Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (2008)

Horn Complements: Towards Horn-to-Horn Belief Revision

Marina Langlois and **Robert H. Sloan**
University of Illinois at Chicago
mirodo1 | sloan@uic.edu

Balázs Szörényi
Hungarian Academy of Sciences & U. of Szeged
szorenyi@inf.u-szeged.hu

György Turán
University of Illinois at Chicago,
Hungarian Academy of Sciences & U. of Szeged
gyt@uic.edu

Abstract

Horn-to-Horn belief revision asks for the revision of a Horn knowledge base such that the revised knowledge base is also Horn. Horn knowledge bases are important whenever one is concerned with *efficiency*—of computing inferences, of knowledge acquisition, etc. Horn-to-Horn belief revision could be of interest, in particular, as a component of any efficient system requiring large commonsense knowledge bases that may need revisions because, for example, new contradictory information is acquired.

Recent results on belief revision for general logics show that the existence of a belief contraction operator satisfying the generalized AGM postulates is equivalent to the existence of a *complement*. Here we provide a first step towards efficient Horn-to-Horn belief revision, by characterizing the existence of a complement of a Horn consequence of a Horn knowledge base. A complement exists if and only if the Horn consequence is not the consequence of a modified knowledge base obtained from the original by an operation called body building. This characterization leads to the efficient construction of a complement whenever it exists.

Introduction

Revising a knowledge base in the presence of new, potentially conflicting information is a basic task facing a commonsense reasoning agent. Belief revision usually approaches this task by identifying postulates that should be satisfied by a rational revision operator, such as the AGM postulates (Alchourrón, Gärdenfors, and Makinson 1985), and characterizing operators that satisfy these postulates (Gärdenfors 1988; Hansson 1999). The basic operators are belief *revision*, when a new, perhaps contradictory belief is to be incorporated into the knowledge base, and belief *contraction*, when an undesirable consequence is to be removed from the knowledge base. One often considers contraction first, and then defines revision in a standard way in terms of contraction. Most of the work in belief revision assumes that the underlying logic includes full propositional logic. On the other hand, motivated in part by efficiency considerations, in many applications one uses a logic based on only a *fragment* of propositional logic.

Adapting belief revision to this more general situation has been initiated by the recent work of (Flouris, Plexousakis, and Antoniou 2004). They study belief revision in *general logics*, and formulate a property called *decomposability* of the logic. It is shown that decomposability is a necessary and sufficient condition for the existence of an AGM-compliant belief contraction operator. This framework is then used in (Flouris, Plexousakis, and Antoniou 2005) to study decomposition properties of description logics, motivated by applications to the Semantic Web.

In order to build a powerful agent capable of commonsense reasoning, one fundamental challenge is to integrate various capabilities, including belief revision, that have been studied so far mostly in isolation. A commonsense reasoning agent should execute its tasks efficiently, and therefore it has to use some tractable knowledge representation. In this paper we explore an approach that takes *tractable knowledge representation as the primary constraint*, and considers other requirements, in particular rationality constraints, as also important, but secondary. For belief revision, this is a departure from the standard framework.

Horn formulas provide a natural candidate for a general framework for the integration process. This fragment of propositional logic is expressive, allows for polynomial time inference, and indeed is generally computationally tractable, which explains its central role in artificial intelligence and computer science. In particular, belief revision with Horn formulas has been studied extensively (see references in the next section). However, the problem of belief revision that maintains a Horn knowledge base throughout the revision process apparently has not been studied. Here we consider such *Horn-to-Horn* revisions. That is, we are interested in the possibilities and limitations of revising Horn formulas such that the revised formula is also Horn. This, then, is a special case of the general framework of (Flouris, Plexousakis, and Antoniou 2004), and thus their general characterization for the existence of an AGM-compliant belief contraction operator applies. Horn logic as a whole turns out not to be decomposable. Thus one must ask the more detailed question of when contraction is possible. More precisely, one is led to the following problem, which may also be of interest in itself.

Problem A. For Horn formulas φ and ψ, where ψ is a consequence of φ, when does there exist a proper Horn consequence χ of φ, such that $\psi \wedge \chi$ is equivalent to φ?

Such a formula χ is called a *φ-complement* of ψ, and it corresponds to the result of a contraction operator, when the formula ψ is contracted from the knowledge base φ. For the decomposability framework, it would be desirable always to have a complement, but unfortunately this is not the case. (See Example 9.) When a complement does not exist, one may try to find an approximate complement.

Our main result (Theorem 7) gives a complete answer to Problem A by giving two characterizations of all those pairs φ and ψ for which ψ has a φ-complement. The characterizations give efficiently decidable criteria and lead to efficient algorithms to construct a complement, if it exists. The complements constructed are only polynomially larger than the original knowledge base (although repeated application could cause an exponential blowup). As a corollary, one obtains a complete description of *decomposable* Horn formulas as well, where a Horn formula is decomposable if all its Horn consequences have a complement. We also present some computational results on the fraction of Horn implicates of a random Horn formula having a complement.

The rest of the paper is structured as follows. The next section gives a brief overview of previous related work. The section after that gives a more detailed review of Flouris, Plexousakis, and Antoniou's AGM-based framework and their basic results together with the formal connection between belief revision theory and the rest of this paper. Then background is provided on Horn formulas. The next three sections contain our results, a proof sketch of the main theorem and the sketch of another proof of the first characterization. The last two sections present the experimental results and some remarks on further work.

Related work

Horn formulas have been considered previously in several papers dealing with the complexity of belief revision, e.g., (Eiter and Gottlob 1992; Gogic, Papadimitriou, and Sideri 1998; Jin and Thielscher 2005; Liberatore 1997; 2000; Nebel 1998). The results obtained in these papers are mostly complexity-theoretic negative results, and they deal with revision methods where the revision of a Horn knowledge base is not necessarily Horn, or they propose revision methods that may be inefficient. A more detailed comparison with our work will be given in the full version of this paper.

If ψ is a single Horn clause implicate C, then Problem A can be reformulated as follows: does φ have an irredundant conjunctive normal form expression containing C? The related question where C is a *prime* implicate and the irredundant conjunctive normal form expression is also assumed to consist of *prime* implicates only, has been studied by (Hammer and Kogan 1995). They call such a prime implicate *non-redundant*, and show that non-redundancy is polynomially decidable for negative clauses, but it is NP-complete for definite clauses.

General logics and belief contraction

In this section we give a brief outline of the formal framework for belief revision in general logics, following (Flouris, Plexousakis, and Antoniou 2004). We also indicate how this very general and abstract framework can be specialized to Horn logic.

A *logic* is specified by a set of expressions L and a consequence operator $Cn : \mathcal{P}(L) \to \mathcal{P}(L)$, where $\mathcal{P}(L)$ is the power set of L. The consequence operator is assumed to satisfy the properties of *inclusion* ($A \subseteq Cn(A)$), *iteration* ($Cn(A) = Cn(Cn(A))$) and *monotonicity* ($A \subseteq B$ implies $Cn(A) \subseteq Cn(B)$) for every $A, B \subseteq L$.

For Horn logic, L is the set of all Horn clauses over a fixed finite set of variables, and for a set of Horn clauses φ, the set $Cn(\varphi)$ contains all Horn clauses implied by φ. The required properties are clearly satisfied.

A *theory* or *knowledge base* K is a subset of L such that $K = Cn(K)$. For Horn logic, a theory can be specified by a set of Horn clauses. The corresponding theory then consists of all their consequences. A *contraction operator*, denoted by $-$, is of the form $- : \mathcal{P}(L) \times \mathcal{P}(L) \to \mathcal{P}(L)$, and maps a knowledge base and a set to be contracted to a new knowledge base, the result of the contraction. The (generalized) AGM postulates for contraction are [1]

closure: $K - A = Cn(K - A)$

inclusion: $K - A \subseteq Cn(K)$

vacuity: $A \not\subseteq Cn(K)$ implies $K - A = Cn(K)$

success: $A \not\subseteq Cn(\emptyset)$ implies $A \not\subseteq Cn(K - A)$

preservation: $Cn(A) = Cn(B)$ implies $K - A = K - B$

recovery: $K \subseteq Cn((K - A) \cup A)$.

A contraction operator is *AGM-compliant* if it satisfies these postulates. A logic is *AGM-compliant* if there exists an AGM-compliant contraction operator for it.

Given a logic, and subsets $K, A \subseteq L$, the set of *complements* of A with respect to K is

$$A^{-}(K) =$$

$$\{B \subseteq L : Cn(B) \subset Cn(K) \text{ and } Cn(A \cup B) = Cn(K)\}$$

if $Cn(\emptyset) \subset Cn(A) \subseteq Cn(K)$, and $A^{-}(K) = \{B \subseteq L : Cn(B) = Cn(K)\}$ otherwise.

A set $K \subseteq L$ is *decomposable* if $A^{-}(K) \neq \emptyset$ for every $A \subseteq L$. A logic is *decomposable* if every K is decomposable. For Horn logic, these definitions specialize to Definitions 3 and 4.

Theorem 1 ((Flouris, Plexousakis, and Antoniou 2004))**.** *A logic is AGM-compliant iff it is decomposable.*

In particular, Flouris, Plexousakis, and Antoniou show that if a logic is decomposable, then by selecting any complement one obtains an AGM-compliant contraction operator. More generally, if K is decomposable, then by selecting

[1](Flouris, Plexousakis, and Antoniou 2004) refers to these as *generalized* versions of the properties, but for simplicity we omit the term 'generalized'.

any complement with respect to K, one obtains a contraction operator that handles contractions from K in an AGM-compliant manner. It also follows from their arguments that if $A^-(K) \neq \emptyset$ then selecting any complement of A with respect to K handles the contraction of A from K in an AGM-compliant manner.

Preliminaries

Let U be the set of propositional variables in our universe. Two clauses *collide* if they contain a pair of complementary literals. A clause is *Horn* if it contains at most one unnegated literal, *definite* if it contains exactly one unnegated literal, and *negative* if it contains only negated literals. A *(definite) Horn formula* is a conjunction—or a set, whichever view is more convenient—of (definite) Horn clauses (an empty conjunction is always true). A Boolean function is a *(definite) Horn function* if it has a (definite) Horn formula.

Lower case Roman alphabet letters f, g, h denote Boolean *functions*; lower case Greek letters φ, ψ, χ denote Boolean *formulas*. If we use φ in a place where a function is expected, then φ stands for the Boolean function represented by φ. Two formulas φ_1 and φ_2 are *equivalent*, denoted $\varphi_1 \sim \varphi_2$, if they represent the same Boolean function.

For a Horn clause C, let $\text{Body}(C)$ be the set of variables corresponding to the negated literals in C, or their conjunction (which will be clear from context). Also, let $\text{Head}(C)$ be the unnegated variable of C if C is a definite clause, and 0 if C is a negative clause. We use $\rightarrow$ to denote the Boolean implication operator, so Horn clause C can be written as $\text{Body}(C) \rightarrow \text{Head}(C)$. For example, if C is the Horn clause $\bar{x} \vee \bar{y} \vee z$, then $\text{Body}(C) = \{x, y\}$, $\text{Head}(C)$ is z, and C can also be written as $x, y \rightarrow z$ or $(x \wedge y) \rightarrow z$. If C is the Horn clause $\bar{x} \vee \bar{y}$ then it can also be written as $x, y \rightarrow 0$ or simply $x, y \rightarrow$.

A Boolean function g is a *consequence* of Boolean function f, denoted $f \Rightarrow g$, if every assignment that satisfies f also satisfies g. A function g is a *proper consequence* of f, denoted $f \overset{\Rightarrow}{\not\Leftarrow} g$, if $f \Rightarrow g$ but not $g \Rightarrow f$. A clause C is an *implicate* of f if $f \Rightarrow C$.

The set of satisfying (resp., falsifying) truth assignments of f is denoted by $T(f)$ (resp., $F(f)$). A function f is *anti-monotone* if $T(f)$ is downward closed, i.e., $f(a) = 1$ and $b \leq a$ imply $f(b) = 1$.

We will use a slight generalization of anti-monotone functions.

Definition 2 (Almost anti-monotone function)**.** A function is almost anti-monotone if it is either anti-monotone, or there is an anti-monotone function g such that $T(f) = T(g) \cup \{\mathbf{1}\}$, where $\mathbf{1}$ is the all 1's assignment.

Every almost anti-monotone function is Horn. Now we formulate the central concept discussed in this paper.

Definition 3 (f-complement)**.** For Horn functions f and g such that $f \Rightarrow g$, a Horn function h is an *f-complement of g* iff $f \overset{\Rightarrow}{\not\Leftarrow} h$ and $f \sim (g \wedge h)$.

According to the definition, no complements exist if $f \sim 1$ (where 1 denotes the identically 1 function). Also according to the definition, $g \sim 1$ can never have a complement, so this case is excluded from consideration in the following definition.

Definition 4 (Decomposable Horn function)**.** A Horn function f is *decomposable* if every Horn consequence $g \not\sim 1$ of f has an f-complement.

One usually works with formulas as opposed to functions, but as the notions of complement and decomposability depend only on the function represented by the formula, the definitions are given in a syntax-independent way.

Results

For a function f and a set of variables $X \subseteq U$, we define the *f-closure of X* to be the set of variables

$$\text{Cl}_f(X) = \{v \in U : \ f \Rightarrow (X \rightarrow v)\} \quad .$$

A direct consequence of this definition is that if a negative clause C is an implicate of f, then $\text{Cl}_f(\text{Body}(C)) = U$.

In order to formulate our main result, we need two definitions. The formula $\hat{\varphi}$ is obtained from φ by adding to the body of each definite clause in φ a variable not contained in the closure of its body, in all possible ways.

Definition 5 (Body-building formula $\hat{\varphi}$)**.** For a Horn formula φ let $\hat{\varphi}$ be the formula

$$\bigwedge_{C \in \varphi \text{ definite}} \ \bigwedge_{v \notin \text{Cl}_\varphi(\text{Body}(C))} (\text{Body}(C), v \rightarrow \text{Head}(C)).$$

We could have defined $\hat{\varphi}$ as a conjunction over all clauses of φ, as negative clauses would make no contribution. Every clause of $\hat{\varphi}$ is definite. It may be the case that $\hat{\varphi}$ is the empty conjunction. This happens, for example, when φ consists of negative clauses only.

Given a Horn formula φ and a Horn clause D, we partition the clauses of φ not colliding with D into two classes.

Definition 6 (Formulas $\mathcal{A}_\varphi(D)$ and $\mathcal{B}_\varphi(D)$)**.** Given a Horn formula φ and a Horn clause D, let

$$\mathcal{A}_\varphi(D) = \{C \in \varphi : C, D \text{ don't collide, } \text{Body}(D) \subseteq \text{Cl}_\varphi(\text{Body}(C))\},$$
$$\mathcal{B}_\varphi(D) = \{C \in \varphi : C, D \text{ don't collide, } \text{Body}(D) \not\subseteq \text{Cl}_\varphi(\text{Body}(C))\}.$$

The existence of a complement can now be characterized.

Theorem 7 (Main theorem)**.** *Let $\varphi \not\sim 1$ be a Horn formula, and ψ be a Horn consequence of φ. Then the following are equivalent:*

1. *ψ has a φ-complement,*
2. *$\hat{\varphi} \not\Rightarrow \psi$,*
3. *for some clause D of ψ it holds that $\mathcal{B}_\varphi(D) \not\Rightarrow D$.*

Although the definition of $\hat{\varphi}$ is given in terms of a formula, it follows from this characterization that it depends on only the function (see also Lemma 14 below). The following corollary gives the algorithmic aspects of Theorem 7.

Corollary 8. *There is a polynomial time algorithm which, given a Horn formula φ and a Horn consequence ψ of φ, decides if ψ has a φ-complement, and if it does, then constructs such a φ-complement.*

The following simple example illustrates the results.

Example 9. Let $U = \{x, y, z\}$, $\varphi = C_1 \wedge C_2$, where $C_1 = (x \to y)$ and $C_2 = (y \to z)$. Then $\mathrm{Cl}_\varphi(x) = U$ and $\mathrm{Cl}_\varphi(y) = \{y, z\}$. So $\hat{\varphi} = (x, y \to z)$.

The clause $(x, y \to z)$ is implied by $\hat{\varphi}$, and so it has no φ-complement. This is also shown by the fact that $\mathcal{B}_\varphi(x, y \to z) = \{y \to z\}$, which implies $(x, y \to z)$.

On the other hand, the clause $(x \to z)$ is not implied by $\hat{\varphi}$, so it does have a φ-complement. This is also shown by the fact that $\mathcal{B}_\varphi(x \to z) = \{y \to z\}$, which does not imply $(x \to z)$. Both constructions mentioned in the paper give the φ-complement $(x, z \to y) \wedge (y \to z)$.

Decomposable Horn functions have the following characterization.

Theorem 10. *For every Boolean function f the following are equivalent:*

1. *f is a decomposable Horn function,*
2. *there is a Horn representation φ of f such that $\hat{\varphi} \sim 1$,*
3. *for every Horn representation φ of f it holds that $\hat{\varphi} \sim 1$,*
4. *for every Horn implicate C of f it holds that $\mathrm{Cl}_f(\mathrm{Body}(C)) = U$,*
5. *f is almost anti-monotone.*

Proof sketch for Theorem 7

We take care of the case of negative clauses first.

Lemma 11. *Let $\varphi, \psi \not\sim 1$ be Horn formulas such that $\varphi \Rightarrow \psi$, and ψ contains a negative clause D. Then*

- *ψ has a φ-complement,*
- *$\hat{\varphi} \not\Rightarrow \psi$,*
- *$\mathcal{B}_\varphi(D) \not\Rightarrow D$.*

For the rest of the proof we may assume that ψ is a definite Horn formula.

The $(1 \Rightarrow 2)$ part of the proof is based on the following lemma.

Lemma 12. *Let f be a Horn function and let $D_1 = (B \to z)$ and $D_2 = (B \to u)$ be definite Horn clauses with the same body B such that $f \Rightarrow D_1$ and $f \not\Rightarrow D_2$. Then*

$$D = (B, u \to z)$$

has no f-complement.

The $(2 \Rightarrow 3)$ part of the proof is omitted. For the $(3 \Rightarrow 1)$ part of the proof, let D be a clause in ψ such that $\bigwedge_{C \in \mathcal{B}_\varphi(D)} C \not\Rightarrow D$. It can be shown that $\mathcal{A}_\varphi(D) \neq \emptyset$.

Now we can define a φ-complement of ψ. For each clause $C \in \mathcal{A}_\varphi(D)$ let

$$\chi'_C = \bigwedge_{z \in \mathrm{Body}(D)} (\mathrm{Body}(C) \to z),$$

$$\chi''_C = (\mathrm{Body}(C), \mathrm{Head}(D) \to \mathrm{Head}(C)),$$

and finally put

$$\chi = \left(\bigwedge_{C \in \mathcal{A}_\varphi(D)} \chi'_C \wedge \chi''_C \right) \wedge \left(\bigwedge_{C \in (\varphi \setminus \mathcal{A}_\varphi(D))} C \right).$$

Thus χ is formed from φ by replacing clauses $C \in \mathcal{A}_\varphi(D)$ by $\chi'_C \wedge \chi''_C$, and leaving the rest of the formula unchanged. Note that in the definition of χ''_C, if C is a negative clause then $\mathrm{Head}(C) = 0$.

Example 13. Consider $\varphi = (x \to y) \wedge z$ and $\psi = z$. Then both clauses of φ are in $\mathcal{A}_\varphi(z)$, and so the φ-complement of ψ provided by the construction (after deleting redundant clauses) is $(x, z \to y)$.

Singleton Horn extensions and $\hat{\varphi}$

A different proof of the equivalence $(1 \Leftrightarrow 2)$ in Theorem 7 provides a semantic characterization of the body building formula. The proof is based on the following lemma. It shows that $T(\hat{\varphi}) \setminus T(\varphi)$ consists of precisely the *singleton Horn extensions* of φ, i.e., of those points which can be added to the set $T(\varphi)$ maintaining the Horn property.

Lemma 14. *Let φ be a Horn formula and $a \in F(\varphi)$. Then $T(\varphi) \cup \{a\}$ is a Horn function iff $\hat{\varphi}(a) = 1$.*

The $\Leftarrow$ direction of Lemma 14 can be proved by constructing a Horn formula χ_a for $T(\varphi) \cup \{a\}$ for every truth assignment $a \in T(\hat{\varphi}) \setminus T(\varphi)$, and this, in turn, gives an alternative construction of complements.

Both constructions for the complement may increase the size of the formula by a linear factor, and it is not known whether this increase is necessary. Similar questions for DNF are studied in (Mubayi, Turán, and Zhao 2006).

Experimental results

Recently there appears to be growing interest in exploring the computational properties of belief revision methods by running experiments (Benferhat et al. 2004; Bessant et al. 2001). The results presented in this paper raise the related question of what fraction of implicates of a random Horn formula have complements. Properties of random CNF expressions, such as their phase transition from almost surely satisfiable to almost surely unsatisfiable have been, and are, much studied (Martin, Monasson, and Zecchina 2001). Similar work has also been done for random Horn formulas (Moore et al. 2007). The results indicate that the choice of the probability distribution on Horn formulas requires care.

We have considered the following probabilistic model to generate a random Horn formula. The parameters are n, m, p and q. Here n is the number of variables, m is the number of clauses and p is the fraction of definite clauses. For each definite clause we pick the head from the uniform distribution over the variables. The bodies of the clauses are generated by determining the clause length using a geometric distribution of parameter q, and then picking the right number of variables without replacement, again using uniform distribution. This model produces Horn formulas with clauses of small, but not uniformly bounded size.

After having generated a Horn formula φ, we used exhaustive testing of all implicates for having a complement. This was done by constructing the formula $\hat{\varphi}$, and checking whether a candidate clause is a consequence of $\hat{\varphi}$. By Theorem 7, these are the implicates that do *not* have a complement. Because of the exhaustive testing, we present results

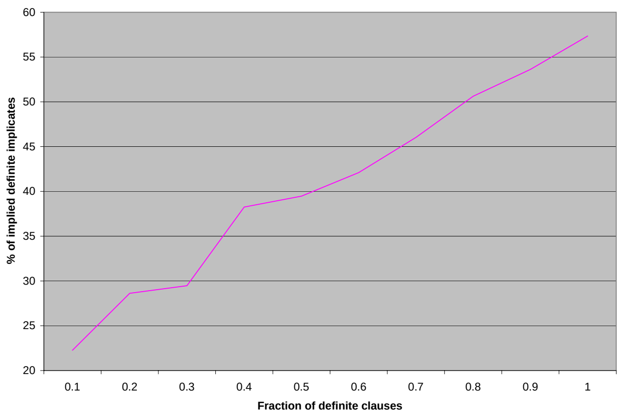

Figure 1: Percent of definite Horn implicates of φ implied by $\hat{\varphi}$ as a function of p. Measured on 100 random formulas with $n = 12$ variables, $m = 24$ clauses, and $q = 1/3$.

for 12 variables. As negative implicates always have complements, the figures show the fraction among definite implicates. For definite implicates, we want to know when we have an exact contraction operator, and when we must turn to some form of approximation or violation of some AGM postulate. The computational results suggest that in certain ranges of the parameters complements are likely to exist. For example, Figure 2 shows that for a random formula of 50 or more clauses over 12 variables (with the clause distribution given in the figure), at least 3/4 of the implicates have complements. It appears to be an interesting problem to obtain theoretical results on the fractions.

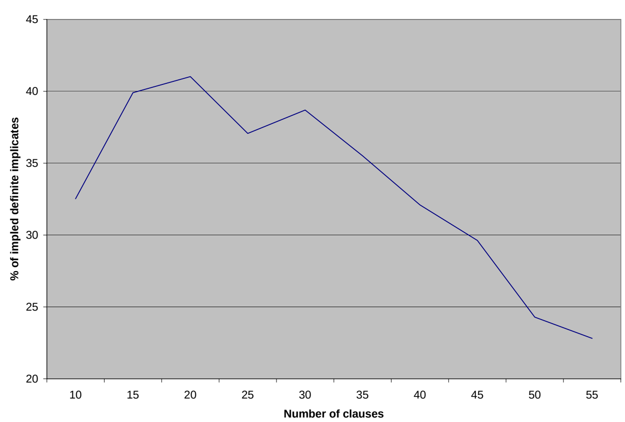

Figure 2: Percent of definite Horn implicates of φ implied by $\hat{\varphi}$ as a function of m. Measured on 100 random formulas with $n = 12$, $p = 1/2$, $q = 1/3$.

Remarks and further work

The results of this paper provide a first step towards Horn-to-Horn belief revision, but much remains to be done. It would be interesting to get a characterization of *all* complements, in the cases when a complement exists. Another next step is to investigate the two supplementary AGM postulates for contraction in the Horn case; Flouris refers to contraction operators also satisfying these postulates as *fully* AGM-compliant (Flouris 2006). The problem of extending the framework of Flouris, Plexousakis, and Antoniou to revision, as opposed to contraction, at least in the Horn case, is also open. But perhaps most importantly, one should study the possibilities of *approximating* complements in cases when a complement does not exist, giving up on adherence to the postulates in order to gain efficiency. As noted in the introduction, this would constitute a departure from current approaches to belief revision theory.

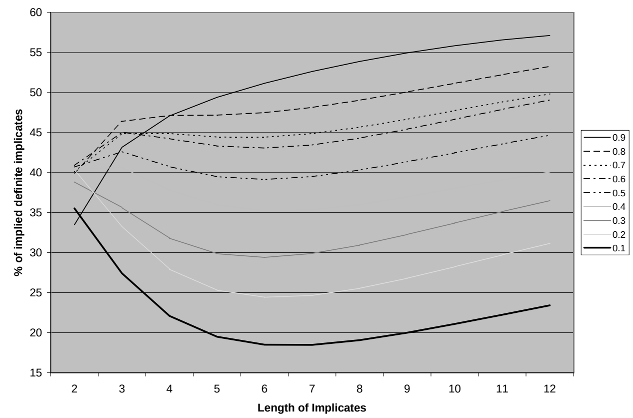

Figure 3: Percent of definite Horn implicates of φ implied by $\hat{\varphi}$, for specific implicate sizes, as a function of implicate size. Measured on 100 random formulas with $n = 12$ variables, $m = 24$ clauses, various p, and $q = 1/3$.

Our long term goal towards the construction of a commonsense reasoning agent is the integration of *belief revision and learning*. A commonsense reasoning agent should not only be able to do both, but to do both efficiently. An important motivation to study this problem is the *interactive acquisition of large commonsense knowledge bases*, such as the *Open Mind Common Sense* (Singh 2002) project. Here it seems reasonable to assume that the knowledge base receives contradictory information from the users, and thus it has to revise its contents, and at the same time should improve its quality in the long run. Also, Horn logic seems to be a reasonable knowledge representation, as inferences need to be done with the knowledge acquired. An interesting application of such knowledge bases is given by the recent work of (Pentney et al. 2007), showing that such knowledge bases could be combined with sensor data in health care and other areas.

The combination of belief revision and learning has been studied by, e.g., (Kelly, Schulte, and Hendricks 1995; Martin and Osherson 1997; Pagnucco and Rajaratnam 2005; Wrobel 1994). We plan to approach the problem of combining belief revision and learning in a formal model of computational learning theory. Due to lack of space, we do not describe the formal modeling details, but simply state the problem one may refer to as *Knowledge Base Learning (KnowBLe)*:

Problem B. Find an efficient algorithm that learns a propositional Horn formula in the model of learning from entailment (with or without queries), and updates its hypotheses in a rational manner.

Acknowledgment

This material is based upon work supported by the National Science Foundation under Grant No. CCF-0431059. The third author was also supported in part by the NKTH grant of Jedlik Ányos R&D Programme 2007 of the Hungarian government (codename TUDORKA7).

References

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

Benferhat, S.; Kaci, S.; Berre, D. L.; and Williams, M. A. 2004. Weakening conflicting information for iterated revision and knowledge integration. *Artificial Intelligence* 153:339–371.

Bessant, B.; Grégoire, E.; Marquis, P.; and Saïs, L. 2001. Interated syntax-based revision in a nonmonotonic setting. In Williams, M.-A., and Rott, H., eds., *Frontiers of Belief Revision*, volume 22 of *Applied Logic*. Kluwer. 369–391.

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

Flouris, G.; Plexousakis, D.; and Antoniou, G. 2004. On generalizing the AGM postulates: Preliminary results and applications. In *Proc. 10th Int. Workshop on Non-Monotonic Reasoning (NMR-04)*, 171–179.

Flouris, G.; Plexousakis, D.; and Antoniou, G. 2005. Updating description logics using the AGM theory. In *7th Int. Symp. Logical Formalizations of Commonsense Reasoning*.

Flouris, G. 2006. *On Belief Change and Ontology Evolution*. Ph.D. Dissertation, Dept. of Computer Science, University of Crete.

Gärdenfors, P. 1988. *Knowledge in Flux*. Detroit/Cambridge, Mass.: Bradford Books/MIT Press.

Gogic, G.; Papadimitriou, C. H.; and Sideri, M. 1998. Incremental recompilation of knowledge. *J. Artif. Intell. Res. (JAIR)* 8:23–27.

Hammer, P. L., and Kogan, A. 1995. Essential and redundant rules in Horn knowledge bases. In *Proc. 28th Hawaii Int. Conf. System Sciences*, volume 3, 209–218.

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

Jin, Y., and Thielscher, M. 2005. Iterated belief revision: A computational approach. In *Belief Change in Rational Agents: Perspectives from Artificial Intelligence, Philosophy, and Economics*, number 05321 in Dagstuhl Seminar Proceedings (online). Downloaded from URL `http://drops.dagstuhl.de/opus/volltexte/2005/359/` in October 2006.

Kelly, K. T.; Schulte, O.; and Hendricks, V. 1995. Reliable belief revision. In *Proc. X Int. Joint Congress for Logic, Methodology and the Philosophy of Science (1995)*. Proceedings published as *Logic and Scientific Methods*, M. Dalla Chiara et al., eds. Kluwer, 1997.

Liberatore, P. 1997. The complexity of iterated belief revision. In *Proc. Sixth Int. Conf. Database Theory (ICDT'97)*, 276–290.

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

Martin, E., and Osherson, D. N. 1997. Scientific discovery based on belief revision. *The Journal of Symbolic Logic* 62(4):1352–1370.

Martin, O. C.; Monasson, R.; and Zecchina, R. 2001. Statistical mechanics methods and phase transitions in optimization problems. *Theoret. Comput. Sci.* 265:3–67.

Moore, C.; Istrate, G.; Demopoulos, D.; and Vardi, M. Y. 2007. A continuous–discontinuous second-order transition in the satisfiability of random Horn-SAT formulas. *Random Struct. Algorithms* 31(2):173–185.

Mubayi, D.; Turán, G.; and Zhao, Y. 2006. The DNF exception problem. *Theoret. Comput. Sci.* 352:85–96.

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.

Pagnucco, M., and Rajaratnam, D. 2005. Inverse resolution as belief change. In *Proc. 19th Int. Joint Conf. on Artificial Intelligence (IJCAI)*, 540–545.

Pentney, W.; Philipose, M.; Bilmes, J.; and Kautz, H. 2007. Learning large scale common sense models of everyday life. In *Proc. 27th Annual Conf. AAAI*, 465–470.

Singh, P. 2002. The public acquisition of commonsense knowledge. In *Proc. AAAI Spring Symposium on Acquiring (and Using) Linguistic (and World) Knowledge for Information Access*.

Wrobel, S. 1994. *Concept Formation and Knowledge Revision*. Kluwer.