To appear on the Proc. of the Fourteenth ACM SIGACT SIGMOD SIGART Symposium on Principles of Database Systems (PODS-95)

The Size of a Revised Knowledge Base

Marco Cadoli **Francesco M. Donini** **Paolo Liberatore**
Marco Schaerf

Dipartimento di Informatica e Sistemistica
Università di Roma "La Sapienza"
via Salaria 113, I-00198 Roma, Italy
email: `<lastname>@assi.dis.uniroma1.it`

Abstract

In this paper we address a specific computational aspect of belief revision: The *size* of the propositional formula obtained by means of the revision of a formula with a new one. In particular, we focus on the size of the smallest formula *equivalent* to the revised knowledge base. The main result of this paper is that not all formalizations of belief revision are equal from this point of view. For some of them we show that the revised knowledge base can be expressed with a formula admitting a polynomial-space representation (we call these results "compactability" results). On the other hand we are able to prove that for other ones the revised knowledge base does not always admit a polynomial-space representation, unless the polynomial hierarchy collapses at a sufficiently low level ("non-compactability" results). The time complexity of query answering for the revised knowledge base has definitely an impact on being able to represent the result of the revision compactly. Nevertheless formalisms with the same complexity may have different compactability properties.

1 Introduction

Belief revision is a well-studied topic in the area of advanced treatment of information. It has to do with evolution of the state of a knowledge base: How does our set of beliefs change when new information arrives? Several researchers attempted at defining formalizations of belief revision, answering to the following question. Suppose we have a set of beliefs, represented by propositional theory T; at some point new information arrives, let's say propositional formula P. If T and P are inconsistent, can we still infer something reasonable from all the knowledge we have? In other words, what are the logical consequences of theory T when it is revised by formula P? In symbols, what is the set of formulae Q such that $T * P \models Q$? Proposed formalizations of belief revision can be very different in spirit, rely on different assumptions and have different goals [Bor85, Dal88, FHV83, For89, Gin86, Neb91, Sat88, Web86, Win90].

Other researchers [EG92, Neb91, Win90] focused on computational properties of belief revision. As an example, they addressed questions such as: Given formulae T, P, Q and a suitable semantics for $*$, what is the time complexity of deciding $T * P \models Q$? In which cases polynomial algorithms exist?

Both aspects of belief revision (semantic and computational) are important from the theoretical as well as the practical point of view.

In this paper we address a new specific computational aspect of belief revision: The *size* of the revised theory $T * P$. An informal description of our work follows. We insist that $T * P$ is represented as a propositional theory, i.e., we want a propositional theory T' such that

$$\{Q \mid T' \models Q\} \equiv \{Q \mid T * P \models Q\}, \tag{1}$$

Q being any formula in which only symbols of T or P occur. We call this property *query equivalence*, and we say that a T' satisfying the above criterion is *query equivalent* to $T * P$.

The reason why we insist on $T * P$ being a propositional theory is twofold. From the *epistemological* point of view, it seems reasonable that our set of beliefs keeps the format of its representation after being revised. From the *computational* point of view, it would be nice to split the task of deciding $T * P \models Q$ into two subtasks:

1. compute T' such that (1) holds;
2. decide $T' \models Q$.

There are two positive aspects in doing this: The first subtask can be done off-line, i.e., not necessarily when the query Q arrives. Moreover we could use the same set of algorithms and heuristics both for subtask 2 and for regular query answering.

A question now naturally arises: What is the size of such a T'? If the size of the smallest T' is super-polynomial in the size of T plus the size of P the above mentioned approach to query answering is clearly not practical. The main result of this paper is that not all formalizations of belief revision are equal from this point of view. For some of them (e.g., Dalal's [Dal88]) we show that T' admits a polynomial-space representation (we call these results "compactability" results). On the other hand we are able to prove that for other ones (e.g., Ginsberg's [Gin86] and Forbus' [For89]) T' does not always admit a polynomial-space representation, unless the polynomial hierarchy collapses at a sufficiently low level ("non-compactability" results). The time complexity of answering $T * P \models Q$ on-line (T, Q and P being the input) has definitely an impact on being able to represent T' compactly, although formalisms with the same complexity may have different compactability properties.

Winslett addresses this problem in [Win90] for the specific case in which the size of P is bounded by a constant, showing several compactability results. We give a complete analysis, proving that some formalisms (e.g., Ginsberg's) are not compactable even in this restricted case, while other ones (e.g., Forbus') are compactable.

A further aspect we address is the representation of a revised knowledge base using a form of equivalence characterized by the following requirement

$$T' \equiv T * P \tag{2}$$

We call this property *logical equivalence*, and we say that a T' satisfying the above criterion is *logically equivalent* to $T * P$. Notice that a T' satisfying logical equivalence (2) satisfies query equivalence (1) as well, but not the other way around. Basically, query equivalence (1) gives the possibility of introducing new propositional letters. This has definitely an impact on compactability: As an example, Dalal's formalization admits compact representations only according to criterion (1).

For what concerns non-compactability results, we use concepts such as Turing machines with advice and non-uniform complexity classes, as well as results relating uniform and non-uniform complexity classes. In fact, our results not only show unlikeliness of *propositional* representations of revised theories, but are valid for any "reasonable" (in the sense of Section 3.1) data structure. As for compactability results, we show effective procedures for obtaining compact representations.

The structure of the paper is as follows: Section 2 contains definitions about theories of belief revision and non-uniform complexity classes. Section 3 contains the analysis for the unbounded-size case; we prove compactability as well as non-compactability results. Section 4 contains the analysis for the bounded-size case, while Section 5 contains conclusions and briefly addresses future research topics.

2 Preliminaries

The *alphabet* of a propositional formula is the set of all propositional letters occurring in it. The special letter $\perp$ denotes falsity. We use $x \neq y$ as a shorthand for $(x \vee y) \wedge (\neg x \vee \neg y)$, and analogously for $x = y$. Another shorthand we use is $x \rightarrow y$ for $\neg x \vee y$. An *interpretation* of a formula is a truth assignment to the letters of its alphabet. A theory T is a set of propositional formulae. Sometimes it also denotes the formula representing the logical "and" of all its elements. A *model* M of a formula P (theory T) is an interpretation that satisfies P (all formulae in T). This is written $M \models P$ ($M \models T$). Interpretations and models of propositional formulae will be denoted as sets of letters (those which are mapped into 1). Given a propositional theory T, we denote with $\mathcal{M}(T)$ the set of its models. The expression $|s|$ denotes the size of s, in a reasonable (in the sense of [Joh90]) encoding.

In the paper we frequently use the notion of substitution of letters in a formula. Where not explicitly stated, we implicitly assume that all formulae are built over the alphabet $X = \{x_1, \ldots, x_n\}$. The notation $P[x/y]$ denotes the formula P where every occurrence of the letter x is replaced by the formula y. This notation is generalized to ordered sets: $P[X'/Y]$ denotes the formula P where all occurrences of letters in X' are replaced by the corresponding elements in Y, where X' is an ordered set of letters (in general, $X' \subseteq X$) and Y is an ordered set of formulae with the same cardinality. That is, $P[X'/Y] = P[x'_1/y_1, \cdots, x'_k/y_k]$. For example, let Q be the formula $x_1 \wedge (x_2 \vee \neg x_3)$. Let $X' = \{x_1, x_3\}$ and $Y = \{y_1, \neg y_3\}$. Then, $Q[X'/Y]$ is the formula $y_1 \wedge (x_2 \vee \neg\neg y_3)$.

2.1 Non-uniform complexity classes

As already pointed out, our proofs use the notion of non-uniform computation. We assume the reader is familiar with (uniform) classes of the polynomial hierarchy, and we just briefly introduce non-uniform classes, following Johnson [Joh90].

Definition 1 *An* advice-taking Turing machine *is a Turing machine that has associated with it a special "advice oracle" A, which can be any function (not necessarily a recursive one). On input s, a special "advice tape" is automatically loaded with $A(|s|)$ and from then on the computation proceeds as normal, based on the two inputs, x and $A(|s|)$.*

Note that the advice is only function of the *size* of the input, not of the input itself.

Definition 2 *An advice-taking Turing machine uses* polynomial *advice if its advice oracle A satisfies $|A(n)| \leq p(n)$ for some fixed polynomial p and all nonnegative integers n.*

Definition 3 *If $\mathcal{C}$ is a class of languages defined in terms of resource-bounded Turing machines, then $\mathcal{C}$/poly is the class of languages defined by Turing machines with the same resource bounds but augmented by polynomial advice.*

Any class $\mathcal{C}$/poly is also known as non-uniform $\mathcal{C}$, where non-uniformity is due to the presence of the advice. Non-uniform and uniform complexity classes are related in [KL80, Yap83].

2.2 Revision operators

We now recall the different approaches to revision and update, classifying them into formula-based and model-based ones. A more thorough exposition can be found in [EG92]. We use the following conventions: the expression card(S) denotes the cardinality of a set S, and symmetric difference between two sets S_1, S_2 is denoted by $S_1 \Delta S_2$. If S is a set of sets, $\cap S$ denotes the set formed intersecting all sets of S, and analogously $\cup S$ for union; $\min_\subseteq S$ denotes the subset of S containing only the minimal (wrt set inclusion) sets in S. The boldface name prefixed to each approach will be used for further reference.

2.2.1 Formula-based approaches

These revisions operate on the formulae syntactically appearing in the theory T. Let $W(T,P)$ be the set of maximal subsets of T which are consistent with the revising formula P:

$$\begin{aligned} W(T,P) \quad = \quad & \{T' \subseteq T \mid T' \cup \{P\} \not\models \bot, \\ & \not\exists U : T' \subset U \subseteq T, U \cup \{P\} \not\models \bot\} \end{aligned}$$

Ginsberg. In [FHV83] and in [Gin86], the revised theory is defined as a set of theories: $T *_G P \doteq \{T' \cup \{P\} \mid T' \in W(T,P)\}$. Each theory of this set has been called "world" by Ginsberg, with no reference to possible worlds in modal logics. Logical consequence in the revised theory is defined as logical consequence in each of the theories, i.e. $T *_G P \models Q$ iff for all $T' \in W(T,P)$, $T' \cup \{P\} \models Q$. If a theory T' is also viewed as the conjunction of its formulae $\wedge T'$, this consequence relation corresponds to consequence from the disjunction of all theories. Hence, we will often write $T *_G P$ as $(\bigvee_{T' \in W(T,P)} (\wedge T')) \wedge P$.

Nebel. The operator $*_N$, proposed in [Neb91], is an extension of Ginsberg's one, in which the new theories are built based on a prioritized partition of the formulae in T.

WIDTIO. Since there may be exponentially many new theories in $T *_G P$, a simpler (but somewhat drastical) approach is the so-called WIDTIO (When In Doubt Throw It Out), which is defined as $T *_{Wid} P \doteq (\cap W(T,P)) \cup \{P\}$.

Note that formula-based approaches are sensitive to the syntactic form of the theory. That is, the revision with the same formula P of two logically equivalent theories T_1 and T_2, may yield different results, depending on the syntactic form of T_1 and T_2. We illustrate this fact through an example.

Example. Consider $T_1 = \{a, b\}$, $T_2 = \{a, a \rightarrow b\}$ and $P = \neg b$. Clearly, T_1 is equivalent to T_2. The only maximal subset of T_1 consistent with P is $\{a\}$, while there are two maximal consistent subsets of T_2, that are $\{a\}$ and $\{a \rightarrow b\}$.

Thus, $T_1 *_G P = \{a, \neg b\}$ while $T_2 *_G P = \{a \vee (a \rightarrow b), \neg b\} = \{\neg b\}$. The WITDIO revision gives the same results.

2.2.2 Model-based approaches

These revisions obey the principle of irrelevance of syntax [Dal88]. They operate by selecting the models of P on the basis of some notion of proximity to the models of T. We distinguish between pointwise proximity and global proximity.

Approaches in which proximity between models of P and models of T is computed pointwise wrt each model of T were proposed as suitable for knowledge update [KM91]. Let M be a model of T; we define $\mu(M,P)$ as the set containing the minimal differences (wrt set inclusion) between each model of P and the given M; more formally, $\mu(M,P) \doteq \min_\subseteq \{M \Delta N \mid N \in \mathcal{M}(P)\}$.

Winslett. In [Win90], Winslett defines the models of the updated theory as $\mathcal{M}(T *_{Win} P) \doteq \{N \in \mathcal{M}(P) \mid \exists M \in \mathcal{M}(T) : M \Delta N \in \mu(M,P)\}$. Borgida's operator $*_B$ is the same as Winslett's, except in the case when P is consistent with T, in which case Borgida's revised theory is just $T \cup \{P\}$.

Forbus. This approach [For89] takes into account cardinality: let $k_{M,P}$ be the minimum cardinality of sets in $\mu(M,P)$. The models of Forbus' updated theory are $\mathcal{M}(T *_F P) \doteq \{N \in \mathcal{M}(P) \mid \exists M \in \mathcal{M}(T) : \mathrm{card}(M \Delta N) = k_{M,P}\}$. Note that by means of cardinality, Forbus can compare (and discard) models which are incomparable in Winslett's approach.

We now recall approaches where proximity between models of P and models of T is defined considering globally *all* models of T. Let $\delta(T,P) \doteq \min_\subseteq \bigcup_{M \in \mathcal{M}(T)} \mu(M,P)$.

Satoh. In [Sat88], the models of a revised theory are defined as $\mathcal{M}(T *_S P) \doteq \{N \in \mathcal{M}(P) \mid \exists M \in \mathcal{M}(T) : N \Delta M \in \delta(T,P)\}$.

Dalal. This approach is similar to Forbus', but global. Let $k_{T,P}$ be the minimum *cardinality* of sets in $\delta(T,P)$; in [Dal88], Dalal defines the models of a revised theory as $\mathcal{M}(T *_D P) \doteq \{N \in \mathcal{M}(P) \mid \exists M \in \mathcal{M}(T) : \mathrm{card}(N \Delta M) = k_{T,P}\}$.

Weber. Let $\Omega \doteq \cup\delta(T,P)$, i.e. Ω contains every letter appearing in at least one minimal difference between a model of T and a model of P. Following [EG92], the models of Weber's revised theory [Web86] are defined as $\mathcal{M}(T *_{Web} P) \doteq \{N \in \mathcal{M}(P) \mid \exists M \in \mathcal{M}(T) : N \Delta M \subseteq \Omega\}$.

To illustrate the differences between different model-based approaches, consider the following example.

Example. Let T and P be defined as:

$$
\begin{aligned}
T &= \{a, b, c\} \\
P &= (\neg a \wedge \neg b \wedge \neg d) \vee (\neg d \wedge \neg b \wedge (a \neq d))
\end{aligned}
$$

Note that T has only two models, which are:

$$
\begin{aligned}
J_1 &= \{a, b, c, d\} \\
J_2 &= \{a, b, c\}
\end{aligned}
$$

while P has four models:

$$
\begin{aligned}
I_1 &= \{a, b\} \\
I_2 &= \{c\} \\
I_3 &= \{b, d\} \\
I_4 &= \emptyset
\end{aligned}
$$

The set differences between each model of T and each model of P are:

Δ	$I_1 = \{a, b\}$	$I_2 = \{c\}$	$I_3 = \{b, d\}$	$I_4 = \emptyset$
J_1	$\{c, d\}$	$\{a, b, d\}$	$\{a, c\}$	$\{a, b, c, d\}$
J_2	$\{c\}$	$\{a, b\}$	$\{a, c, d\}$	$\{a, b, c\}$

Hence, the minimal differences between J_1 and models of P are $\mu(J_1, P) = \{\{c, d\}, \{a, b, d\}, \{a, c\}\}$; the minimal differences between J_2 and models of P are $\mu(J_2, P) = \{\{c\}, \{a, b\}\}$.

The cardinalities of set differences between each model of T and each model of P are:

	$I_1 = \{a, b\}$	$I_2 = \{c\}$	$I_3 = \{b, d\}$	$I_4 = \emptyset$
J_1	2	3	2	4
J_2	1	2	3	3

Winslett. The minimal differences in $\mu(J_1, P)$ correspond to the models I_1, I_2, I_3 of P, while those in $\mu(J_2, P)$ correspond to the models I_1, I_2 of P. Therefore, the models of $T *_{Win} P$ are I_1, I_2, I_3. The same result holds for Borgida's revision, since T and P are inconsistent.

Forbus. From the table with cardinalities: the minimal cardinality of differences between J_1 and each model of P is $k_{J_1,P} = 2$, corresponding to models I_1 and I_3; while $k_{J_2,P} = 1$, corresponding to I_1. Therefore, $T *_F P$ has models I_1 and I_3.

We now turn to global proximity approaches, where also entries in different rows of the above tables are compared for minimality.

Satoh. The minimal differences between any model of T and any model of P are $\delta(T, P) = \{\{c\}, \{a, b\}\}$. These minimal differences correspond to models I_1 and I_2 of P, which therefore are the models of $T *_S P$.

Dalal. The minimum cardinality of all set differences is $k_{T,P} = 1$, corresponding to I_1. As a result, $T *_D P$ selects the model I_1 only.

Weber. Consider the union of all minimal global differences, that is $\Omega = \cup\delta(T, P) = \{a, b, c\}$. In Weber's revision, one selects the models of P for which there exists a model of T whose difference is included in Ω. Since all models of P have this property, they are all selected. Thus, the revision coincides with P in this case.

The complexity of deciding $T * P \models Q$ (T, P and Q being the input) was studied in [EG92]: in Dalal's approach, the problem is $\Delta_2^p[\log n]$-complete, while in all other approaches it is Π_2^p-hard (in some cases, Π_2^p-complete).

3 General Case

The purpose of this section is to show an analysis on the *size* of formulae T' such that either criterion (1) or (2) holds. We consider revision operators mentioned in Section 2, and show both compactability and non-compactability results. There is no assumption on the size of the incoming formula P; the bounded-size case will be addressed in the next section. Notice that WIDTIO semantics always admits a compact representation according to both criteria, since it immediately follows from its definition that the size of $T *_{Wid} P$ is always less than or equal to $|T| + |P|$.

3.1 Query equivalence

In this subsection we investigate the size of a propositional representation of the result of revising a knowledge base T with a new formula P, when the representation satisfies the query-equivalence criterion (1).

We begin our investigation focusing on Ginsberg's operator. Other researchers have already noticed that the explicit representation of the result of revising a knowledge base under Ginsberg's semantics might lead to some difficulties.

We introduce this problem with an example presented by Nebel in [Neb94]. Let

$$T = \{x_1, \ldots, x_m, y_1, \ldots, y_m\}$$

$$P = \bigwedge_{i=1}^{m} (x_i \neq y_i)$$

Applying Ginsberg's revision method we have that $W(T, P)$ contains 2^m distinct theories, each one containing, for each i, exactly one of x_i and y_i. If we represent $T *_G P$ as the disjunction of all theories in $W(T, P)$, conjoined with P, the size of this representation is exponential in $|T| + |P|$.

The problem of the explosion of the size of the revised knowledge base has also been pointed out by Winslett in [Win90], where she shows another example:

$$
\begin{array}{lcllll}
T & = & \{ \; x_1, & y_1, & z_1 = (\neg x_1 \vee \neg y_1), & \\
 & & \quad \vdots & \vdots & \vdots & \\
 & & \quad x_i, & y_i, & z_i = (z_{i-1} \wedge (\neg x_i \vee \neg y_i)), & \\
 & & \quad \vdots & \vdots & \vdots & \\
 & & \quad x_m, & y_m, & z_m = (z_{m-1} \wedge (\neg x_m \vee \neg y_m)) \; \} & \\
P & = & \quad z_m & & &
\end{array}
$$

Again, the cardinality of the set $W(T,P)$ is exponential in m. These two examples show that the "obvious" representation of $T *_G P$ might have size exponential in $|T| + |P|$. However, they do not rule out the existence of a different representation of polynomial size. As Winslett notes in [Win90, pg. 34] the exponential increase in the size is proven if we "assume a completely naive storage organization, where the theories are written out as we would write them down on paper". Later on she also conjectures that "these bounds hold even for clever storage schemes".

We now show that her conjecture on Ginsberg's revision is indeed true. This result will later be generalized to the other revision operators mentioned above. In order to achieve this result on Ginsberg's revision operator ($*_G$) we first prove a stronger one based on the following idea:

Let L be the alphabet of T and P; by doing some off-line computation, we want to find a data structure D with the following characteristics:

1. for some fixed polynomial p, the *size* of D is bounded by $p(|T| + |P|)$;

2. there exists a relation $ASK(\cdot,\cdot)$, such that given any clause Q on the alphabet L, $ASK(D,Q)$ is true iff $T *_G P \models Q$;

3. deciding the relation $ASK(\cdot,\cdot)$ is a problem in coNP, where the inputs are its arguments.

Intuitively, this means that we are trying to "compile" $T *_G P$ in such a way that deciding $T *_G P \models Q$ (a Π_2^p-complete problem) becomes a problem in coNP. One way to do that would be to rewrite $T *_G P$ into an equivalent propositional formula T' of size polynomial in $|T| + |P|$, where ASK corresponds to classical consequence relation, i.e. $ASK(T', Q) = true$ iff $T' \models Q$.

We are able to show that it is very unlikely that such a data structure D may exist. To do that we resort on the notion of non-uniform computation. In what follows a relation R such that deciding R is a problem in coNP will be called coNP-relation.

Theorem 1 *Suppose there exists a polynomial p such that given any revised knowledge base $T *_G P$, there exists a data structure $D_{T,P}$ and coNP-relation $ASK(\cdot,\cdot)$ such that $|D_{T,P}| < p(|T| + |P|)$, and for any clause Q, $T *_G P \models Q$ iff $ASK(D_{T,P}, Q)$ is true. Then NP $\subseteq$ coNP/poly.*

Proof. Since the proof is rather long, we first give an outline to improve its readability. The proof consists of the following steps:

1. Choice of a NP-complete problem π;

2. Showing that for any integer n there exists a pair of CNF formulae (T_n, P_n) (depending only on n and of size $O(n^3)$) such that for any instance F of π there exists a clause Q_F such that $T_n *_G P_n \models Q_F$ iff the answer to F is "yes".

3. Showing that if for each pair (T,P) there exists a D with the properties required in the statement of the theorem, then NP is contained in coNP/poly.

Step 1: A complete problem (wrt many-one transformations) for the class NP is 3-SAT. Let F be a 3CNF formula with $|F| = n$. We observe that the number of propositional letters contained in F is at most equal to n. We choose an "oversized" set $X = \{x_1, \ldots, x_n\}$ as the set of letters, even if F uses only a proper subset of it. In such a way we have $|F| = \mathrm{card}(X)$.

Step 2: We show that for any integer n, there exists a formula P_n and a set of atomic facts T_n, both depending only on n, of polynomial size wrt n, such that given any set of 3CNF clauses F, there exists a query Q_F such that F is satisfiable if and only if $T_n *_G P_n \models Q_F$.

Let L be the alphabet $X \cup C \cup D \cup \{r\}$, where C is a set of new atoms one-to-one with possible three-literals clauses of X, i.e., $C = \{c_i \mid \gamma_i \text{ is a three-literals clause of } X\}$, D is a set of new atoms one-to-one with atoms in C and r is a new distinct atom.

We define T_n and P_n on the alphabet L according to the following rules:

$$T_n = C \cup D \cup X \cup \{r\} \tag{3}$$

$$P_n = [(\bigwedge\{\neg x_i \mid x_i \in X\} \wedge \neg r) \vee \bigwedge(c_i \rightarrow \gamma_i)] \wedge \bigwedge(c_i \neq d_i) \tag{4}$$

Given a generic 3CNF formula F, we define Q_F as follows:

$$Q_F = \bigwedge(\{c_i \mid \gamma_i \in F\} \cup \{d_i \mid \gamma_i \notin F\}) \rightarrow r \tag{5}$$

Given an interpretation η of the atoms of C, we can associate to it a theory W_η:

$$W_\eta = \{c_i \mid c_i \in \eta\} \cup \{d_i \mid c_i \notin \eta\}$$

Note that no c_i or d_i can be added consistently to this set, otherwise, the set $W_\eta \cup \{c_i\}$ would contain both c_i and d_i, whereas P_n imposes that $c_i \neq d_i$. Furthermore, the theory $W_\eta \cup P_n$ is consistent, because the interpretation η', that extends η to L by making true all atoms d_i such that $\eta \not\models c_i$ and false all the other atoms in L makes both W_η and P_n true.

Given a generic 3CNF formula F, in order to find maximal subsets of T_n that, when joined with P_n, do not satisfy Q_F we must focus on those satisfying $\bigwedge(\{c_i \mid \gamma_i \in F\} \cup \{d_i \mid \gamma_i \notin F\})$. Hence, all worlds not satisfying Q_F must include W_η. Thus, it could be that $W_1 \not\models Q_F$ only if $W_\eta \subseteq W_1$.

Suppose F unsatisfiable. We show that $W_\eta \in W(T_n, P_n)$ and $W_\eta \cup P_n \not\models Q_F$. For each interpretation ϕ of the atoms of X, there exists a $\gamma_i \in F$ such that $\phi \not\models \gamma_i$. Let $\psi = \eta \cup \phi$ be an interpretation on L: we have that $\psi \not\models \bigwedge(c_i \rightarrow \gamma_i)$. As a consequence, $W_\eta \cup P_n \models (\wedge\{\neg x_i \mid x_i \in X\}) \wedge \neg r \wedge \bigwedge(c_i \neq d_i)$ and, therefore, $W_\eta \in W(T_n, P_n)$. In fact, adding to it any other element of T_n would make it inconsistent with P_n. Since $W_\eta \cup P_n \not\models Q_F$, then $T_n *_G P_n \not\models Q_F$.

On the converse, suppose F satisfiable. We show that all theories in $W(T_n, P_n)$ imply Q_F. Note that all theories not including W_η imply Q_F, hence we concentrate on those containing W_η. These theories are obviously uniquely characterized by their intersection with $X \cup \{r\}$. For any interpretation ϕ of the atoms in X, let W_ϕ be defined as follows:

$$W_\phi = W_\eta \cup \{x_i \mid x_i \in \phi\} \cup \{r\}$$

If ϕ does not satisfy F then W_ϕ is inconsistent with P_n, and therefore, $W_\phi \notin W(T_n, P_n)$. On the other side, if ϕ satisfies F, the interpretation $\psi = \eta \cup \phi \cup \{r\}$ of L satisfies W_ϕ, P_n and Q_F. Since there exists at least one interpretation ϕ satisfying F, the corresponding $W_\phi \in W(T_n, P_n)$, and therefore $W_\eta \notin W(T_n, P_n)$. As a consequence, for all elements W_ϕ of $W(T_n, P_n)$ we have that $W_\phi \cup P_n \models Q_F$.

Step 3: Let us assume that there exists a polynomial p with the properties claimed in the statement of Theorem 1. Then, for each revised knowledge base $T_n *_D P_n$ there exists a data structure D_n, with $|D_n| < p(|T_n| + |P_n|)$, and a coNP-relation $ASK(\cdot,\cdot)$ such that given any query Q, $T *_G P \models Q$ iff $ASK(D, Q)$ is true. We could define an advice-taking Turing machine, solving unsatisfiability of propositional formulae in nondeterministic polynomial time, in this way: Given a generic propositional formula F, with $|F| = n$, the machine loads the advice D_n, computes Q_F, and then decides whether $ASK(D_n, Q_F)$ is *false* in nondeterministic polynomial time (since $ASK(\cdot,\cdot)$ is a coNP-relation, this machine exists by definition). Since $|T_n| + |P_n| = O(n^3)$, the advice D_n has size $O(p(n^3))$, hence we would have shown that unsatisfiability of propositional formulae is in non-uniform NP. Since unsatisfiability of propositional formulae is a coNP-complete problem, this implies coNP $\subseteq$ NP/poly. The results of Yap [Yap83] imply that coNP $\subseteq$ NP/poly if and only if NP $\subseteq$ coNP/poly, hence the claim follows. □

The above theorem shows the unfeasibility, under certain conditions, of making explicit the effect of the revision such that the new version is more effective in answering queries. Notice that no bound is imposed on the time spent in the compilation process. An immediate consequence of this result is the following corollary:

Corollary 2 *Let T and P be CNF formulae. Unless NP $\subseteq$ coNP/poly, there is no polynomial p and formula T' satisfying criterion (1), where $|T'| \leq p(|T| + |P|)$, when the revision operator is $*_G$.*

In fact, if such a T' exists, then there exists a coNP-relation (namely, logical consequence) such that, given a clause Q, $T *_G P \models Q$ can be decided by checking $T' \models Q$. In particular, the corollary implies that there is no formula T' satisfying criterion (2), whose size is polynomial in $|T| + |P|$. Since Nebel's revision operator is a generalization of Ginsberg's one, Corollary 2 also holds for $*_N$.

The above corollary completes our analysis of Ginsberg's revision operator. We now come to consider the model-based revision operators. Even though their semantics is very different from Ginsberg's one, we obtain a similar result for most of them:

Theorem 3 *Suppose there exists a polynomial p such that given any revised knowledge base $T * P$, where $*$ is one of $\{*_B, *_S, *_{Win}\}$, there exists a data structure $D_{T,P}$ and coNP-relation $ASK(\cdot,\cdot)$ such that $|D_{T,P}| < p(|T| + |P|)$, and for any clause Q, $T * P \models Q$ iff $ASK(D_{T,P}, Q)$ is true. Then NP $\subseteq$ coNP/poly.*

Proof. Eiter and Gottlob's result [EG92, Lemma 6.1] implies that $T *_G P \models Q$ iff $T *_B P \models Q$ iff $T *_S P \models Q$ iff $T *_{Win} P \models Q$ when T is a set of literals and all the letters in Q also appear in T. Note that these conditions hold for the T defined in (3) and Q_F as defined in (5). □

Using a proof similar to the one of Theorem 1, we can show that the same result also holds for Forbus' revision operator.

Theorem 4 *Suppose there exists a polynomial p such that given any revised knowledge base $T *_F P$, there exists a data structure $D_{T,P}$ and coNP-relation $ASK(\cdot,\cdot)$*

such that $|D_{T,P}| < p(|T| + |P|)$*, and for any clause* Q*,* $T *_F P \models Q$ *iff* $ASK(D_{T,P}, Q)$ *is true. Then* $NP \subseteq$ *coNP/poly.*

Proof (sketch). The proof is similar to Ginsberg's one, but we have now to consider closeness between models.

Let C be an $m \times (n+2)$ boolean matrix, where m is the number of the possible 3CNF clauses on the alphabet X. Let

$$U_n = \bigwedge_{j=1}^{m} \bigwedge_{i=2}^{n+2} (c_i^j = c_1^j)$$

This formula forces all rows of the matrix C to be equal. Let T_n and P_n be defined as follows:

$$\begin{aligned} T_n &= \{U_n\} \cup X \cup \{r\} \\ P_n &= U_n \wedge [\bigwedge_{j=1}^{m} (c_1^j \rightarrow \gamma_j) \vee (\wedge\{\neg x_i \mid x_i \in X\} \wedge \neg r)] \end{aligned}$$

Given a generic 3CNF formula E, we define:

$$Q_E = \bigvee_{i=1}^{n+2} \{c_i^j \mid \gamma_j \notin E\} \vee \bigvee_{i=1}^{n+2} \{\neg c_i^j \mid \gamma_j \in E\} \vee \bigvee X \vee r$$

We show that $T_n *_F P_n \models Q_E$ if and only if E is satisfiable. Note that any interpretation, but $M_E = \bigcup_{i=1}^{n+2} \{c_i^j \mid \gamma_j \in E\}$, satisfies Q_E. Thus, we only need to show that $M_E \models T_n *_F P_n$ if and only if E is unsatisfiable.

First of all, note that T_n has a model for each set of 3CNF clauses. We denote $I_E = \bigcup_{i=1}^{n+2} \{c_i^j \mid \gamma_j \in E\} \cup X \cup \{r\}$ that model.

Observe that the distance between M_E and I_E is exactly $n+1$, while any other $J \in \mathcal{M}(P_n)$, with a different valuation of atoms in C, is at least $n+2$ far from I_E. Hence, the models of P that are closest to I_E must have the same valuation of atoms in C.

Suppose E unsatisfiable. The closest model of I_E is now M_E, since there is no other model of P_n with the same valuation of C. Hence $M_E \models T_n *_F P_n$.

On the converse, suppose E satisfiable. Let ϕ be the interpretation of the atoms in X satisfying E. The model $J_E = \bigcup_{i=1}^{n+2} \{c_i^j \mid \gamma_j \in E\} \cup \{x_i \mid x_i \in \phi\} \cup \{r\}$ is now in $\mathcal{M}(P_n)$, and it is at most n letters far from I_E. Hence, M_E is not a model of $T_n *_F P_n$. □

The above results of incompactability of belief revision into a propositional formula are conditioned on NP not being included in coNP/poly. Using the results of Yap in [Yap83, pg. 292 and Theorem 2] it follows that if NP $\subseteq$ coNP/poly then $\Pi_3^p = \Sigma_3^p =$ PH, i.e. the polynomial hierarchy collapses at the third level. Such an event is considered very unlikely by most researchers in structural complexity.

We now turn our attention to operators that admit a compact representation under the query-equivalence criterion (1). In particular, we show that a knowledge base revised using Dalal's or Weber's operators can be expressed in a compact way, if we go for criterion (1), i.e. the initial language is extended.

Let X be the alphabet of the initial knowledge base and the revising formula, and Y be another set of (distinct) letters, one-to-one with X. Let $EXACTLY(m, X, Y, W)$ denote a formula containing letters of X and Y, and possibly other letters W, which is true iff the Hamming distance between the values assigned to X and Y is exactly m. Such a formula can be constructed in several ways. For example, first compute the number of true exclusive-or between each x_i and y_i. The binary representation of this number can be computed with a suitable adding circuit, that requires $O(n^2)$ half-adders. This circuit can be expressed as a formula using $O(n^2)$ new letters. Secondly, write the formula that equals the bits of this number with those of the binary representation of m. The revised theory $T *_D P$ can be expressed as $T[X/Y] \wedge P \wedge EXACTLY(k, X, Y, W)$, where k is the minimum distance between the models of P and T, denoted as $k_{T,P}$ in the definition of Dalal's revision.

Theorem 5 *Let* k *be the minimal distance between the models of* P *and* T*, every model of* $T[X/Y] \wedge P \wedge EXACTLY(k, X, Y, W)$ *is a model of* $T *_D P$*, and conversely for every model* M *of* $T *_D P$ *there is an extension* M' *of* M *to* Y *such that* M' *is a model of* $T[X/Y] \wedge P \wedge EXACTLY(k, X, Y, W)$*.*

Proof. If. Let M be a model of $T[X/Y] \wedge P \wedge EXACTLY(k, X, Y, W)$. Then $M \cap X$ satisfies P, and since it satisfies $EXACTLY(k, X, Y, W)$ it has a distance k from $M \cap Y$, and $M \cap Y$ satisfies $T[X/Y]$. Since k is by definition the minimal distance between a model of T and a model of P, $M \cap X$ is also a model of $T *_D P$.

Only If. Let M be a model of $T *_D P$. This model satisfies P. Let $MT(X)$ be a model of T having distance k from M. Define M' as $M \cup MT(Y)$, where $MT(Y)$ is a model interpreting each y_i as the corresponding x_i in $MT(X)$. Obviously, M' satisfies both P and $T[X/Y]$, and by definition of $MT(X)$ it can be extended to an assignment to W so that it also satisfies $EXACTLY(k, X, Y, W)$. □

From the theorem, we conclude that $T *_D P$ and $T[X/Y] \wedge P \wedge EXACTLY(k, X, Y, W)$ are equivalent wrt queries over the original alphabet X.

Corollary 6 $T[X/Y] \wedge P \wedge EXACTLY(k, X, Y, W)$ *is query equivalent (1) to* $T *_D P$*.*

Now we show how Weber's revision can be compactly represented. Winslett, in [Win90] gives a similar proof, but only if the new formula has a size bounded by a constant. Our result, instead, doesn't need such a restriction. Consider the theory T where the letters of Ω are replaced by the letters of a new set Z, one-to-one with Ω. We denote such a new theory with $T[\Omega/Z]$.

Theorem 7 *$T[\Omega/Z] \wedge P$ is query equivalent (1) to $T *_{Web} P$.*

We also note that this representation of $T *_{Web} P$ increases the size of T only by the length of P whereas the "compact" representation of $T *_D P$ requires a formula whose size is quadratic in the number of the letters.

3.2 Logical equivalence

In this section we investigate the size of revised knowledge bases satisfying logical equivalence (2). In particular, we show that Dalal's and Weber's revisions, which admit polynomial-size representations wrt query equivalence (1), are uncompactable wrt logical equivalence (2) unless the condition NP$\subseteq$ P/poly holds (which implies $\Sigma_2^p = \Pi_2^p =$PH [KL80]).

Theorem 8 *Suppose there exists a polynomial p such that given any revised knowledge base $T *_D P$, there exist a data structure $D_{T'}$ and P-relation $ASK(\cdot,\cdot)$ such that $|D_{T'}| \leq p(|T| + |P|)$, and for any interpretation M, $ASK(D_{T'}, M)$ is true iff $M \models T *_D P$. Then NP$\subseteq$ P/poly. The same holds for $*_{Web}$ in place of $*_D$.*

Proof. The proof follows the same steps as proof of Theorem 1.

Step 1: We choose the NP-complete problem 3-SAT. Let F be a 3CNF formula with $|F| = n$. We observe that the number of propositional letters contained in F is at most equal to n. We choose an "oversized" set $X = \{x_1, \ldots, x_n\}$ as the set of letters, even if F uses only a proper subset of it. In such a way we have $|F| = \text{card}(X)$.

Step 2: We show that for any integer n, there exists a formula P_n and a knowledge base T_n, both depending only on n, of polynomial size wrt n, such that given any 3CNF formula F over an alphabet of n atoms, there exists an interpretation M_F such that F is satisfiable iff M_F is a model of $T_n *_D P_n$. The proof contains implicitly a reduction showing that model checking in Dalal- and Weber-revised knowledge bases is NP-hard.

Let $Y = \{y_1, \ldots, y_n\}$ be a set of new letters in one-to-one correspondence with letters of X, and let C be a set of new letters one for each three-literals clause over X, i.e., $C = \{c_i \mid \gamma_i \text{ is a three-literals clause of } X\}$. Finally, let L be the set $X \cup Y \cup C$, where r is a distinguished new letter. Notice that $\text{card}(L) \in O(n^3)$. We define T_n as the conjunction of two formulae:

$$T_n \equiv \Delta_n \wedge \Gamma_n \tag{6}$$

The 2CNF formula Δ_n states non-equivalence between atoms in X and their correspondent in Y:

$$\Delta_n \equiv \bigwedge_{x_i \in X} (x_i \neq y_i).$$

Γ_n codes every possible 3CNF formula over X, using the atoms in C as "enabling gates". The formula is defined as:

$$\Gamma_n \equiv \bigwedge_{c_i \in C} \gamma_i \vee \neg c_i.$$

Γ_n is a 4CNF formula and it contains $O(n^3)$ clauses.

We define P_n as:

$$P_n \equiv \bigwedge_{i=1}^{n} (\neg x_i \wedge \neg y_i) \tag{7}$$

Note that the size of T_n and P_n is $O(n^3)$. Moreover, T_n and P_n do not depend on a specific 3CNF formula F, but only on the size n of its alphabet.

Given a 3CNF formula F over X, we denote $C_F = \{c_i \in C \mid \gamma_i \text{ is a clause of } F\}$. We first show that F is satisfiable iff C_F is a model of $T_n *_D P_n$, then we show that F is satisfiable iff C_F is a model of $T_n *_{Web} P_n$.

Dalal's revision

We observe that in every model of T_n exactly n atoms from $X \cup Y$ are true, while in every model of $T_n *_D P_n$ all atoms from $X \cup Y$ are false. Hence, recalling the definition of Dalal's revision, $k_{T_n,P_n} \geq n$. Moreover, $k_{T_n,P_n} = n$ since X is a model of T_n and $\emptyset$ is a model of P_n, and the cardinality of their difference is n. Now C_F is a model of P_n, hence it is also a model of $T_n *_D P_n$ iff there exists a model of T_n such that $\text{card}(C_F \Delta M) = k_{T_n,P_n} = n$.

If part. Let F be satisfiable, and let X_F be a model of F. Let $\overline{Y_F} = \{y_i \mid x_i \notin X_F\}$, and let $M = C_F \cup X_F \cup \overline{Y_F}$. We show that M is a model of T_n. In fact, M satisfies Δ_n by construction of $\overline{Y_F}$, and also M satisfies Γ_n, because for each clause $\gamma_i \vee \neg c_i$ of Γ_n, either $c_i \notin C_F$ or γ_i is satisfied by X_F. Now observe that $\text{card}(C_F \Delta M) = \text{card}(X_F \cup \overline{Y_F}) = n$. Hence, C_F is a model of $T_n *_D P_n$.

Only if part. Suppose C_F is a model of $T_n *_D P_n$. Then there exists a model M of T_n such that $\text{card}(C_F \Delta M) = n$, that is C_F and M differ on exactly n atoms. Since M satisfies Δ_n, the difference $C_F \Delta M$ contains exactly n atoms from $X \cup Y$. Hence, M and C_F agree on the truth assignment to atoms of C, that is, $M \cap C = C_F$. We claim that $M \cap X$ is a model of F. In fact, M satisfies

$\Gamma_n = \bigwedge_{c_i \in C} \gamma_i \vee \neg c_i$. Simplifying Γ_n with truth values of $M \cap C = C_F$, we conclude that M satisfies $\bigwedge_{c_i \in C_F} \gamma_i$, which is exactly formula F. Since the formula contains only atoms from X, the interpretation $M \cap X$ satisfies F, hence F is satisfiable.

We now show that F is satisfiable iff C_F is a model of $T_n *_{Web} P_n$.

Weber's revision
First observe that $\mathcal{M}(P_n) = 2^C$, i.e. every subset of C is a model of P_n. Since X and Y are models of T_n, $\mu(X, P_n) = X$, $\mu(Y, P_n) = Y$. This implies that both X and Y are in $\delta(T_n, P_n) \doteq \bigcup_{M \in \mathcal{M}(T_n)} \mu(M, P_n)$. Recall that $\Omega \doteq \cup\delta(T, P)$. Hence, $X \cup Y \subseteq \Omega$. Moreover, for every model M of T_n, $\mu(M, P_n)$ contains no atom from C because $\mathcal{M}(P_n) = 2^C$. Hence, $\Omega = X \cup Y$. From the definition of Weber's revision, C_F is a model of $T_n *_{Web} P_n$ iff there exists a model M of T_n such that $M - \Omega = C_F - \Omega$. Since $\Omega = X \cup Y$, the last condition is equivalent to $M \cap C = C_F$.

If part. Let F be satisfiable, and let X_F be a model of F. Let $\overline{Y_F} = \{y_i \mid x_i \notin X_F\}$, and let $M = C_F \cup X_F \cup \overline{Y_F}$. From the above proof for Dalal's revision, we know that M is a model of T_n. Since $M \cap C = C_F$, C_F is a model of $T_n *_{Web} P_n$.

Only if part. Suppose C_F is a model of $T_n *_{Web} P_n$. Then there exists a model M of T_n such that $M \cap C = C_F$. We claim that $M \cap X$ is a model of F. In fact, as for Dalal's revision, M satisfies $\Gamma_n = \bigwedge_{c_i \in C} \gamma_i \vee \neg c_i$. Simplifying Γ_n with truth values of $M \cap C = C_F$, we conclude that M satisfies $\bigwedge_{c_i \in C_F} \gamma_i$, which is exactly formula F. Since the formula contains only atoms from X, $M \cap X$ satisfies F, hence F is satisfiable.

Step 3: is analogous to Step 3 of Theorem 1. Let us assume that there exists a polynomial p with the properties claimed in the statement of Theorem 8. Then, for each revised knowledge base $T_n *_D P_n$ there exists a data structure D_n, with $|D_n| < p(|T_n| + |P_n|)$, and a P-relation $ASK(\cdot, \cdot)$ such that given any interpretation M for $T_n *_D P_n$, $ASK(D_n, M)$ is true iff M is a model of $T_n *_D P_n$. We can define an advice-taking Turing machine solving satisfiability of 3CNF formulae in this way: given a generic 3CNF formula F, with $|F| = n$, the machine loads the advice D_n, computes C_F, and then decides $ASK(D_n, C_F)$ in polynomial time. Since $|T_n| + |P_n| = O(n^3)$, the advice D_n has polynomial size wrt n, hence we would have shown that 3-SAT is in non-uniform P. Since 3-SAT is an NP-complete problem, this implies NP$\subseteq$ P/poly. □

Theorem 9 *Let T and P be CNF formulae. Unless NP$\subseteq$ P/poly, there is no polynomial p such that there always exists a T' satisfying criterion (2), of size less than $p(|T| + |P|)$, where $*$ is one revision defined above but WIDTIO.*

Proof. For Dalal's and Weber's revision, the claim is a corollary of the above Theorem 8. For all other operators, the claim follows from Theorems 1, 3 and the fact that the non-existence of a representation satisfying criterion (1) implies the non-existence of a representation satisfying criterion (2). □

4 Bounded revision

In the previous section we investigated the issue of the existence of compact representations of revised knowledge bases. As it turned out, for most of the operators it does not exist a compact explicit representation of the result of revising a knowledge base with a new formula. From an analysis of the proofs, it turns out that this behavior depends on the new formula being very complex. However, in database applications it is reasonable to assume that the size of the new formula is very small wrt the size of the knowledge base. In this section we investigate which impact this assumption has on the existence of compact representations. In particular, throughout this section we assume that the size of the new formula P is bounded by a constant (k in the sequel).

We first show that Ginsberg's revision remains incompressible even under the above assumption. This theorem is a trivial consequence of the results of Eiter and Gottlob [EG92, Theorem 8.2]. Note that in the following theorem the size of the data structure representing $T *_G P$ depends only on $|T|$.

Theorem 10 *Suppose there exists a polynomial p such that given any revised knowledge base $T' *_G P'$, there exists a data structure $D_{T',P'}$ and coNP-relation $ASK(\cdot, \cdot)$ such that $|D_{T',P'}| < p(|T'|)$, and for any clause Q, $T' *_G P' \models Q$ iff $ASK(D, Q)$ is true. Then NP$\subseteq$ coNP/poly.*

Proof. Let s be a new propositional variable. Define T', P' from T, P in (3) and (4) by $T' = T \cup \{\neg s, \neg s \vee P\}$, $P' = s$. It immediately follows that $T' *_G P' \models Q$ iff $T *_G P \models Q$. Now suppose there is a polynomial p satisfying the hypotheses of the theorem, in particular, there is a data structure $D_{T',P'}$ such that $|D_{T',P'}| < p(|T'|)$. Since $|T'| = O(|T| + |P|)$, there exists a (slightly greater) polynomial satisfying the hypotheses of Theorem 1. Hence, NP $\subseteq$ coNP/poly. □

The situation for model-based revision operators is more complex. As it turns out all of them admit a compact representation when the size of P is bounded.

Since all the representations are very similar, we only show the representation for Winslett's operator.

Without loss of generality, assume that the alphabet of P is included in the alphabet of T. Let X^1 be the set of all letters appearing in P, and $X = X^1 \cup X^2$ be the alphabet T. Because of the assumption of $|P|$ being bounded by k, it follows that $\mathrm{card}(X^1) \leq k$. Without loss of generality, we assume $\mathrm{card}(X^1) = k$, and denote letters in X^1 as $\{x_1^1, \ldots, x_k^1\}$. Let Y be a set of letters one-to-one with X^1. Since $T \wedge P$ may be inconsistent, we replace in T the letters X^1 of P with the new letters Y. This yields the formula

$$T[X^1/Y] \wedge P \tag{8}$$

This formula is satisfiable – if both T and P are – but it is not query equivalent to $T *_{Win} P$, since any model of P can be suitably extended to a model of (8). Hence we want to impose further constraints to this formula.

Let M be a model of (8), i.e. an assignment to $X^1 \cup X^2 \cup Y$. We partition M as $M_{|X^1} \cup M_{|X^2} \cup M_{|Y}$, with $M_{|X^1} = M \cap X^1$, $M_{|X^2} = M \cap X^2$ and $M_{|Y} = M \cap Y$. Observe M is a model of (8) if and only if $M_{|X^1} \cup M_{|X^2}$ is a model of P, and $M'_{|X^1} \cup M_{|X^2}$ is a model of T, where $M'_{|X^1} = \{x_i^1 \in X_1 \mid y_i \in M_{|Y}\}$.

Recall that in Winslett's revision, proximity between models is defined using set difference. Since P does not impose constraints to letters in X^2, the models of P which are closest to a model M of T – i.e. having minimal set difference with M – agree with M on letters in X^2.

Therefore, a model $M = M_{X^1} \cup M_{X^2} \cup M_{|Y}$ of (8) should be discarded if there exists another model of P which is closer to the model $M'_{X^1} \cup M_{X^2}$ of T. Let Z be a set of letters one-to-one with X^1. We impose that an assignment to Z is a model of P with the formula

$$V_1 \quad = \quad P[X^1/Z]$$

The assignment $M_{|Z}$ on Z defines an assignment $M''_{|X^1}$ on X^1 as $M''_{|X^1} = \{x_i^1 \in X_1 \mid z_i \in M_{|Z}\}$. This assignment is closer than the assignment M'_{X^1} if the following two formulae are both satisfied:

$$V_2 \quad = \quad \bigwedge_{j=1}^{k} ((z_j = x_j^1) \rightarrow (y_j = x_j^1))$$

$$V_3 \quad = \quad \bigvee_{j=1}^{k} ((z_j = x_j^1) \wedge (y_j \neq x_j^1))$$

Since a model of (8) must be considered only if there does not exist a closer model, the whole revision can be reformulated as the following quantified boolean formula:

$$T[X^1/Y] \wedge P \wedge \neg \exists Z.(V_1 \wedge V_2 \wedge V_3)$$

where $\exists Z$ is a shorthand for $\exists z_1 \cdots \exists z_k$. This formula could be turned into an (unquantified) propositional formula by replacing the existential quantification with a disjunction over all assignments to Z, which are 2^k (a constant, since $k = |P|$ is a constant). Hence the total size of $T *_{Win} P$ is $O(|T| + |P| + 2^{|P|})$. Observe that the assignments to Z which are not models of P can be discarded (they do not satisfy V_1), and the simplified formula is linear in the number of models of P, hence it could be significantly smaller than $2^{|P|}$. Note that this explicit representation introduces new letters, hence it does not preserve logical equivalence (2). A linear-size representation with new letters was also shown by Winslett herself in [Win90].

It is also possible to find a compact representation which uses exactly the same alphabet of T and P. Again, we exploit the fact that we can explicitly represent all assignments to X^1 in constant space, since $\mathrm{card}(X^1) = k$.

Let $S \subseteq X^1$. Define the substitution $R_S = \{x_i/\neg x_i \mid x_i \in S\}$, replacing each letter in S with its negation. The formula

$$P \wedge \bigvee_{S \subseteq X^1} \left(T[R_S] \wedge V\right)$$

where $\bigvee_{S \subseteq X^1}$ means a disjunction for all possible subsets S of X^1, is true iff there is a model in T (whose distance from the model of P is given by the set S), satisfying condition V.

Now S represents the distance between these models. We have to specify that there is no other model in P whose distance is less than S.

$$V = \neg \bigvee_{C \neq \emptyset\ ,\ C \subseteq S} P[R_C]$$

where $R_C = \{x_i/\neg x_i \mid x_i \in C\}$, and $\bigvee_{C \neq \emptyset\ ,\ C \subseteq S}$ means a disjunction for all possible subsets C of S, except for $\emptyset$. This formula imposes a condition over the distance C between two models of P. Namely, it forbids that C is between the model of T and the model of P.

The whole formula is

$$P \wedge \bigvee_{S \subseteq X^1} (T[R_S] \wedge \neg \bigvee_{C \neq \emptyset\ ,\ C \subseteq S} P[R_C])$$

Notice that this formula uses only letters of X, and has size linear in $|T|$, but exponential in $|P|$ (i.e. $O(2^{2|P|})$). We found similar representations for the other operators. As a result of the above analysis we obtain the following property:

Theorem 11 *Let T and P be propositional formulae. There exist formulae T' of size polynomial in $|T|$ satisfying criterion (2) where the revision operator is one of $\{*_B, *_D, *_F, *_S, *_{Web}, *_{Win}\}$.*

5 Conclusions

In this paper we presented several results about the size of a propositional theory T' representing the revision of a knowledge base and satisfying either criterion (1) or (2). We proved that some formalizations of belief revision (e.g., Forbus' and Ginsberg's) lead to propositional theories T' which are intrinsically not representable in polynomial space (unless there is a collapse in the polynomial hierarchy). Dalal's and Weber's formalizations have an interesting behavior: T' has no polynomial-size representation if we insist on logical equivalence (2), but such a representation does exist if we ask only query equivalence (1).

Restricting our attention to the bounded-size case also lead to interesting results. Some formalizations retain their non-compactability, while others admit a compact representation. In this case no differences arise when we choose a different equivalence criterion. Results are summarized in Table 1, where YES stands for compactable, while NO stands for not compactable.

While the analysis presented in the paper covers a number of systems and situations, many other interesting cases are yet to be considered. We are currently addressing other restricted cases. In particular, we will focus on situations where both T and P admit polynomial-time algorithms for inference (e.g., they are Horn). We want to find out when the representation T' of $T * P$ is such that $T' \models Q$ can be decided in polynomial time.

Another issue that we want to address is the increase of the size of a revised knowledge base when this is composed with a series of revisions, not only one. Take for example Dalal's revision and query equivalence. If we are to apply the representation proposed in the paper, after a series of n updates the size of the result is exponentially larger than the original theories. Nevertheless, it might well be the case that more compact representations can be found when a series of revisions is applied.

References

[Bor85] A. Borgida. Language features for flexible handling of exceptions in information systems. *ACM Transactions on Database Systems*, 10:563–603, 1985.

[Dal88] M. Dalal. Investigations into a theory of knowledge base revision: Preliminary report. In *Proc. of the 7th Nat. Conf. on Artificial Intelligence (AAAI-88)*, pages 475–479, 1988.

[EG92] T. Eiter and G. Gottlob. On the complexity of propositional knowledge base revision, updates and conterfactuals. *Artificial Intelligence Journal*, 57:227–270, 1992.

[FHV83] R. Fagin, J. Y. Halpern, and M. Y. Vardi. On the semantics of updates in databases. In *Proc. of the 2nd ACM SIGACT SIGMOD Sym. on Principles of Database Systems (PODS-83)*, pages 352–365, 1983.

[For89] K. D. Forbus. Introducing actions into qualitative simulation. In *Proc. of the 11th Int. Joint Conf. on Artificial Intelligence (IJCAI-89)*, pages 1273–1278, 1989.

[Gin86] M. L. Ginsberg. Conterfactuals. *Artificial Intelligence Journal*, 30:35–79, 1986.

[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.

[KL80] R. M. Karp and R. J. Lipton. Some connections between non-uniform and uniform complexity classes. In *Proc. of the 12th ACM Sym. on Theory of Computing (STOC-80)*, pages 302–309, 1980.

[KM91] H. Katsuno and A. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proc. of the 2nd Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR-91)*, pages 387–394, 1991.

[Neb91] B. Nebel. Belief revision and default reasoning: Syntax-based approaches. In *Proc. of the 2nd Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR-91)*, pages 417–428, 1991.

[Neb94] B. Nebel. Base revision operations and schemes: Semantics, representation and complexity. In *Proc. of the 11th European Conf. on Artificial Intelligence (ECAI-94)*, pages 341–345, 1994.

[Sat88] K. Satoh. Nonmonotonic reasoning by minimal belief revision. In *Proc. of the Int. Conf. on Fifth Generation Computer Systems (FGCS-88)*, pages 455–462, 1988.

[Web86] A. Weber. Updating propositional formulas. In *Proc. of First Conf. on Expert Database Systems*, pages 487–500, 1986.

[Win90] M. Winslett. *Updating Logical Databases*. Cambridge University Press, 1990.

[Yap83] C. K. Yap. Some consequences of non-uniform conditions on uniform classes. *Theoretical Computer Science*, 26:287–300, 1983.

Formalism	General case: Logical equiv. (2)	General case: Query equiv. (1)	Bounded case: Logical equiv. (2)	Bounded case: Query equiv. (1)
Ginsberg, Nebel	NO th. 9	NO cor. 2	NO th. 10	NO th. 10
Borgida, Satoh Winslett	NO th. 9	NO th. 3	YES th. 11	YES th. 11, [Win90]
Forbus	NO th. 9	NO th. 4	YES th. 11	YES th. 11
Dalal	NO th. 8	YES cor. 6	YES th. 11	YES cor. 6, th. 11
Weber	NO th. 8	YES th. 7	YES th. 11	YES th. 7, th. 11, [Win90]
WIDTIO	YES –	YES –	YES –	YES –

Table 1: Is the revised knowledge base compressible?