$ of Φ; the latter can be easily achieved by adding to Φ sufficiently many clauses $y_i \vee \overline{y}_i$, where the y_i are fresh atoms.

For a formula Φ satisfying (i'), (ii) and (iii), it follows from the construction that the characteristic models $v \in C^*(\Sigma) \setminus S$ correspond 1-1 to the models of Φ. Hence, it follows that Φ is satisfiable, if and only if $C^*(\Sigma)$ is exponential in $|\Phi|$, and that Φ is unsatisfiable, if and only if $S = C^*(\Sigma)$, which is polynomial in $|\Phi|$.

Suppose then an algorithm $\mathcal{A}$ as hypothesized exists, whose running time is bounded by a polynomial $r(I, O)$, where I and O are the input and output length, respectively. We use $\mathcal{A}$ to solve ADD-CMODEL in polynomial time as follows. We run $\mathcal{A}$ on $\Sigma_1, \ldots, \Sigma_l$ for at most $r(I, q(p(|S|)))$ many steps; this is the maximum running time if $C^*(\Sigma) = S$ holds. Since $|C^*(\Sigma) \setminus S|$ is exponential in $|S|$ if $S \neq C^*(\Sigma)$, it follows that $S = C^*(\Sigma)$, if $\mathcal{A}$ halts within this time, and that $S \neq C^*(\Sigma)$, if $\mathcal{A}$ does not. Consequently, ADD-CMODEL can be decided in polynomial time, which implies P=NP; the result follows. □

Theorem 5.6 *Algorithm* ALL-MODELS *is a polynomial delay algorithm for problem* ALL-MODELS, *where the delay is $O(n^2 \sum_{i=1}^{l} |M_i|)$, i.e., number of atoms times input length.*

Proof. The correctness of algorithm ALL-MODELS follows from that fact that it is an instance of the general enumeration scheme described in [12]; we omit the details.

For the time complexity, we note that by Corollary 3.1, MODEL+ finds a model of Σ within time $O(n \sum_{i=1}^{l} |M_i|)$. Furthermore, until the first successful call of PART-MODEL and between two successful calls of PART-MODEL, at most $n - 1$ failing calls of PART-MODEL may occur; since Lemma 6.1 implies that the run time of PART-MODEL is $O(n \sum_{i=1}^{l} |M_i|)$, it follows that the delay between consecutive outputs is bounded by $O(n^2 \sum_{i=1}^{l} |M_i|)$. Finally, at most $n - 1$ failing calls of PART-MODEL may occur until the algorithm halts, and hence it stops within time $O(n^2 \sum_{i=1}^{l} |M_i|)$ after the last output.

Consequently, ALL-MODELS outputs the models in Σ with $O(n^2 \sum_{i=1}^{l} |M_i|)$ delay. □

Theorem 6.3 *Given the characteristic sets $M_i = C^*(\Sigma_i) \subseteq \{0,1\}^n$ of Horn theories Σ_i, $i = 1, \ldots, l$, an assumption set $A \subseteq \{x_1, \ldots, x_n\}$, and an atom q from $x_1, \ldots, x_n$, deciding whether q has an explanation from $\Sigma = \bigcap_{i=1}^{l} \Sigma_i$ and A is* NP*-complete.*

Proof. (continued) We claim that q has an explanation from Σ and A if and only if Φ is satisfiable.

Prior to a proof, we first observe the following useful lemma.

Lemma A.1 *A letter q has an explanation from a Horn theory Σ and assumptions A, if and only if there exists a model v in Σ such that $v \models q$ and $\Sigma \models E \supset q$, where E is the set (seen as conjunction) of all literals ℓ over A such that $v \models \ell$.*

Proof (of Lemma A.1). The if direction is trivial; for the only-if direction, suppose E' is an explanation. Then, there exists a model v in Σ such that $v \models E' \wedge q$. Let E as described; then, since $E' \subseteq E$ and $\widehat{\Sigma} \cup E' \models q$, we have $\widehat{\Sigma} \cup E \models q$, and thus $\Sigma \models E \supset q$. $\diamond$

To prove the only-if direction of the claim, suppose an explanation E exists. We may assume that E has the form as in Lemma A.1 for some model $v \in \Sigma'$. Then, since component 0 of v has value 1, v must be the intersection of vectors from $T'_{1,1}$. Moreover, this intersection must correspond to the choice of a literal from each clause, such that no two opposite literals are selected, i.e., $v = x^B$ such that $B \cap \{k, \overline{k}\} \neq \emptyset$, for all $k \in V_L$. For, otherwise for some model $v^{(k)} = x^{V_L \setminus \{k, \overline{k}\}} \in \Sigma'$, we would have that $w = v \bigwedge v^{(k)}$ would satisfy $w \models E$ but $w \not\models q$, which contradicts that E is an explanation. (From v, we obtain a model of formula Φ as in the proof of Theorem 5.1).

For the if-direction, suppose Φ is satisfiable. Then, from any model of Φ, we construct similar as in the proof of Theorem 5.1 a model v in Σ' which is the intersection of models from $T_{1,1}$ and has no two components $k, \overline{k}$ set to 0, for any $k \in V_L$; observe that v has value 1 at component 0. Let E be as in Lemma A.1; then, E is an explanation for q. Indeed, any model $w \in \Sigma$ which has value 0 at component 0, i.e., $w \models \neg q$, must have value 0 at some components $k, \overline{k}$ where $k \in V_L$. It follows that $w \models \neg E$, and hence clearly $\Sigma \models E \supset q$. Thus, by Lemma A.1, E is an explanation of q. This proves the claim and the result. $\square$

Theorem 7.1 *Problem* MODEL *for class $\mathcal{C}_{EH}$ is NP-hard, even if $l = 2$.*

Proof. (continued) (i): Let us consider M_1. By (7.7), we have

$$M_1 \quad = \quad \Gamma^{B_H}(\Sigma_1) \cup \Gamma^{B_{RH}}(\Sigma_1).$$

Since $\max(Q_1) = Q_1$ and $\Sigma_1 = Cl_\wedge(Q_1)$, we have $C^*(\Sigma_1) = Q_1$. Thus, by (7.8) we have $\Gamma^{B_H}(\Sigma_1) = Q_1$.

Concerning $\Gamma^{B_{RH}}(\Sigma_1)$, let $z = \bigwedge_{w \in Q_1} w$ and $z(b) = \bigwedge_{w \in Q_1 : w \geq b} w$ for any b with $\| b \| = 1$. Then, since $z \leq_b v$ (resp., $z(b) \leq_b v$) holds for all $v \in \Sigma_1$ with $v_j = 0$ (resp., $v_j = 1$), where j denotes an index such that $b_j = 1$, it follows that $\min_{(00\ldots0)}(\Sigma_1) = z$ and $\min_b(\Sigma_1) \subseteq \{z, z_b\}$. This implies that also $\Gamma^{B_{RH}}(\Sigma_1)$ is computable from Q_1 in polynomial time. Consequently, M_1 is computable from Q_1 in polynomial time. The set M_2 can be obtained in a similar manner; this proves (i).

(ii): Any model $v \in \Sigma_1 \cap \Sigma_2$ must satisfy

$$v_j = 1, \qquad \text{for all } \; j = 1, 2, \ldots, m \tag{1.23}$$

$$v_j = 0, \qquad \text{for all } \; j = m + r + 1, m + r + 2, \ldots, m + r + m \tag{1.24}$$

To prove the only-if-part of (ii), assume that some model $v \in \Sigma_1 \cap \Sigma_2$ exists. Then, $v = \bigwedge_{w \in Q'_1} w = \bigvee_{w \in Q'_2} w$ holds for some nonempty sets $Q'_1 \subseteq Q_1$ and $Q'_2 \subseteq Q_2$. We show that $H = \{h \,|\, x^{\{i, m+h\}} \in Q'_2\}$ $(\subseteq S)$ forms an exact hitting set of $\mathcal{S}$. By (1.23), for each $i = 1, 2, \ldots, m$ there is an h such that $x^{\{i, m+h\}} \in Q'_2$. This means that H satisfies

$$|H \cap S_i| \geq 1 \tag{1.25}$$

for all i. Furthermore, by (1.24), there are for each i elements k and l such that $x^{V\setminus\{m+k,m+l,m+n+i\}} \in Q'_1$. which implies $k, l \notin H$. Thus we have

$$|H \cap S_i| \leq 1 \tag{1.26}$$

for all i. By (1.25) and (1.26), we conclude that H is an exact hitting set of $\mathcal{S}$.

For the if-direction, assume that H is an exact hitting set of $\mathcal{S}$. Then define

$$\begin{aligned} Q'_1 &= \{x^{V\setminus\{m+k_i,m+l_i,m+r+i\}} \mid \{k_i, l_i\} = S_i \setminus H, i = 1, 2, \ldots, m\} \\ Q'_2 &= \{x^{\{i,m+h_i\}} \mid \{h_i\} = S_i \cap H, i = 1, 2, \ldots, m\}. \end{aligned}$$

We can see that $\bigwedge_{w\in Q'_1} w = \bigvee_{w\in Q'_2} w\,(= v)$ holds and hence $v \in \Sigma_1 \cap \Sigma_2$; this proves (ii) and the theorem. $\square$

B Appendix: Computing a Maximum Model

Theorem B.1 *Given sets of characteristic models $M_i \subseteq \{0,1\}^n$, representing Horn theories Σ_i, $i = 1, 2, \ldots, l$, the problem of finding a maximum model $v \in \Sigma\,(= \bigcap_{i=1}^{l} \Sigma_i)$ (i.e., with the maximum $\| v \|$) is* NP-*hard, even if $l = 2$.*

Proof. We shall reduce problem SET-COVER to our problem, where SET-COVER is the following integer programming problem and is known to be NP-hard [19].

$$\begin{array}{ll} \textbf{minimize} & \sum_{j=1}^{n} y_j \\ \textbf{subject to} & Ay \geq 1 \\ & y_j \in \{0,1\}, \quad j = 1, 2, \ldots, n, \end{array}$$

where $A = [a_{ij}]$ is an $m \times n$ 0-1 matrix, $y = (y_1, y_2, \ldots, y_n)^T$ and $1 = (1, 1, \ldots, 1)^T$.

Let $V_1 = \{1, 2, \ldots n\}$, $V_2 = \{n+1, n+2, \ldots n+m\}$ and $V = V_1 \cup V_2$, and let $A_i = \{j \mid a_{ij} = 1\}$. Given an instance of SET-COVER, i.e., a matrix A, we construct two sets $M_1, M_2 \subseteq \{0,1\}^{n+m}$ of characteristic models of Horn theories as follows.

$$M_1 = \{x^{V_1\setminus\{j\}} \mid j = 1, 2, \ldots, n\} \tag{2.27}$$

$$M_2 = \{x^{V\setminus\{n+i,j_i\}} \mid i = 1, 2, \ldots, m, j_i \in A_i\}, \tag{2.28}$$

where x^B denotes the model such that $x^B_i = 1$ if $i \in B$ and 0 if $i \notin B$. It is easy to see that M_1 and M_2 are in fact sets of characteristic models of Horn theories Σ_1 and Σ_2, respectively, i.e., $M_1 = C^*(\Sigma_1)$ and $M_2 = C^*(\Sigma_2)$ hold for some Horn theories Σ_1 and Σ_2. This is because no pair of models $v, w \in M_i$ satisfies $v \leq w$.

Let us first consider the models in Σ_1. Since all models x^B in M_1 satisfy $B \subseteq V_1$, all models $x^B \in \Sigma_1$ also have the same property. For any set $B \subset V_1$, where $\subset$ denotes the proper inclusion, $x^B = \bigwedge_{j\in V_1\setminus B} x^{V_1\setminus\{j\}}$ clearly holds, and hence $x^B \in \Sigma_1$. Furthermore, since $x^{V_1} \notin \Sigma_1$ is clear, we have

$$\Sigma_1 = \{x^B \mid B \subset V_1\}. \tag{2.29}$$

Next we consider those models v in Σ_2 which can be written as $v = x^B$ for some $B \subset V_1$. By the definition of M_2, we can see that $x^{V_1\setminus B} \in \Sigma_2$ holds if and only if the y defined by $y_j = 1$ if $j \in V_1 \setminus B$

and 0 otherwise is a feasible solution of the instance of SET-COVER. We call such a $V_1 \setminus B$ *feasible*. Thus, $\Sigma\,(= \Sigma_1 \cap \Sigma_2)$ can be written as

$$\Sigma \;=\; \{x^B \mid B \subset V_1, V_1 \setminus B \text{ is feasible}\}. \tag{2.30}$$

We can then conclude that x^B is a maximum model in Σ if and only if the corresponding y is an optimal solution of the instance of SET-COVER. $\Box$

References

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

[2] C. Baral, S. Kraus, J. Minker, and V. S. Subrahmanian. Combining knowledge bases consisting of first order theories. In M. Ras, Z.W., Zemankova (eds), *Proc. 6th International Symposium on Methodologies for Intelligent Systems (ISMIS '91)*, LNAI 542, pp. 92–101, 1991.

[3] C. Beeri, M. Down, R. Fagin, and R. Statman. On the Structure of Armstrong Relations for Functional Dependencies. *Journal of the ACM*, 31(1):30–46, January 1984.

[4] E. Boros, T. Ibaraki, and K. Makino. Error-free and Best-fit Extensions of Partially Defined Boolean Functions. *Information and Computation*, 140:254–283, 1998.

[5] G. Brewka, J. Dix, and K. Konolige. *Nonmonotonic Reasoning – An Overview*. CSLI Lecture Notes 73. CSLI Publications, Stanford University, 1997.

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

[7] R. Burke and K. J. Hammond. Combining Databases and Knowledge Bases for Assisted Browsing. In *AAAI Spring Symposium on Information Gathering*, 1995.

[8] M. Cadoli. The Complexity of Model Checking for Circumscriptive Formulae. *Information Processing Letters*, 44:113–118, 1992.

[9] M. Cadoli. Semantical and Computational Aspects of Horn Approximations. In *Proc. IJCAI-93*, pp. 39–44, 1993.

[10] M. Cadoli and M. Schaerf. Tractable Reasoning via Approximation. *Artificial Intelligence*, 74(2):249–310, 1995.

[11] Y. Crama, P. Hammer, and T. Ibaraki. Cause-Effect Relationships and Partially Defined Boolean Functions. *Annals of Operations Research*, 16:299–326, 1988.

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

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

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

[15] T. Eiter and G. Gottlob. The Complexity of Logic-Based Abduction. *Journal of the ACM*, 42(1):3–42, 1995.

[16] T. Eiter, K. Makino, and T. Ibaraki. Double Horn Functions. *Information and Computation*, 142, 1998.

[17] D. Gabbay, C. Hogger, and J. Robinson, editors. *Handbook of Logic in Artificial Intelligence and Logic Programming*. Clarendon Press, Oxford; New York, 1993.

[18] P. Gärdenfors. *Knowledge in Flux*. Bradford Books, MIT Press, 1988.

[19] M. Garey and D. S. Johnson. *Computers and Intractability – A Guide to the Theory of NP-Completeness*. W. H. Freeman, New York, 1979.

[20] G. Gogic, C. Papadimitriou, and M. Sideri. Incremental Recompilation of Knowledge. *Journal of Artificial Intelligence Research*, 8:23–37, 1998.

[21] A. Itai and J. A. Makowsky. Unification as a Complexity Measure for Logic Programming. *Journal of Logic Programming*, 4:105–117, 1987.

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

[23] H. Katsuno and A. O. Mendelzon. Propositional Knowledge Base Revision and Minimal Change. *Artificial Intelligence*, 52:253–294, 1991.

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

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

[26] D. Kavvadias, C. Papadimitriou, and M. Sideri. On Horn Envelopes and Hypergraph Transversals. In W. Ng (ed), *Proc. 4th International Symposium on Algorithms and Computation (ISAAC-93)*, LNCS 762, pp. 399–405, 1993.

[27] R. Khardon. Translating between Horn Representations and their Characteristic Models. *Journal of Artificial Intelligence Research*, 3:349–372, 1995.

[28] R. Khardon, H. Mannila, and D. Roth. Reasoning with Examples: Propositional Formulae and Database Dependencies. manuscript, Acta Informatica.

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

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

[31] H. Levesque. Making Believers out of computers. *Artificial Intelligence*, 30:81–108, 1986.

[32] P. Liberatore and M. Schaerf. The Complexity of Model Checking for Belief Revision and Update. In *Proc. AAAI-96*, pp. 556–561, 1996.

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

[34] J. McKinsey. The Decision Problem for Some Classes of Sentences Without Quantifiers. *Journal of Symbolic Logic*, 8:61–76, 1943.

[35] C. S. Peirce. Abduction and induction. In J. Buchler, editor, *Philosophical Writings of Peirce*, chapter 11. Dover, New York, 1955.

[36] P. Z. Revesz. On the Semantics of Theory Change: Arbitration between Old and New Information. In *Proc. ACM Symposium on Principles of Database Systems (PODS-93)*, pp. 71–79, 1993.

[37] D. Roth. On the Hardness of Approximate Reasoning. *Artificial Intelligence*, 82(1/2):273–302, 1996.

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

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

[40] L. Valiant. The Complexity of Enumeration and Reliability Problems. *SIAM Journal of Computing*, 8:410–421, 1979.

[41] M. Winslett. *Updating Logical Databases*. Cambridge University Press, 1990.