Compact Representations of Revision of Horn Clauses

Paolo Liberatore
Università di Roma "La Sapienza" - Dipartimento di Informatica e Sistemistica
via Salaria 113 - Roma, Italy
email: liberato@assi.dis.uniroma1.it

Abstract

Several methods have been proposed, in the past years, as an attempt to define how to deal with a dynamically-changing scenario. From a computational point of view, different definitions lead to different computational properties, namely to a complexity between Δ_2^p and Δ_3^p in the polynomial hierarchy.

When we restrict the analysis to cases where some operations become tractable (Horn formulas), other problems arise. We show that, providing some assumption on the polynomial hierarchy, many update operators are uncompactable, i.e. it is impossible to find a representation of the result without an exponential explosion of its size. We supply also the explicit representations of the updates for which a small representation is possible.

1 Introduction

The importance of the dynamic treatment of information in artificial intelligence, databases and philosophy leads many researchers to focus their attention on its computational aspects.

From a computational viewpoint, the definitions given result to be at the second level of the polynomial hierarchy, namely Π_2^p or almost. Although this can be viewed as a discouraging result, it ought to be compared with the implicit cost of using the knowledge base.

Practical systems use a knowledge base T as a formalization of the "state of affairs" of the world; the most important problem is to decide which facts derive from this piece of information. The problem of deciding whether a formula Q is implied by T is in general a coNP complete problem, thus intractable. For a practical use, we have to restrict our attention to special cases in which derivation becomes tractable i.e. there exists an efficient algorithm. The most used limitation is to consider only Horn formulas.

The problem we address in this paper is the following: given a set T and a formula P, both Horn, whose size are $|T|$ and $|P|$ respectively, is it always possible to represent the revised knowledge base (k.b.) with a propositional formula whose size is polynomial w.r.t. $|T| + |P|$?

Clearly, if this does not happen, it can be considered impossible, from a practical point of view (limitation of computers memory), to represent the revised knowledge with a propositional formula that can be effectively stored.

This paper is organized as follows: in the next section we remind definitions of revision given in the literature. In section 3 we prove that several methods proposed are uncompactable, while in section 4 we show the explicit polynomial-size representation of the others.

2 Preliminaries

Let $X = \{x_1, \ldots, x_n\}$ be a set of propositional variables. Given a variable x_i, a literal is a formula x_i or $\neg x_i$. We denote with l_i a literal corresponding to the variable x_i (i.e. either x_i or $\neg x_i$).

A clause is a disjunction of literals: $C = l_{i_1} \vee \cdots \vee l_{i_m}$. A Horn clause is a clause with at most one positive literal. A Horn formula is a conjunction of Horn clauses. It is well-known [DG84] that restricting our attention to Horn formulas only, problems such satisfiability, unsatisfiability and implication can be resolved by polynomial-time algorithms. This means that given two Horn formulas, deciding if they are satisfiable, or if one implies the other can be done by an efficient algorithm. A 3CNF formula is a conjunction of clauses, each composed of at most three literals. Deciding the satisfiability of a 3CNF formula is an NP-complete problem.

Let T and P be propositional formulas. An interpretation is a truth assignment to the variables, i.e. a function from X to $\{\mathit{true}, \mathit{false}\}$. We denote a model with the set of its positive variables. An interpretation M is a model of T (and is denoted with $M \models T$) if it satisfies T. We use $\mathit{mod}(T)$ to denote the set of the models of T i.e. $\mathit{mod}(T) = \{M \mid M \models T\}$. Given a set of models A, the (multi-valued) function $form(A)$ gives a formula whose set of models is A.

Consider T as a piece of information that we consider true until (at least) time t_1. Suppose that we know (we observe) a fact represented by the formula P, at time t_2. What should be our knowledge after t_2? If T and P are mutually consistent (that is $T \cup \{P\}$ is satisfiable), we can simply assume $T \cup \{P\}$ as our new knowledge base (we can put together the old k.b. and the new observation). However, it is possible that $T \cup \{P\}$ is inconsistent. In such cases, we have to specify how T and P concur to create the new knowledge base.

We briefly remind revision and update operators proposed in the last years. For a survey of several methods see [KM91]. Relations between these operators and the well-known non-monotonic operator of circumscription have been investigated in [LS95].

2.1 Syntax Model-based Revisions

Full meet revision. It is possible to define a simple revision satisfying the AGM postulates (see [KM91]) as $T * P = T \cup P$ is $T \cup P$ is satisfiable, $T * P = P$ otherwise.

Ginsberg revision. Let $W(T,P)$ be the set of the maximal subsets of T consistent with P, that is $W(T,P) = max_{\subseteq}(\{T' \subseteq T \mid T \cup \{P\} \not\vdash \bot\})$. Ginsberg revision is defined as a skeptical union of these subsets: $T *_G P = \vee(W(T,P) \cup \{P\})$.
WIDTIO. (When In Doubt, Throw It Out) is very similar: $T *_{Wit} P = \cap W(T,P) \cup \{P\}$.

2.2 Existential Model-based Revisions.

Although these revisions had been defined in a different form by each author, it is possible to reformulate them with a single expression.

$$mod(T *_L P) = \begin{cases} mod(T \cup \{P\}) & \text{if non-empty} \\ \{J \in mod(P) \mid \exists I \in mod(T) \ . \ \langle I, J\rangle \in min_{\leq_L}(mod(T) \times mod(P))\} \end{cases}$$

that is, J is a model of the result if and only if there exists in $mod(T)$ a model I such that I and J are sufficiently close each other, where $\leq_L$ measures the closeness between I and J.

Let $diff(I,J)$ be the set of the variables whose value is different between I and J. Dalal's, Forbus', Borgida's and Satoh's revisions can be expressed by suitable definitions of $\leq_L$.

Dalal	$\langle I,J\rangle \leq_D \langle L,M\rangle$	iff	$\|diff(I,J)\| \leq \|diff(L,M)\|$
Satoh	$\langle I,J\rangle \leq_S \langle L,M\rangle$	iff	$diff(I,J) \subseteq diff(L,M)$
Forbus	$\langle I,J\rangle \leq_F \langle L,M\rangle$	iff	$I = L$ and $\|diff(I,J)\| \leq \|diff(L,M)\|$
Borgida	$\langle I,J\rangle \leq_B \langle L,M\rangle$	iff	$I = L$ and $diff(I,J) \subseteq diff(L,M)$

2.3 Ignored Variables Revisions

When a new piece of information P leads us to consider T wrong, we may suppose that T was wrong because of an error during the "observation" of a certain set of variables $\Omega \subseteq X$. Such an assumption can be formalized as

$$T *_\Omega P = \{J \in mod(P) \mid \exists I \in mod(T) \ . \ I - \Omega = J - \Omega\}$$

Standard Semantic Update. Simply define Ω as the set of the variables of P.
Weber's update. The set Ω is defined as $\cup\{diff(I,J) \mid \langle I,J\rangle$ is minimal w.r.t. $\leq_S \}$.

3 What do we mean with "compactable"?

The problem we address in this paper, as said above, is the following: given a Horn set T and a Horn formula P, is it possible to express $T * P$ with a formula whose size is polynomial in the size of the inputs T and P?

More precisely, we want to determine if there is a formula T_1, equivalent to $T * P$, such that the size of T_1 is polynomial w.r.t. the $|T| + |P|$. We must be more precise on what we mean with "equivalent", since there are two kinds of equivalence.

Logical Equivalence $T_1 = T * P$

Query Equivalence $T_1 \models Q$ iff $T * P \models Q$

The difference depends on the use of new variables: using new variables, the size of a formula can be consistently reduced. For example, the formula

$$T_1 = [(a \vee b) \vee (c \wedge \neg(a \vee b))] \wedge \neg(d \wedge \neg(a \vee b))$$

can be rewritten using the observation that $a \vee b$ occur several times:

$$T_2 = (e = a \vee b) \wedge ([e \vee (c \wedge \neg e)] \wedge \neg(d \wedge \neg e)$$

(we impose e to be a renaming of $a \vee b$, and then use e instead of any occurence of $a \vee b$).

Note that T_1 is not logically equivalent to T_2 because the model $\{\neg e, a, b\}$ satisfies T_1 but not T_2. However, if Q does not contains e, then $T_1 \models Q$ iff $T_2 \models Q$.

Criterion (3) states that T_2 is built over the same alphabet of T_1 (ie there are no new variable in it). Criterion (3) is a weaker condition, since the formula T_2 can contain new variables in order to make the formula T_2 shorter.

Unfortunately, up to now, there is no method to prove that an operator does not allow a "short" representation wrt criterion (3).

About criterion (3), we note in the example that the new variable e is used to represent a subformula, that is $e = a \vee b$ imposes to e a value determined by the value of a and b. When the introduction of a new variable is always constrained by a formula like

$$new\ variable = formula(old\ variables)$$

we say that we have a circuit.

It can be proved that removing this constraint on the introduction of new variables we obtain a more powerful method of representation: there are operators that cannot admit a polynomial-sized circuit representation, but adding new variables without constraints such a representation exists.

The details of the method of proving that there is no polynomial-size circuit can be found in [KS92, CDS95, CDLS95], and it is based on a polynomial reduction from an NP complete problem. We use the concept of P/poly algorithm. Following [Joh90] we define

Definition 1 *A P/poly algorithm is a pair (A, R) where A is a function from strings to strings such that $A(s)$ has a size polynomial in $|s|$, and R is a polinomial time algorithm that decide if a string s is in a language L using s and $A(|s|)$.*

Furthermore, P/poly is the class of decision problems that can be solved by a P/poly algorithm.

4 Uncompactable Cases

Due to the lack of space, we show the proof of uncompactability for Borgida's operator only. See the table at the end of the paper for the compactability results of the other operators.

Theorem 1 *It is impossible (unless* $\Pi_3^p = \Sigma_3^p$ *=PH) to find a circuit* T' *that is equivalent to* $T *_B P$*, whose size is polynomial w.r.t.* $|T| + |P|$*.*

Proof. We prove the claim in two steps.

1. We show that given a 3CNF formula Π, it is possible to find in polynomial time two formula $T_{|\Pi|}$ and P_Π and a model M_Π, such that

$$\Pi \text{ satisfiable } \Leftrightarrow M_\Pi \models T_{|\Pi|} *_B P_{|\Pi|}$$

   Notice that $T_{|\Pi|}$ and P_Π depend only on the size of Π rather than Π itself.

2. We prove that if a compact representation existed, then it would be NP $\subseteq$ P/poly, that implies $\Pi_3^p = \Sigma_3^p =$ PH. This is considered very unlikely by complexity researchers.

First step. Borgida's revision, when $T \cup \{P\} \models \perp$, can be defined as follows:

$$mod(T *_B P) = \bigcup_{I \in mod(T)} mod(form(\{I\}) *_B P)$$

where $mod(form(I) *_B P)$ is the set of the models $J \in mod(P)$ whose symmetric difference $(I - J) \cup (J - I)$ is minimal w.r.t. set inclusion $\subseteq$.

Given a model $I \in mod(T_n)$, its closest models of P_n have the same value for $\{c_i\}$, since P_n does not contain any c_i. Hence, the model M_Π can be in the result if and only if there is a model $I \in mod(T_n)$, with the same value for $\{c_i\}$, and such that M_Π is one of its closest models of P_n.

If Π is unsatisfiable, all the models of T_n with c_i iff $\gamma_i \in \Pi$ have at least one i such that $x_i = y_i = false$. Their closest models of P_n have $x_i = y_i = false$ too.

On the contrary, if Π is satisfiable, there is in T_n a model I such that $\{c_i\}$ corresponds to the clauses of Π , and $\wedge(x_i \neq y_i)$ is true. The difference between I and M_Π is the set of the negative variables $X \cup Y$ of I; thus I and M_Π differ exactly on n variables. The model M_Π would not be in the result only if there was a model M' closer to I than M_Π. But such a M' would be at most $n-1$ far from I, so it must have at least one i such that $x_i \neq y_i$.

Second step. Suppose that a compact representation exists. Then, given Π it is possible to decide whether it is satisfiable by evalutating $M_\Pi \models T_{|\Pi|} *_B P_\Pi$.

Then, we can resolve this problem using an oracle that computes a T' equivalent to $T_{|\Pi|} *_B P_\Pi$, and a polynomial time algorithm that computes $M_\Pi \models T'$. This is a P/poly algorithm to resolve an NP complete problem thus NP $\subseteq$ P/poly. □

In a similar way, it is possible to prove that there exist proofs of NP hardness in the form Π satisfiable $\Leftrightarrow M_\Pi \models T_n * P_n$, where $*$ is either $*_S$ or $*_D$. These results allows us to conclude that Dalal's, Borgida's and Satoh's revisions cannot be compactly represented with a circuit (and thus neither with a formula over the same alphabet).

For Forbus' revision the result is stronger: a polynomial-size representation does not exist even if we introduce new free variables (without constraints).

5 Compactable cases

In this section we show which revisions are compactable, and how it can be done. First of all, the trivial results: since $T *_{FM} P \subseteq T \cup \{P\}$ and $T *_{Wit} P \subseteq T \cup \{P\}$, we have that these two revisions are compactable.

From an analysis of the general (non-Horn) case, follows that Dalal's, Weber's and SSU operators are compactable with new variables (without constraints). We show here that this property holds, in the Horn case, for Weber's and SSU operators, even if we limit our attention to equivalence with a circuit. Furthermore, Ginsberg's revision is always compactable, while the same happens for all existential revisions (but Forbus') only introducing new free variables.

We need first a result about circuits. Consider two arrays of variables $B = \{b^i_j\}$ and $H = \{h^i_j\}$. We can use a set of variables like this to store a set of clauses in this way: Let $T = \{C_1, \ldots, C_m\}$: the value of B and H will be

$$\begin{aligned} b^i_j = \mathit{true} \quad &\text{iff} \quad \neg x_j \vdash C_i \\ h^i_j = \mathit{true} \quad &\text{iff} \quad x_j \vdash C_i \end{aligned}$$

Roughly speaking, each row b^i is the characteristic array of the set of negative literals in C_i, while h^i is the representation of its positive literals. If C_i is a Horn clause, at most one h^j_i may be true.

Note that the *value* of B and H represent the syntactic form of the set T. Thus, any truth assignment to $B \cup C$ represents a set of clauses.

Let SAT be the function from the set of truth assignment over $B \cup H$ to $\{\mathit{true}, \mathit{false}\}$, defined as follows: SAT gives true on an assignment to $B \cup H$ if and only if the Horn set represented by the value of $B \cup H$ is satisfiable.

For example, if $\{b^1_1, h^2_1\}$ are all the variables assigned true, the function SAT gives false, since the corresponding set of clauses is $\{\neg x_1, x_1\}$. On the truth assignment whose true variables are $\{b^1_1, h^1_2, h^2_1\}$ the function SAT gives true, because the Horn set is now $\{\neg x_1 \vee x_2, x_1\}$, that is satisfiable.

Now, SAT is a function from the set of truth assignment on $B \cup H$ to $\{true, false\}$, thus it can be viewed as a propositional formula over the alphabet $B \cup H$. One can wonder if it is representable with a formula whose size is polynomial w.r.t. $|B| + |H|$.

Theorem 2 *It is possible to write a circuit* SAT(B, H) *of size polynomial w.r.t.* $|B| + |H|$, *whose value is true if and only if the Horn set represented by* $B \cup H$ *is satisfiable.*

The ignored variables revisions are defined using $T *_\Omega P$.

$$\mathit{mod}(T *_\Omega P) = \{J \in \mathit{mod}(P) \mid \exists I \in \mathit{mod}(T) \ . \ I - \Omega = J - \Omega\}$$

Given a set Ω, it is easy to find a representation with new variables: let $T *_\Omega P = T[\Omega/Y] \wedge P$ where $T[\Omega/Y]$ denotes the formula obtained from T by replacing every $x_i \in \Omega$ with a corresponding new variable $y_i \in Y$. Now, we prove that a compact representation exists even if we consider the equivalence w.r.t. a circuit.

Theorem 3 *Given T and P (Horn), there is a polynomial-size circuit equivalent to $T *_\Omega P$, for every set of variables Ω.*

Proof. A model of P is also a model of $T *_\Omega P$ if and only if there is in T a model that agrees with it at least on the variables not in Ω.

Given a model $J \in mod(P)$, this is equivalent to saying:

$$J \models T *_\Omega P \quad \Leftrightarrow \quad \text{the Horn set } \left[T \cup \bigcup_{x_i \notin \Omega \;\; J \models x_i} x_i \cup \bigcup_{x_i \notin \Omega \;\; J \models \neg x_i} \neg x_i \right] \text{ is satisfiable}$$

Using the polynomial-size circuit SAT, it is possible to express this definition without introducing new free variables.

Let BT and HT be two constant arrays representing the Horn set T, and B, H two $n \times n$ arrays of new variables. We want the rows $b^i h^i$ to represent the clause x_i if $J \models x_i$, the clause $\neg x_i$ otherwise. The formula REPR formalizes this idea:

$$\text{REPR}(b^i h^i, x_i) = (b^i_i = \neg x_i) \wedge (h^i_i = x_i) \wedge \bigwedge_{i \neq j} (b^i_j = h^i_j = \mathit{false})$$

where the new variables B and H can be easily determined from the old X. Hence, this is a circuit.

The whole formula is thus

$$T *_\Omega P = P \wedge \text{SAT}(BT \cup B, HT \cup H) \wedge \bigwedge_{x_i \notin \Omega} \text{REPR}(b^i h^i, x_i) \wedge \bigwedge_{x_i \in \Omega} (b^i = h^i = \mathit{true})$$

This is a polynomial-size circuit representing the result of every ignored-variables revision. □

We can give a similar theorem for Ginsberg's revision. For the existential revisions (but Forbus' one) it is possible to give representations with new free variables.

6 Conclusions

In this paper we have investigated the compactability of some revision operators, finding also compact representations where possible. Other operators have been shown uncompactable, that is a propositional formula representing the result must be of size exponential.

In order to prove uncompactability of operators, we use some polynomial-time reductions, like in the complexity analysis of languages. Nevertheless, this kind of reduction is different: it may be (see for example [CDLS95]) that two operators with the same computational complexity have different compactability properties: the complexity of an operator is not enough to characterize its compactability.

The results given in this paper can be summarized in the following table.

	circuit	new variables
WIDTIO, FM, Ginsberg	YES	YES
Weber, SSU	YES	YES
Satoh, Borgida, Dalal	NO	YES
Forbus	NO	NO

Is the revised knowledge base compressible?

References

[CDLS95] M. Cadoli, F. M. Donini, P. Liberatore, and M. Schaerf. The size of a revised knowledge base. In *Proceedings of the Fourteenth ACM SIGACT SIGMOD SIGART Symposium on Principles of Database Systems (PODS-95)*, pages 151–162, 1995.

[CDS95] M. Cadoli, F. M. Donini, and M. Schaerf. On compact representations of propositional circumscription. In *Proceedings of the Twelfth Symposium on Theoretical Aspects of Computer Science (STACS-95)*, pages 205–216, 1995. Extended version as RAP.14.95 DIS, Univ. of Roma "La Sapienza", July 1995.

[DG84] W. P. Dowling and J. H. Gallier. Linear-time algorithms for testing the satisfiability of propositional Horn formulae. *Journal of Logic Programming*, 1:267–284, 1984.

[Joh90] D. S. Johnson. A catalog of complexity classes. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A, chapter 2. Elsevier Science Publishers (North-Holland), Amsterdam, 1990.

[KM91] H. Katsuno and A. O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence Journal*, 52:263–294, 1991.

[KS92] H. A. Kautz and B. Selman. Forming concepts for fast inference. In *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92)*, pages 786–793, 1992.

[LS95] P. Liberatore and M. Schaerf. Relating belief revision and circumscription. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 1557–1563, 1995.