Horn Belief Contraction: Remainders, Envelopes and Complexity[*]

Kira Adaricheva
Yeshiva University

Robert H. Sloan
University of Illinois at Chicago

Balázs Szörényi
Hungarian Academy of Sciences
and University of Szeged
Research Group on Artificial Intelligence

György Turán
University of Illinois at Chicago,
Hungarian Academy of Sciences and University of Szeged
Research Group on Artificial Intelligence

Abstract

Belief change studies how to update knowledge bases used for reasoning. Traditionally belief revision has been based on full propositional logic. However, reasoning with full propositional knowledge bases is computationally hard, whereas reasoning with Horn knowledge bases is fast. In the past several years, there has been considerable work in belief revision theory on developing a theory of belief contraction for knowledge represented in Horn form.

Our main focus here is the *computational complexity* of belief contraction, and, in particular, of various methods and approaches suggested in the literature. This is a natural and important question, especially in connection with one of the primary motivations for considering Horn representation: efficiency.

The problems considered lead to questions about Horn envelopes (or Horn LUBs), introduced earlier in the context of knowledge compilation. This work gives a *syntactic characterization* of the remainders of a Horn belief set with respect to a consequence to be contracted, as the Horn envelopes of the belief set and an elementary conjunction corresponding to a truth assignment satisfying a certain explicitly given formula. This gives an *efficient algorithm* to generate all remainders, each represented by a truth assignment.

On the negative side, examples are given of Horn belief sets and consequences where Horn formulas representing the result of contraction, based either on remainders or on weak remainders, must have *exponential size* for almost all possible choice functions (i.e., different possible choices of partial meet contraction). Therefore using the Horn framework for belief contraction does not by itself give us computational efficiency. Further work is required to explore the possibilities for efficient belief change methods.

1 Introduction

Belief revision attempts to answer the question of how to update a set of beliefs when new information is obtained that may be inconsistent with the current beliefs (Fermé and Hansson 2011; Hansson 1999; Peppas 2008). The standard approach, called the AGM approach after (Alchourrón, Gärdenfors, and Makinson 1985), is to formulate postulates that need to be satisfied by rational agents performing belief revision, and then to characterize all possible operations that satisfy these postulates, with the usual set of postulates being the AGM postulates. Until recently, work on AGM-style belief revision focused on logics at least as rich as full propositional logic, and assumed a language that was closed under the basic operations of propositional logic: negation, disjunction, and conjunction. For example, a central result of AGM for contraction—that a contraction operator obeys the AGM postulates if and only if that contraction operator consists of returning some intersection of remainder sets (i.e., partial meet contraction)—relies on the language being closed under those operations.

Evolving knowledge bases and ontologies appear to be interesting potential application areas for belief revision. These applications require *tractable* knowledge representation formalisms, such as Horn logic or some versions of description logic, which do not contain full propositional logic. Thus it is of interest to develop a belief change theory for these particular logics, and, furthermore, for arbitrary logics in general.

In recent years, there have been a number of papers considering logics that are *not* necessarily closed under the basic operations. As far as we know, (Flouris, Plexousakis, and Antoniou 2004) were the first to look into this question. Flouris et al. wanted to develop a theory of belief revision that would apply to description logic. In the case with closure under the basic propositional logic operators, a contraction operator obeying the AGM postulates always exists. This is not always the case for logics not so closed. Flouris et al. formulate a property called *decomposability* of the logic, and show that decomposability is a necessary and sufficient

condition for the existence of an AGM-compliant belief contraction operator.

Starting in 2008, there has been a flurry of papers considering specifically the case of *belief contraction for Horn logic*, that is, the subset of propositional logic consisting of Horn formulas (Booth et al. 2011; Booth, Meyer, and Varzinczak 2009; Delgrande 2008; Delgrande and Peppas 2011; Delgrande and Wassermann 2010; Fotinopoulos and Papadopoulos 2009; Langlois et al. 2008; Zhuang and Pagnucco 2010a; 2010b; 2011). The problem was actually studied much earlier, independently of the AGM framework (Kleine Büning and Lettmann 1987).

Horn logic is a fragment of (propositional and predicate) logic that is of central importance in AI. Horn clauses express *rules* that are natural and easy to understand for humans. Another main reason for interest in Horn logic is that reasoning in propositional logic is computationally intractable, but reasoning in Horn logic is efficient. Incidentally, Horn logic is one of the six tractable cases of the constraint satisfaction problem in the sense of Schaefer's famous dichotomy theorem (Schaefer 1978). [1]

(Langlois et al. 2008) consider the framework of (Flouris, Plexousakis, and Antoniou 2004) and show that Horn logic is not decomposable. (See also (Ribeiro 2010).) They characterize the Horn belief sets that have an AGM-compliant contraction operator, and give a *polynomial-time algorithm* to compute such a contraction when one exists. Their results use the notion of *complement* introduced by (Flouris, Plexousakis, and Antoniou 2004), which, in standard belief revision terms, is a remainder of a belief set with respect to itself.

The papers (Booth et al. 2011; Booth, Meyer, and Varzinczak 2009; Delgrande 2008; Delgrande and Wassermann 2010; Zhuang and Pagnucco 2010a; 2010b; 2011) ask the important question, "What sort of contraction operator should we use for the full Horn logic?" They consider variations of both the definition of remainder set and of how to combine remainder sets to obtain contraction operators that are defined for *all* Horn knowledge bases K and consequences φ, and give results relating those contraction operators to various sets of postulates. A common feature is that the *recovery postulate*, which is assumed in (Flouris, Plexousakis, and Antoniou 2004), is replaced by other postulates. Computational efficiency issues are not considered in these papers.

Since one major advantage of restricting oneself to a Horn knowledge base is computational efficiency, it is important to understand the computational complexity of contracting or revising for such knowledge bases. Here we give initial results on the efficiency of some of the Horn belief contraction methods from several of the papers discussed in the previous paragraph. Previous, mostly negative, complexity theoretic results on classical belief revision are given in (Eiter and Gottlob 1992; Liberatore 2000; Nebel 1998). These papers contain results on revising Horn belief sets. However, in the classical framework the result

of a contraction or a revision of a Horn belief set is not necessarily Horn. Notice that while one of the motivations for considering Horn logic is to gain in efficiency, for example in reasoning with the belief set, it may be the case that one has to pay a price in the sense that some tasks become *more difficult* due to the restricted nature of the logic. Some of our results show that such a phenomenon does indeed occur.

Results

In this paper we focus on remainder sets and weak remainder sets. Given a Horn belief set K and a Horn formula φ to be contracted, one can consider enlarging the set of truth assignments satisfying K by a single truth assignment falsifying φ. This sometimes produces a remainder set. However, as noted in (Delgrande and Wassermann 2010), not all such truth assignments produce a remainder set, only those which, when intersected componentwise with the truth assignments satisfying K, do not produce any other truth assignments falsifying φ. Weak remainder sets are defined in (Delgrande and Wassermann 2010) similarly to remainder sets, except that now an arbitrary truth assignment falsifying φ can be added to the satisfying truth assignments of K. Both remainders and weak remainders lead to interesting, but different, theories of Horn belief contraction.

As not every truth assignment falsifying φ can be used to build a remainder set of K with respect to φ, the first task is to understand which truth assignments can in fact be added. In Section 3, we give a logical characterization of those truth assignments that do lead to a remainder set, using a generalization of the *body building* formula from (Langlois et al. 2008).[2] This result is essentially different from (Delgrande and Wassermann 2010), as it leads to efficient algorithms. (See Remark 5 for a detailed comparison.) The result gives a characterization of *quasi-closed sets* of a closure system, which appears to be different from those previously given in the literature (Caspard and Monjardet 2003). Using this characterization, an efficient listing algorithm is given to list all such truth assignments. Thus, this algorithm produces all possible remainder sets, represented by their 'generating' truth assignment.

In Sections 4 and 5, we consider the question of finding a Horn formula representation for Horn belief sets produced by the various contraction operators. We construct Horn belief sets and Horn clause consequences to be contracted with the property that for certain contraction operators *every* Horn formula representing the new belief set must be exponentially larger than the original belief set. In fact, this holds for full meet contraction and, in an asymptotic sense, for most maxichoice and most partial meet contractions, both based on remainders and on weak remainders.[3] Our result is based on a new blowup result for computing Horn envelopes (also called Horn LUB's), which is connected to the study of *D-bases of closure systems* (Adaricheva, Nation, and Rand 2011). A related earlier blowup result was given in (Kleine

---

[1] Belief revision for some of the other tractable cases is discussed by Creignou et al. (2012).

[2] In other words, an *efficiently computable syntactic criterion* is given to distinguish remainders from weak remainders.

[3] For our example, the two actually coincide as every weak remainder happens to be a remainder.

Büning and Lettmann 1987). We also give some positive results on cases where such a blowup cannot occur.

As it is indicated by the connection to bases of closure systems, Horn formulas are closely related to concepts in algebra (lattices, closure systems, closure operators and implicational systems (Davey and Priestley 1990)). There is a large body of work on problems related to the ones studied in this paper (Bertet and Monjardet 2010; Burosch, Demetrovics, and Katona 1987; Caspard and Monjardet 2003; Freese 1995), and the study of belief contraction for general logics in (Flouris, Plexousakis, and Antoniou 2004) also uses a lattice theoretic framework. We plan to give an account of these useful connections in the full version of this paper.

2 Preliminaries

We assume a fixed finite set of propositional variables. We use 0 and 1 for representing truth values. The set of truth assignments satisfying (resp., falsifying) a propositional formula ψ is denoted by $T(\psi)$ (resp., $F(\psi)$). For formulas ψ, φ it holds that $\psi \models \varphi$ (i.e., φ is a consequence of ψ) iff $T(\psi) \subseteq T(\varphi)$. For a truth assignment a and a variable x we sometimes write $x(a)$ for the value of the x-component of a.

Truth assignments are partially ordered by the relation $a \leq b$, which holds for $a = (a_1, \ldots, a_n)$ and $b = (b_1, \ldots, b_n)$ iff $a_i \leq b_i$ for every $i = 1, \ldots, n$. We write $a < b$ if $a \leq b$ and $a \neq b$. The (componentwise) intersection of $a = (a_1, \ldots, a_n)$ and $b = (b_1, \ldots, b_n)$ is $a \wedge b = (a_1 \wedge b_1, \ldots, a_n \wedge b_n)$.

The elementary conjunction E_a corresponding to a truth assignment a contains a variable (resp., the negation of a variable) iff the component corresponding to that variable in a is set to 1 (resp., 0). Thus, for example, $E_a = x_1 \wedge \bar{x}_2 \wedge x_3$ for $a = (1, 0, 1)$. Clearly $T(E_a) = \{a\}$.

A clause is a disjunction of literals (unnegated and negated variables). A clause is Horn if it contains at most one unnegated variable, and it is definite if it contains exactly one unnegated variable. (See, e.g., (Crama and Hammer 2011) for background on Horn formulas). A definite Horn clause C is also written as $\text{Body}(C) \rightarrow \text{Head}(C)$, where $\text{Body}(C)$, resp., $\text{Head}(C)$, are the body, resp., the head of the clause. For example, the definite Horn clause $C = \bar{x} \vee \bar{y} \vee z$ can be written as $x, y \rightarrow z$, with $\text{Body}(C) = \{x, y\}$ and $\text{Head}(C) = \{z\}$. A (definite) Horn formula is a conjunction of (definite) Horn clauses.

A clause C is an implicate of a formula ψ iff $\psi \models C$, and it is a prime implicate if none of its subclauses is an implicate. Every prime implicate of a (definite) Horn formula is a (definite) Horn clause. *Forward chaining* is an efficient procedure to decide $\psi \models C$, where ψ is a definite Horn formula and C is a definite Horn clause. It starts by marking all variables in the body of C. While there is a clause in ψ with all its body variables marked, the head variable of that clause is marked as well. Then $\psi \models C$ iff the head of C gets marked. Forward chaining is the basis of efficient satisfiability, equivalence, and inference algorithms for Horn formulas (Kleine Büning and Lettmann 1999).

A Boolean function f can be represented by a Horn formula iff $T(f)$ is closed under intersection (Horn 1951;

McKinsey 1943). Given an arbitrary propositional formula ψ, its *Horn envelope* $\text{Env}(\psi)$ is the conjunction of all Horn implicates of ψ. (In fact it is easily seen that $\text{Env}(\psi)$ depends only on the function that ψ represents, and not the particular formula ψ.) The Horn envelope is also referred to as the Horn LUB (least upper bound) or the Horn closure of ψ (Selman and Kautz 1996). It holds that $T(\text{Env}(\psi))$ is the closure of $T(\psi)$ under intersection.

The closure $\text{Cl}_\psi(S)$ of a set of variables S with respect to a definite Horn formula ψ is the set of all variables that must be true in every truth assignment satisfying ψ and having all variables in S set to 1, in other words

$$\text{Cl}_\psi(S) = \left\{ v \; : \; \psi \models \left(\bigwedge_{x \in S} x \rightarrow v \right) \right\} \; .$$

This is a closure operator on the set of all variables that can be computed by forward chaining. Note that $\text{Cl}_\psi(S)$ depends only on the function represented by ψ and not on the particular representation of ψ.

A *Horn belief set* K is a set of definite Horn clauses closed under implication.[4] As we are working with a fixed finite set of variables, and we may assume without loss of generality that clauses do not contain any repeated literals, belief sets are finite. A finite set of clauses in the belief set can be also thought of as the conjunction of the clauses in the set. For representational and computational purposes we may represent K by a subset of its clauses that imply all the others. Different logically equivalent formulas are considered to represent the *same* belief set. This is different from the *belief base* approach where clauses explicitly represented in the base have a distinguished role, and different logically equivalent representations are considered to be *different* as belief bases.

3 Body building: a characterization of remainders

The main result of this section is Theorem 4, which gives a syntactic characterization of remainders.

If K is a Horn belief set and Horn formula φ is a consequence of K, then an *e-remainder set* or, briefly, a *remainder* of K with respect to φ is a maximal subset $K' \subset K$ not implying φ. Since this paper is devoted exclusively to Horn belief sets, and, as (Delgrande and Wassermann 2010) points out, the general definition of remainder set restricted to the Horn case gives e-remainder, we will just write "remainder" in the rest of this paper. We will denote the set of all remainders of K with respect to φ by $K \downarrow \varphi$.

The following proposition, which follows directly from the definitions, is implicit in (Delgrande and Wassermann 2010).

Proposition 1. $K \downarrow \varphi$ *is equal to*

$$\{\text{Env}(K \vee E_a) \; : \; T(\text{Env}(K \vee E_a)) \cap F(\varphi) = \{a\}\} \; .$$

---

[4]In this paper we restrict our attention to definite Horn belief sets for simplicity. The extension of the results to the general case will be discussed in the final version.

Proposition 1 gives a description of all remainders, but it does not give an efficient algorithm to find any remainders as it does not tell how to find truth assignments a with the required property. In order to provide a constructive description we introduce the following definition.

Definition 2 (Body building formula).

$$K^\varphi = \bigwedge_{C \in \varphi} \bigwedge_{v \notin \mathrm{Cl}_K(\mathrm{Body}(C))} (\mathrm{Body}(C), v \to \mathrm{Head}(C)).$$

Note that, by definition, K^φ is a consequence of φ.

Proposition 3. *The formula K^φ can be computed in polynomial time in the size of the representations of K and φ.*

The definition of K^φ generalizes the notion of a body building formula introduced in (Langlois et al. 2008). The definition in that paper corresponds to K^K in the current notation. Using Definition 2, remainders can be characterized as follows.

Theorem 4. *If K is a Horn belief set and definite Horn formula φ is a consequence of K then*

$$K \downarrow \varphi = \{\mathrm{Env}(K \vee E_a) \,:\, a \in T(K^\varphi) \cap F(\varphi)\}.$$

Remark 5. While Theorem 4 looks quite similar to Proposition 1 in terms of its syntax, the two statements are quite different. Proposition 1 is in terms of Horn envelopes that depend on a, and so it gives no hint on how to find a's for which the condition holds. Furthermore, as we will see, envelopes may be hard to compute. Theorem 4, on the other hand, is in terms of the body building formula which is independent of a and thus can be used to find all suitable a's using standard methods (see Theorem 6). Also, as noted in Proposition 3, the body building formula can be computed efficiently.

Proof of Theorem 4. In order to prove the "⊆" part of the theorem we show that if some $a \in F(\varphi)$ falsifies K^φ then

$$|T(\mathrm{Env}(K \vee E_a)) \setminus T(\varphi)| \geq 2, \tag{1}$$

and thus by Proposition 1 it is not a remainder of K with respect to φ.

If $K^\varphi(a) = 0$ then there is a definite clause C in φ and a variable $v \notin \mathrm{Cl}_K(\mathrm{Body}(C))$ such that a falsifies $\mathrm{Body}(C), v \to \mathrm{Head}(C)$. Thus $\mathrm{Body}(C)(a) = 1$, $v(a) = 1$ and $\mathrm{Head}(C)(a) = 0$.

As $v \notin \mathrm{Cl}_K(\mathrm{Body}(C))$, there is a $b \in T(K)$ such that $\mathrm{Body}(C)(b) = 1$ and $v(b) = 0$. But as $K \models \varphi$ and $b \in T(K)$, it holds that $b \in T(\varphi)$, and thus b must satisfy C. Hence $\mathrm{Head}(C)(b) = 1$.

Now consider the truth assignment $d = a \wedge b$. Claim (1) follows if we show that $d \in T(\mathrm{Env}(K \vee E_a)) \setminus T(\varphi)$ and $d \neq a$.

We know that $\mathrm{Env}(K \vee E_a)$ is closed under intersection and so $b \in T(K)$ implies that $d \in T(\mathrm{Env}(K \vee E_a))$. As $\mathrm{Body}(C)(a) = \mathrm{Body}(C)(b) = 1$ it follows that $\mathrm{Body}(C)(d) = 1$. On the other hand, $\mathrm{Head}(C)(a) = 0$ implies $\mathrm{Head}(C)(d) = 0$. Thus d falsifies clause C and so it falsifies φ as well. Furthermore, $v(b) = 0$ implies $v(d) = 0$, thus from $v(a) = 1$ we get $a \neq d$.

Thus a and d are two different points in $\mathrm{Env}(K \vee a)$ outside $T(\varphi)$ and so $\mathrm{Env}(K \vee a)$ is not a remainder set of K with respect to φ.

This proves the first half of the theorem.

For the "⊇" direction we show that if $\mathrm{Env}(K \vee E_a)$ is not a remainder set for $a \in F(\varphi)$ then a falsifies K^φ.

If $\mathrm{Env}(K \vee E_a)$ is not a remainder set for $a \in F(\varphi)$ then by Proposition 1 there are at least two points satisfying $\mathrm{Env}(K \vee E_a)$ but falsifying φ. The set $T(\mathrm{Env}(K \vee E_a))$ is obtained from $T(K)$ and a by adding all truth assignments of the form $d = a \wedge b$ for every $b \in T(K)$. By our assumption, there is such a $d \neq a$ with $\varphi(d) = 0$ for some $b \in T(K)$. The theorem is proved by using d to find a clause of K^φ falsified by a.

Now $T(K) \subseteq T(\varphi)$ implies $b \in T(\varphi)$ and so $d \neq b$. Thus a and b are incomparable and so there is a variable v such that $v(a) = 1$ and $v(b) = 0$. Let C be a clause of φ falsified by d. Then

$$\mathrm{Body}(C)(d) = 1 \text{ and } \mathrm{Head}(C)(d) = 0 \,. \tag{2}$$

We are done if we show

$$v \notin \mathrm{Cl}_K(\mathrm{Body}(C)) \tag{3}$$

and

$$(\mathrm{Body}(C), v \to \mathrm{Head}(C))(a) = 0 \,. \tag{4}$$

Claim (3) follows by considering the truth assignment b. We have all the ingredients needed: $b \in T(K)$,

$$\mathrm{Body}(C)(b) = 1 \tag{5}$$

(as $\mathrm{Body}(C)(d) = 1$ and $d < b$), and $v(b) = 0$.

We now prove (4). First note that $d < a$ and (2) together imply

$$\mathrm{Body}(C)(a) = 1 \,. \tag{6}$$

Now $\varphi(b) = 1$ implies that $C(b) = 1$ which, combined with (5), gives that $\mathrm{Head}(C)(b) = 1$. Since $\mathrm{Head}(C)(b) = 1$, it must be that $\mathrm{Head}(C)(a) = 0$, as otherwise, using the fact that $d = a \wedge b$, we would get $\mathrm{Head}(C)(d) = 1$, contradicting (2). This, together with (6) and that v was chosen to satisfy $v(a) = 1$, proves (4), completing the proof of the theorem. □

We claim that Theorem 4 can be used to find remainders efficiently, though precisely what we mean by efficiently needs further discussion. A remainder can be obtained by finding a truth assignment $a \in T(K^\varphi) \cap F(\varphi)$. Such a truth assignment can be found by running an efficient Horn satisfiability algorithm on $K^\varphi \wedge \neg C$ for each clause C in φ. Actually, an even stronger statement is true: *all* remainders can be listed efficiently. As the number of remainders can be large, i.e., superpolynomial in the size of the belief set, we have to explain what is meant by efficient listing in general.

A *listing algorithm* (sometimes also called an enumeration algorithm) is an algorithm to produce a list of objects. For instance, a basic task in data mining is to produce a list of potentially interesting association rules in a transaction database, for some specific definition of interestingness. Using this list, the user is supposed to select those rules which are found to be truly interesting. For belief change, such algorithms could be used to produce a list of possible results

of contraction operators, again, letting the user decide which one is preferred. Another possible application is in the experimental study of belief change algorithms, suggested in Section 6 as a topic for further research, where a list of possible contractions could be necessary to compute various statistics.

Different efficiency criteria for listing algorithms are described in (Goldberg 1993). Here we define only one. An algorithm listing a set of objects works with *polynomial delay* if the time spent before outputting the first object and the time spent between outputting two successive objects (and between the final output and termination) is bounded by a polynomial function of the input size.

Theorem 6. *There is a polynomial delay algorithm which, given a Horn belief set K and a consequence φ of K, outputs a list of all truth assignments a such that $\mathrm{Env}(K \vee E_a)$ is in $K \downarrow \varphi$.*

Theorem 6 is a direct consequence of Theorem 4 and standard algorithms for Horn formulas. The algorithm does backtracking for subproblems obtained by restricting variables to constants, and it uses Horn satisfiability to check whether a new subtree contains any remainders.

The algorithm in Theorem 6 produces a list of all remainders, where each remainder is *represented by a truth assignment*, and not by a Horn formula for $\mathrm{Env}(K \vee E_a)$. This begs the question, considered in the next section, whether Horn formulas for $\mathrm{Env}(K \vee E_a)$ can be computed efficiently?

4 Horn envelopes

As Proposition 1 shows, remainders are closely related to Horn envelopes, which were introduced by (Selman and Kautz 1996) in the context of knowledge compilation. Selman and Kautz showed that Horn envelopes can blow up in size or can be hard to compute. Computational aspects of Horn envelopes are studied in (Langlois, Sloan, and Turán 2009). The special case of computing the Horn envelope of the disjunction of two Horn formulas has been considered in (Eiter, Ibaraki, and Makino 2001; Eiter and Makino 2008). They showed negative results analogous to the general case. Proposition 1 suggests considering the special case where one of the two Horn formulas is an elementary conjunction, i.e., it is satisfied by a single truth assignment.

Definition 7. *The* Singleton Horn Extension (SHE) problem *is to compute a formula for the Horn envelope $\mathrm{Env}(K \vee E_a)$ given definite Horn belief set K and a truth assignment a falsifying K,*

Later in this section we show that the SHE problem is intractable in general, as it may be the case that *every* Horn formula representing the output must be exponentially large compared to the input size. This negative result is perhaps surprising, as it says that adding a single additional true point may result in an exponential blowup in formula size. The result does not depend on any unproven complexity theoretic assumptions. On the other hand, it assumes that the output has to be represented as a Horn formula. In view of the negative result it is of interest to identify cases where the prob-lem has an efficient solution. We begin with such positive results.

Positive results

In this subsection, we give some simple observations on cases when the blowup outlined above cannot occur.

The following proposition shows that $\mathrm{Env}(K \vee E_a)$ has an explicit description in terms of the prime implicates of K. Let $PI(K)$ denote the set of prime implicates of K, and $PI^1(K, a)$, resp., $PI^0(K, a)$ be the set of prime implicates of K satisfied, resp., falsified by a. We write $a = (a_1, \ldots, a_n)$ and $x^1 = x, x^0 = \bar{x}$.

Proposition 8. $\mathrm{Env}(K \vee E_a)$ *can be written as*

$$\left(\bigwedge_{C \in PI^1(K,a)} C \right) \wedge \left(\bigwedge_{C \in PI^0(K,a)} \bigwedge_{\{i : x_i(a) = 0\}} (C \vee \bar{x}_i) \right).$$

Proof. We show that $\mathrm{Env}(K \vee E_a)$ is logically equivalent to

$$\bigwedge_{C \in PI(K), \, C \vee x_i^{a_i} \text{ definite}} (C \vee x_i^{a_i}) . \qquad (7)$$

This, then, can be rewritten in the form stated in the proposition. The '$\models$' direction follows by noting that by distributivity, every clause in (7) is a Horn implicate of $K \vee E_a$, and thus it is an implicate of $\mathrm{Env}(K \vee E_a)$.

For the other direction, consider a Horn prime implicate D of $\mathrm{Env}(K \vee E_a)$. Then D is an implicate of K. As K is definite, D is satisfied by the all 1's vector and so D is also definite. Let $D' \subseteq D$ be a prime implicate of K. Again, D' is definite and so it contains the head of D. As $E_a(a) = 1$, it holds that $D(a) = 1$ and so D contains a literal $x_i^{a_i}$. If this literal is the head of D then $D' \vee x_i^{a_i}$ is equivalent to D' and so D' occurs in (7). Otherwise $a_i = 0$, so $D' \vee x_i^{a_i}$ is definite, and it occurs in (7). In both cases we get that D is an implicate of (7) □

Proposition 8 does *not* lead to an efficient algorithm for computing $\mathrm{Env}(K \vee E_a)$ in general, as K can have exponentially many prime implicates compared to its size. An example is given in (Khardon 1995), and a similar example is given in the next subsection. Nevertheless, one can draw positive algorithmic consequences, and we formulate two of those. We use the fact that the prime implicates of a Horn formula can be listed efficiently (Boros, Crama, and Hammer 1990).

Corollary 9. *The SHE problem can be solved in time polynomial in the size of K and the number of prime implicates of K.*

Proof. The algorithm first runs Boros et al.'s algorithm to generate the prime implicates of K and then uses Proposition 8 to produce $\mathrm{Env}(K \vee E_a)$. □

Corollary 9 provides efficient algorithms for any class of belief sets with small number of prime implicates. One such class is quadratic Horn formulas. A definite Horn formula is *quadratic* if all its clauses are of size two, i.e., they are of the form $x \rightarrow y$.

Corollary 10. *The SHE problem can be solved efficiently for quadratic belief sets K.*

Proof. The resolvent of size-two clauses is again of size two, and hence an n-variable quadratic belief set has $O(n^2)$ prime implicates. $\square$

Quadratic Horn formulas are one of the tractable subclasses of Horn formulas for problems which are hard for Horn formulas in general, such as minimization. The other tractable class is acyclic formulas: a definite Horn formula is *acyclic* if the directed graph over the set of variables, obtained by adding a directed edge from every body variable to the head variable, has no directed cycles (Boros et al. 2010). It is tempting to conjecture that the SHE problem might also be tractable for acyclic formulas. The next subsection shows that this is not the case.

Variants of Proposition 8 can be formulated for formulas over other sets of clauses satisfying some general conditions, along the lines of (del Val 2005). This may be of interest for belief change over other restricted logics.

A negative result

We now give a negative result that is more general than just showing an example where the SHE problem is difficult, because we give a blowup result for a *set* of extensions (Theorem 11). The blowup result for the SHE problem itself is an immediate corollary. We will use Theorem 11 to prove hardness results for Horn contraction in Section 5.

Consider variables $u_{i,j}, v_i$ and w, where $1 \le i \le n, 1 \le j \le 2$, and let the Horn belief set K_n be given by the $2n+1$ Horn clauses

$$u_{i,1} \to v_i$$
$$u_{i,2} \to v_i \quad 1 \le i \le n, \quad \text{and}$$
$$v_1, \ldots, v_n \to w.$$

Note that K_n is acyclic. The belief set K_3 is shown in Figure 1.

Let A be a set of truth assignments to the variables $u_{i,j}, v_i$ and w (where $1 \le i \le n, 1 \le j \le 2$) such that in every $a \in A$ the u-variables and w are set to 1 and v_1 is set to 0. Assume, furthermore, that there are altogether $k \ge 1$ variables each of which is 0 in at least one $a \in A$. It can be assumed *w.l.o.g.* that these variables are $v_1, \ldots, v_k$.

Theorem 11. *Every Horn formula representing*

$$\mathrm{Env}\left(K_n \vee \bigvee_{a \in A} E_a\right)$$

contains at least 2^k clauses.

Proof. Let χ be a Horn formula representing $\mathrm{Env}\left(K_n \vee \bigvee_{a \in A} E_a\right)$. Note that K_n is definite, so it is satisfied by the all ones vector. Thus χ is also satisfied by the all ones vector and so it is definite as well.

For every $s = (s_1, \ldots, s_k)$ such that $1 \le s_i \le 2$ for $i = 1, \ldots, k$ consider the clause

$$\psi_s = (u_{1,s_1}, \ldots, u_{k,s_k}, v_{k+1}, \ldots, v_n \to w).$$

The theorem follows if we show that χ contains ψ_s.

The clause ψ_s is an implicate of K_n, as after marking its body variables, we can mark $v_1, \ldots, v_k$ and then we can mark w. Also, ψ_s is satisfied by all $a \in A$. Thus ψ_s is an implicate of χ as well.

It also holds that ψ_s is a prime implicate of K_n. Indeed, if $C' = \psi_s \setminus \{u_{i,s_i}\}$ for some $1 \le i \le k$, or if $C' = \psi_s \setminus \{v_i\}$ for some $k + 1 \le i \le n$, then the truth assignment setting $u_{i,1}, u_{i,2}, v_i$ and w to 0, and setting all other variables to 1, satisfies K_n and falsifies C'. If $C' = \psi_s \setminus \{w\}$ then the all ones truth assignment satisfies K_n and falsifies C'.

Let b be the truth assignment for which every body variable in ψ_s is set to 1 and all other variables are set to 0. As b falsifies ψ_s, it also falsifies χ, and so there is a clause C in χ such that $C(b) = 0$. Thus it must be the case that

$$\mathrm{Body}(C) \subseteq \{u_{1,s_1}, \ldots, u_{k,s_k}, v_{k+1}, \ldots, v_n\}.$$

Furthermore, C is an implicate of K_n thus its head is in

$$\mathrm{Cl}_{K_n}(\mathrm{Body}(C)) \setminus \mathrm{Body}(C) \subseteq \{v_1, \ldots, v_k, w\}.$$

Now C is satisfied by every $a \in A$. But every $a \in A$ satisfies the body of C, and every $v_i (1 \le i \le k)$ is falsified by some $a \in A$. So the head of C cannot be a v-variable and thus it must be w. So C must be a subclause of ψ_s, and as ψ_s was shown to be a prime implicate of K_n, it must be equal to ψ_s. $\square$

Corollary 12. *There is a family of instances of the SHE problem where every Horn formula representation of the result requires exponentially many clauses.*

Proof. Let $A = \{a_n\}$, where a_n is the truth assignment setting all u-variables to 1, all v-variables to 0 and w to 1. Clearly a_n falsifies K_n. $\square$

A remark on characteristic models

For every Horn formula ψ the set $T(\psi)$ of satisfying truth assignments is closed under intersection (Horn 1951; McKinsey 1943). Those satisfying truth assignments which cannot be obtained as the intersection of others are called the *characteristic models* or *characteristic vectors* of ψ (Kautz, Kearns, and Selman 1995; Khardon and Roth 1996). Representing a Horn function by its set of characteristic vectors is an alternative to the standard clausal representation. This representation has various advantages and disadvantages. The clausal and characteristic set representations are incomparable in the sense that there are examples where one has polynomial size and the other has exponential size (Khardon and Roth 1996).

(Delgrande and Wassermann 2010) discuss possible connections of characteristic models to Horn belief contraction. Let us assume that the Horn belief set K is represented by its set of characteristic vectors $\mathrm{Char}(K)$. Then every truth assignment $b \in T(K \vee E_a)$ is obtained as an intersection of vectors in $\mathrm{Char}(K) \cup \{a\}$, hence $\mathrm{Char}(K \vee E_a) \subseteq \mathrm{Char}(K) \cup \{a\}$. The set $\mathrm{Char}(K \vee E_a)$ can be found efficiently by eliminating those vectors from $\mathrm{Char}(K) \cup \{a\}$ which can be obtained as the intersection of vectors above them in the set.

Figure 1: The belief set K_3

In view of the negative results of the previous subsection one may ask whether $\mathrm{Env}(K \vee E_a)$ has a short Horn formula representation if, in addition, it holds that $\mathrm{Char}(K)$ is small. Unfortunately, this is not the case as the characteristic set of the belief set K_n considered above turns out to be small.

Proposition 13. $|\mathrm{Char}(K_n)| = \Theta(n)$.

Proof. The characteristic models are the following: those with a single $u_{i,j} = 0$ and all other variables set to 1, those with $u_{i,1} = u_{i,2} = v_i = w = 0$ for a single i and all other variables set to 1, those with $u_{i,1} = u_{i,2} = v_i = 0$ for a single i and all other variables set to 1, and the all 1's vector. □

5   Complexity of Horn belief contraction

In this section we draw conclusions from Theorem 11 for Horn belief contraction.

A *partial meet contraction* $K \dot{-} \varphi$ is an intersection of remainders, thus by Theorem 4 it is of the form

$$\mathrm{Env}\left(K \vee \bigvee_{a \in A} E_a \right), \tag{8}$$

where

$$A \subseteq T(K^\varphi) \cap F(\varphi). \tag{9}$$

A *maxichoice contraction* corresponds to a singleton subset in (9) and *full meet contraction* corresponds to equality in (9).

If K is a Horn belief set and Horn formula φ is a consequence of K then a *weak remainder* is a belief set of the form $\mathrm{Env}(K \vee E_a)$ for any $a \in F(\varphi)$ and so the set of weak remainders is

$$K \Downarrow \varphi = \{\mathrm{Env}(K \vee E_a) : a \in F(\varphi)\}.$$

A *partial meet contraction based on weak remainders* $K \dot{-}_w \varphi$ is of the form (8) where

$$A \subseteq F(\varphi). \tag{10}$$

A *maxichoice contraction based on weak remainders* corresponds to a singleton subset in (10) and *full meet contraction based on weak remainders* corresponds to equality in (10).

In order to make use of Theorem 11 in the context of contractions we also need to specify a consequence to be contracted. Let the implicate to be contracted be

$$\varphi_n = u_{1,1}, \ldots, u_{i,j}, \ldots, u_{n,2}, w \to v_1.$$

It is clear that φ_n is an implicate because after marking $u_{1,1}$ we can already mark v_1.

Proposition 14.

$$K_n \downarrow \varphi_n = K_n \Downarrow \varphi_n = \{\mathrm{Env}(K_n \vee E_a) : a \in F(\varphi_n)\}$$

Proof. This follows from Theorem 4 noting that $\mathrm{Cl}_{K_n}(\mathrm{Body}(\varphi_n))$ is the set of all variables, thus the body building formula K^φ is the empty conjunction, and so it is identically true. □

Let us consider a partial meet contraction K' of the form

$$K' = \mathrm{Env}\left(K_n \vee \bigvee_{a \in A} E_a \right), \tag{11}$$

where

$$A \subseteq F(\varphi_n) \tag{12}$$

and assume that there are k variables v_i such that $v_i(a) = 0$ for some $a \in A$. It follows from Theorem 11 that:

Corollary 15. *Every representation of* any *partial meet contraction or partial meet contraction based on weak remainders of the form of K' in Equation* (11) *has at least 2^k clauses.*

In the following theorem a size lower bound is said to hold for *almost all* contractions if the fraction of contractions with at least that size approaches 1 as n grows.

Theorem 16. *Consider contractions, or weak remainder based contractions, of the consequence φ_n from the belief set K_n.*

a) *Every Horn formula representation of the full meet contraction contains at least 2^n clauses.*

b) *For every $\epsilon > 0$ and for almost all maxichoice contractions, every Horn representation contains at least $2^{((1/2)-\epsilon)n}$ clauses.*

c) *For almost all partial meet contractions, every Horn representation contains at least 2^n clauses.*

Proof. Part *a)* follows from Corollary 15 with $A = F(\varphi_n)$.

Part *b)* also follows from Corollary 15 by noting that there A is a single truth assignment, with all u variables and w set to 1 and v_1 set to 0. The remaining $n-1$ v variables are arbitrary. It follows from standard probability estimates that for every $\epsilon > 0$ there are $o(2^n)$ truth assignments with fewer than $((1/2) - \epsilon)n$ zeros. For the remaining $(1 - o(1))2^{n-1}$ truth assignments Corollary 15 implies a $2^{((1/2)-\epsilon)n}$ lower bound.

In part *c)* we consider random partial meet contractions, i.e., a random subsets of all truth assignments for which all u

variables and w are set to 1 and v_1 is set to 0. The probability that there is a variable v_i which is never set to 0 is at most $(n-1)2^{2^{n-2}}/2^{2^{n-1}} \to 0$. For the other choices Corollary 15 implies a 2^n lower bound. $\qquad\square$

6 Further remarks

We have shown that every Horn representation of the full meet contraction, and of most maxichoice and partial meet contractions of the Horn belief set K_n with respect to its consequence φ_n must be exponentially large. This belief set is simple and natural in the sense that it can be thought of as consisting of observable propositions $u_{i,j}$, intermediate conclusions v_i and a final conclusion w, where each $u_{i,j}$ is sufficient to cause v_i and all v_i's are necessary to cause w. The fact that contractions of such a simple belief set may blow up in size may indicate that this phenomenon occurs more often than just for an artificially constructed pathological example. It would be interesting to perform experiments exploring this, and, more generally, to gather computational experience about various other aspects of Horn belief contractions. Some initial results are given in (Langlois et al. 2008).

In view of the fact that fully AGM-compliant Horn-contractions do not exist in the sense of (Flouris, Plexousakis, and Antoniou 2004), it was suggested in (Langlois et al. 2008) that one might study the possibilities of *approximating* them. The negative results in the current paper give a different motivation for such an approach: When the result of a contraction is too large, approximate it with a smaller one. This may be related to *anytime* belief revision algorithms (Williams 1997).

The results of this paper also suggest several specific questions for further study, such as considering the complexity of infra-remainders (Booth et al. 2011) and package contraction (Booth, Meyer, and Varzinczak 2009), and proving complexity-theoretic hardness results for the SHE problem analogous to the results of (Eiter and Makino 2008).

Acknowledgments

References

Adaricheva, K.; Nation, J. B.; and Rand, R. 2011. Ordered direct implicational basis of a finite closure system. *arXiv:1110.5805*.

Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *J. Symb. Logic* 50(2):510–530.

Bertet, K., and Monjardet, B. 2010. The multiple facets of the canonical direct unit implicational basis. *Theoret. Comput. Sci.* 411(22-24):2155–2166.

Booth, R.; Meyer, T.; Varzinczak, I.; and Wassermann, R. 2011. On the link between partial meet, kernel, and infra contraction and its application to Horn logic. *J. Artif. Intell. Research (JAIR)* 42:31–53.

Booth, R.; Meyer, T.; and Varzinczak, I. 2009. Next steps in propositional Horn contraction. In *Proc. IJCAI*, 702–707.

Boros, E.; Cepek, O.; Kogan, A.; and Kucera, P. 2010. Exclusive and essential sets of implicates of Boolean functions. *Discrete Appl. Math.* 158(2):81 – 96.

Boros, E.; Crama, Y.; and Hammer, P. L. 1990. Polynomial-time inference of all valid implications for Horn and related formulae. *Annals of Mathematics and Artificial Intelligence* 1(1):21–32.

Burosch, G.; Demetrovics, J.; and Katona, G. O. H. 1987. The poset of closures as a model of changing databases. *Order* 4(2):127–142.

Caspard, N., and Monjardet, B. 2003. The lattices of closure systems, closure operators, and implicational systems on a finite set: a survey. *Discrete Appl. Math.* 127(2):241–269.

Crama, Y., and Hammer, P. L. 2011. *Boolean Functions: Theory, Algorithms, and Applications*. Cambridge University Press.

Creignou, N.; Papini, O.; Pichler, R.; and Woltran, S. 2012. Belief revision within fragments of propositional logic. Technical Report DBAI-TR-2012-75, Technische Universität Wien.

Davey, B. A., and Priestley, H. A. 1990. *Introduction to Lattices and Order*. Cambridge University Press.

del Val, A. 2005. First order LUB approximations: characterization and algorithms. *Artificial Intelligence* 162(1-2):7–48.

Delgrande, J., and Peppas, P. 2011. Revising Horn theories. In *Twenty-Second International Joint Conference on Artificial Intelligence*.

Delgrande, J. P., and Wassermann, R. 2010. Horn clause contraction functions: Belief set and belief base approaches. In *Proc. 12th Int. Conf. Principles of Knowledge Representation and Reasoning (KR)*.

Delgrande, J. P. 2008. Horn clause belief change: Contraction functions. In *Proc. 11th Int. Conf. Principles of Knowledge Representation and Reasoning (KR)*, 156–165.

Eiter, T., and Gottlob, G. 1992. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence* 57:227–270.

Eiter, T., and Makino, K. 2008. New results for Horn cores and envelopes of Horn disjunctions. In *Proc. ECAI: 18th European Conference on Artificial Intelligence*, 60–64.

Eiter, T.; Ibaraki, T.; and Makino, K. 2001. Disjunctions of Horn theories and their cores. *SIAM Journal on Computing* 31(1):269–288.

Fermé, E., and Hansson, S. 2011. AGM 25 years. *Journal of Philosophical Logic* 1–37.

Flouris, G.; Plexousakis, D.; and Antoniou, G. 2004. Generalizing the AGM postulates: Preliminary results and applications. In *Proc. 10th Int. Workshop on Non-Monotonic Reasoning (NMR)*, 171–179.

Fotinopoulos, A. M., and Papadopoulos, V. 2009. Semantics for Horn contraction. In *Proc. 7th Panhellenic Logic Symposium*, 42–47. Patras University Press.

Freese, R. 1995. *Free Lattices*. American Mathematical Society.

Goldberg, L. A. 1993. *Efficient Algorithms for Listing Combinatorial Objects*. Distinguished Dissertations in Computer Science. Cambridge University Press.

Hansson, S. O. 1999. *A Textbook on Belief Dynamics: Theory Change and Database Updating*. Kluwer.

Horn, A. 1951. On sentences which are true on direct unions of algebras. *J. Symb. Logic* 16:14–21.

Kautz, H. A.; Kearns, M. J.; and Selman, B. 1995. Horn approximations of empirical data. *Artificial Intelligence* 74(1):129–145.

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

Khardon, R. 1995. Translating between Horn representations and their characteristic models. *J. Artif. Intell. Research (JAIR)* 3:349–372.

Kleine Büning, H., and Lettmann, T. 1987. Representation independent query and update operations on propositional definite Horn formulas. In *Computation Theory and Logic*, volume 270 of *LNCS*. Springer. 208–223.

Kleine Büning, H., and Lettmann, T. 1999. *Propositional Logic: Deduction and Algorithms*. Cambridge University Press.

Langlois, M.; Sloan, R. H.; Szörényi, B.; and Turán, G. 2008. Horn complements: Towards Horn-to-Horn belief revision. In *Proc. AAAI 2008*, 466–471.

Langlois, M.; Sloan, R. H.; and Turán, G. 2009. Horn upper bounds and renaming. *JSAT: Journal on Satisfiability, Boolean Modeling and Computation* 7:1–15.

Liberatore, P. 2000. Compilability and compact representations of revision of Horn knowledge bases. *ACM Trans. Computational Logic* 1(1):131–161.

McKinsey, J. C. C. 1943. The decision problem for some classes without quantifiers. *J. Symb. Logic* 8:61–76.

Nebel, B. 1998. How hard is it to revise a belief base? In *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, volume 3. Kluwer Academic. 77–145.

Peppas, P. 2008. Belief revision. In *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence*. Elsevier. chapter 8, 317–359.

Ribeiro, M. 2010. *Belief Revision in Description Logics and Other Non-Classical Logics*. Ph.D. Dissertation, University of São Paulo.

Schaefer, T. 1978. The complexity of satisfiability problems. In *Proc. Tenth Annual ACM Symposium on Theory of Computing*, 216–226.

Selman, B., and Kautz, H. 1996. Knowledge compilation and theory approximation. *Journal of the ACM* 43:193–224.

Williams, M. A. 1997. Anytime belief revision. In *Proc. IJCAI*, volume 15, 74–81.

Zhuang, Z., and Pagnucco, M. 2010a. Horn contraction via epistemic entrenchment. In *Logics in Artificial Intelligence*, 339–351.

Zhuang, Z., and Pagnucco, M. 2010b. Two methods for constructing Horn contractions. In *Proc. AI 2010: Advances in Artificial Intelligence*, volume 6464 of *LNCS*, 72–81.

Zhuang, Z., and Pagnucco, M. 2011. Transitively relational partial meet Horn contraction. In *Twenty-Second International Joint Conference on Artificial Intelligence*.