The Compactness of Belief Revision and Update Operators

Paolo Liberatore and Marco Schaerf
Dipartimento di Informatica e Sistemistica
Università di Roma "La Sapienza"
Via Salaria 113 - Roma, Italy
Email: {liberato,schaerf}@dis.uniroma1.it
WWW: http://www.dis.uniroma1.it/~{liberato,schaerf}

September 8, 2000

Abstract

One of the main challenges in the formal modeling of common-sense reasoning is the ability to cope with the dynamic nature of the world. Among the approaches put forward to address this problem are *belief revision* and *update*. Given a knowledge base T, representing our knowledge of the "state of affairs" of the world of interest, we are often led to trust another piece of information P, possibly conflicting with the old one T. The aim of revision and update operators is to characterize the revised knowledge base T' that incorporates the new formula P into the old one T while preserving consistency and, at the same time, avoiding the loss of too much information. In this paper we study the ability of belief revision semantics in compactly representing the set of models of the updated knowledge base.

1 Introduction

During the last years, many formalisms have been proposed in the AI literature to model common-sense reasoning. Particular emphasis has been put in the formal modeling of the dynamic nature of common-sense reasoning. The AI goal of providing a logic model of human agents' capability of reasoning in the presence of incomplete and changing information has proven to be a very hard one. Nevertheless, many important formalisms have been put forward.

Given a knowledge base T, representing our knowledge of the "state of affairs" of the world of interest, it is possible that we are lead to trust another piece of information P, possibly inconsistent with the old one T. The aim of revision operators is to incorporate the new formula P into the old one while preserving consistency and, at the same time, avoiding the loss of too much information. This process has been called *belief revision* and the result of revising T with P is denoted as $T * P$.

This "minimal change assumption" was followed by the introduction of a large number of specific revision operators. Among the others, we mention Fagin, Ullman and Vardi [FUV83], Ginsberg [Gin86] and Dalal [Dal88]. A general framework for belief revision has been proposed by Alchourron, Gärdenfors and Makinson [AGM85, Gär88]. A close variant of revision is *update*. The general framework for update has been studied by Katsuno and Mendelzon [KM89, KM91] and specific operators have been proposed, among the others, by Winslett [Win90] and Forbus [For89].

While most of the early work aimed at defining the appropriate semantics for revision and update, more recently some researchers investigated the computational complexity of reasoning with the operators introduced in the literature. The most complete complexity analysis has been done by Eiter and Gottlob in [EG92].

The computational complexity of model checking has been studied by Liberatore and Schaerf [LS96]. This problem is strictly related to another computational aspect of knowledge representation formalisms: their *representational compactness*. We present here a comparison of the compactness of the various belief revision formalisms in representing knowledge.

2 Preliminaries

In this section we (very briefly) present the background and terminology needed to understand the results presented later in the paper. Throughout this paper we restrict our attention to a (finite) propositional language.

The *alphabet* of a propositional formula is the set of all propositional atoms occurring in it. Formulae are built over a finite alphabet of propositional letters using the usual connectives $\neg$ (not), $\vee$ (or) and $\wedge$ (and). Additional connectives are used as shorthands, $\alpha \rightarrow \beta$ denotes $\neg\alpha \vee \beta$, $\alpha = \beta$ is a shorthand for $(\alpha \wedge \beta) \vee (\neg\alpha \wedge \neg\beta)$ and $\alpha \neq \beta$ denotes $\neg(\alpha = \beta)$.

An *interpretation* of a formula is a truth assignment to the atoms of its alphabet. A *model* M of a formula F is an interpretation that satisfies F (written $M \models F$). Interpretations and models of propositional formulae will be denoted as sets of atoms (those which are mapped into 1). A theory T is a set of formulae. An interpretation is a model of a theory if it is a model of every formula of the theory. Given a theory T and a formula F we say that T entails F, written $T \models F$, if F is true in every model of T. Given a propositional formula or a theory T, we denote with $\mathcal{M}(T)$ the set of its models. We say that a knowledge base T *supports* a model M if $M \in \mathcal{M}(T)$, or equivalently $M \models T$. A knowledge base T is consistent, written $T \not\models \bot$, if $\mathcal{M}(T)$ is non-empty. Moreover, we denote with $\mathcal{F}$ the inverse operator of $\mathcal{M}$, that is, given a set of models A, $\mathcal{F}(A)$ denotes one of the formulae that have exactly A as set of models.

In order to make the formulae more compact and readable, we overload the boolean connectives to apply to sets of letters. For example, given three disjoint sets of letters W, S and R with the same number of elements k, we use the notation $(\neg S)$ as a shorthand for the formula $\bigwedge\{\neg s_i | s_i \in S\}$, $(S \equiv R)$ to denote $\bigwedge\{s_i \equiv r_i | 1 \leq i \leq k\}$, $(S \equiv \neg R)$ to denote $\bigwedge\{s_i \equiv \neg r_i | 1 \leq i \leq k\}$ and $(W \equiv (S \equiv \neg R))$ for $\bigwedge\{w_i \equiv (s_i \equiv \neg r_i) | 1 \leq i \leq k\}$. For example, the formula

$$T \wedge (W \equiv (X \equiv \neg Y)$$

where $T = (x_1 \wedge (\neg x_3 \vee x_2))$ is a shorthand for:

$$x_1 \wedge (\neg x_3 \vee x_2) \wedge [w_1 \equiv (x_1 \equiv \neg y_1)] \wedge [w_2 \equiv (x_2 \equiv \neg y_2)] \wedge [w_3 \equiv (x_3 \equiv \neg y_3)]$$

2.1 Belief Revision and Update

Belief revision is concerned with the modeling of accommodating a new piece of information (the revising formula) into an existing body of knowledge (the knowledge base), and they might contradict each other. A slightly different perspective is taken by knowledge update. For an interesting discussion on the differences between belief revision and update we refer the reader to the work of Katsuno and Mendelzon [KM91]. We assume that both the revising formula and the knowledge base can be either a single formula or a theory.

We now recall the different approaches to revision and update, classifying them into formula-based and model-based ones. 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 (w.r.t. set inclusion) sets in S.

Formula-based approaches

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

$$W(T, P) \;\; = \;\; \{T' \subseteq T \;\;|\;\; T' \cup \{P\} \not\models \bot, \neg\exists U : T' \subset U \subseteq T, U \cup \{P\} \not\models \bot\}$$

The set $W(T, P)$ contains all the plausible subsets of T that we may retain when inserting P.

Ginsberg. Fagin, Ullman and Vardi in [FUV83] and, independently, Ginsberg in [Gin86] define the revised knowledge base as a set of theories: $T *_G P \doteq \{T' \cup \{P\} \mid T' \in W(T,P)\}$. That is, the result of revising T is the set of all maximal subsets of T consistent with P, plus P. Logical consequence in the revised knowledge base is defined as logical consequence in each of the theories, i.e., $T *_G P \models Q$ iff for all $T' \in W(T,P)$, $T' \cup \{P\} \models Q$. In other words, Ginsberg considers all sets in $W(T,P)$ equally plausible and inference is defined skeptically, i.e., Q must be a consequence of each set.

A more general framework has been defined by Nebel in [Neb91]. Here we do not analyze his definitions.

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 $T *_{Wid} P \doteq (\cap W(T,P)) \cup \{P\}$ [Win89].

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 WIDTIO revision gives the same results.

Model-based approaches

These operators select the models of P on the basis of some notion of proximity to the models of T. Model-based approaches assume T to be a single formula, if T is a set of formulae it is implicitly interpreted as the conjunction of all the elements. Many notions of proximity have been defined in the literature. We distinguish them between pointwise proximity and global proximity.

We first recall approaches in which proximity between models of P and models of T is computed pointwise w.r.t. each model of T. That is, they select models of T one-by-one and for each one choose the closest model of P. These approaches are considered to be more suitable for knowledge update [KM91]. Let M be a model, we define $\mu(M,P)$ as the set containing the minimal differences (w.r.t. 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 knowledge base as $\mathcal{M}(T *_W P) \doteq \{N \in \mathcal{M}(P) \mid \exists M \in \mathcal{M}(T) : M \Delta N \in \mu(M,P)\}$. In other words, for each model of T it chooses the closest (w.r.t. set-containment) model of P.

Borgida. Borgida's operator $*_B$ [Bor85] coincides with Winslett's one, except in the case when P is consistent with T, in which case Borgida's revised theory is simply $T \wedge P$.

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 $\mathcal{M}(T *_F P) \doteq \{N \in \mathcal{M}(P) \mid \exists M \in \mathcal{M}(T) : card(M \Delta N) = k_{M,P}\}$. 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. In other words, these approaches consider at the same time all pairs of models $M \in \mathcal{M}(T)$ and $N \in \mathcal{M}(P)$ and find all closest pairs. Let $\delta(T, P)$ be defined as follows:

$$\delta(T, P) \doteq min_{\subseteq} \bigcup_{M \in \mathcal{M}(T)} \mu(M, P)$$

Satoh. In [Sat88], the models of the revised knowledge base are defined as $\mathcal{M}(T *_S P) \doteq \{N \in \mathcal{M}(P) \mid \exists M \in \mathcal{M}(T) : N \Delta M \in \delta(T, P)\}$. That is, Satoh selects all closest pairs (by set-containment of the difference set) and then projects on the models of 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 $\mathcal{M}(T *_D P) \doteq \{N \in \mathcal{M}(P) \mid \exists M \in \mathcal{M}(K) : card(N \Delta M) = k_{T,P}\}$. That is, Dalal selects all closest pairs (by cardinality of the difference set) and then projects on the models of P.

Wrong Variables Revisions

These two revision operators are based upon the hypothesis that the interpretation of a subset of the variables, denoted with Ω, was wrong in the old knowledge base T. The difference between them is based on a different definition of Ω. In Hegner's revision, Ω is the set of the variables of P. The underlying idea is that the original knowledge base T was completely inaccurate w.r.t. everything mentioned in P.

Hegner. Let Ω be the variables of P. The models of Hegner's revised theory are defined as $\mathcal{M}(T *_H P) \doteq \{N \in \mathcal{M}(P) \mid \exists M \in \mathcal{M}(T) \ : \ N \Delta M \subseteq \Omega\}$. For further details see [Win89].

Weber's revision is slightly less drastic. It assumes that the letters whose interpretation was wrong are a subset of the letters of P, i.e., only those occurring in a minimal difference between models of T and P.

Weber. Same definition as Hegner's update, but $\Omega \doteq \cup \delta(T, P)$.

3 Compactness

In this section we summarize some definitions and results proposed in [CDLS96b, CDLS96a] adapting them to the context and terminology of belief revision formalisms. In those papers we introduced a new complexity measure for decision problems, called *compilability*.

We consider a finite alphabet of propositional symbols $L = \{a, b, c, \ldots\}$, possibly with subscripts. Since we restrict to propositional formalisms, we admit neither variable symbols, nor quantifiers. A *model-theoretic semantics* for a formalism is a description of how to extend an interpretation for L to well-formed knowledge bases and queries.

A *model* of a revised knowledge base $T * P$ in a formalism F is an interpretation M that maps $T * P$ to true (written $M \models T *_F P$). Sometimes models will be denoted as sets of letters which are mapped into true.

Assumption 1. Throughout the paper, we assume that *a revised knowledge base in a belief revision formalism is used to represent a set of models.*

In this way, all belief revision formalisms can be compared on the basis of which sets of models they can represent, and how compact is the representation. Notice that a revised knowledge base is represented by a pair of formulae (T and P).

3.1 Reductions

We now define the kind of reduction between belief revision formalisms that we analyze in the following sections. A formula (as well as pairs of formulae) can always be represented as a string over an alphabet Σ, hence from now on we consider translations as functions transforming strings. The *length* of a string $x \in \Sigma^*$ is denoted by $|x|$.

A function f is called *poly-size* if there exists a polynomial p such that for all x it holds $|f(x)| \leq p(|x|)$. Moreover, a function g is called *poly-time* if there exists a polynomial q such that for all x, $g(x)$ can be computed in time less than or equal to $q(|x|)$. These definitions extend to binary functions considering the sum of the sizes of their arguments. In the sequel we assume that f takes as input a pair of formulae and outputs a pair of formulae.

Let F_1 and F_2 be two belief revision formalisms. There exists a *poly-size reduction* from F_1 to F_2, denoted as $f : F_1 \mapsto F_2$, if f is a poly-size function such that for any given pair of formulae $\langle T, P \rangle$ in F_1, $f(\langle T, P \rangle)$ is a pair of formulae in F_2.

Clearly some restrictions must be imposed on such functions. In particular, we define a form of reduction that preserves the models of the original revised theory:

Definition 1 (Model Preservation) *A poly-size reduction* $f : F_1 \mapsto F_2$ *satisfies* model-preservation *if for each pair* $\langle T, P \rangle$ *in* F_1 *there exists a poly-time function* $g_{\langle T,P \rangle}$ *such that for every interpretation* M *of the variables of* $\langle T, P \rangle$ *it holds that* $M \models T *_{F_1} P$ *iff* $g_{\langle T,P \rangle}(M) \models (T' *_{F_2} P')$ *where* $\langle T', P' \rangle = f(\langle T, P \rangle)$.

Notice that the constraints we impose on the reductions, in order to satisfy model preservation, are very weak. We only require that, for any pair $\langle T, P \rangle$, there exists a pair in the other formalism and a polynomial time function such that models are preserved. Even though our constraints are weak we show in this paper that reductions satisfying model preservation do not always exist.

Using the above definition, we want to characterize a formalism in terms of its capacity to compactly represent information. In order to do this we briefly recall the definition of compilability hierarchy introduced in [CDLS96b].

Compilability classes have been defined for problems whose input can be divided into two parts: one part is *fixed* or *accessible off-line* ($\langle T, P \rangle$), and the second one is *varying*, or *accessible on-line* (the model M). Compilability aims at capturing the *on-line complexity* of solving a problem composed of such inputs, i.e., complexity w.r.t. the second input when the first one can be preprocessed in an arbitrary way. In the next section we show the close connection between compilability and the space efficiency of belief revision formalisms.

We define a *language of pairs* S as a subset of $\Sigma^* \times \Sigma^*$. This is necessary to represent the two inputs to a belief revision reasoning problem, i.e., the revised knowledge base ($T * P$), and the interpretation (M). The *non-uniform compilability classes* are defined as follows.

Definition 2 (nu-comp-C **Classes)** *A language of pairs* $S \subseteq \Sigma^* \times \Sigma^*$ *belongs to* nu-comp-C *iff there exists a binary poly-size function* f *and a language of pairs* S' *such that for all* $\langle x, y \rangle \in S$ *it holds:*

1. $\langle f(x, |y|), y \rangle \in S'$ *iff* $\langle x, y \rangle \in S$;
2. $S' \in$ C.

Notice that the poly-size function f takes as input both x (that is the pair $\langle T, P \rangle$) and the size of y (the interpretation). If we want to rewrite off-line x into a new string (pair of formulae), our definition requires that we know in advance the *size* of y. Since y is an interpretation of the letters of x, its size is bounded by $|x|$ and therefore it is known in advance.

We now introduce a notion of reduction between problems.

Definition 3 (Non-uniform comp-reducibility) *Given two problems* A *and* B, A *is* non-uniformly comp-reducible *to* B *(denoted as* $A \leq_{nu-comp} B$*) iff there exist two poly-size binary functions* f_1 *and* f_2*, and a polynomial-time binary function* g *such that for every pair* $\langle x, y \rangle$ *it holds that* $\langle x, y \rangle \in A$ *if and only if* $\langle f_1(x, |y|), g(f_2(x, |y|), y) \rangle \in B$.

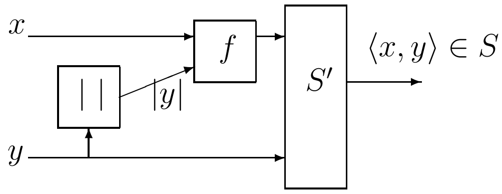

Figure 1: The nu-comp-C Machine

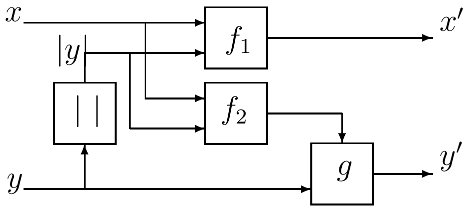

Figure 2: The $\leq_{nu-comp}$ reductions

The $\leq_{nu-comp}$ reductions can be pictorially represented as in Figure 2.
Such reductions satisfy all important properties of a reduction:

Theorem 1 (Thm. 5 of [CDLS96b]) *The reductions $\leq_{nu-comp}$ satisfy transitivity and are compatible [Joh90, pg. 79] with the class* nu-comp-C *for every complexity class* C *in the polynomial hierarchy.*

Therefore, it is possible to define a notions of *hardness* and *completeness*.

Definition 4 (nu-comp-C**-completeness)** *Let S be a language of pairs and* C *a complexity class in the polynomial hierarchy. S is* nu-comp-C-hard *iff for all problems $A \in$* nu-comp-C *we have that $A \leq_{nu-comp} S$. Moreover, S is* nu-comp-C-complete *if S is in* nu-comp-C *and is* nu-comp-C-hard*.*

We close the section by giving a rationale for these complexity classes. nu-comp-NP-hard problems are "not compilable to P": indeed, if there exists a preprocessing of their fixed part that makes them on-line solvable in polynomial time, then the polynomial hierarchy collapses. The same holds for nu-comp-coNP-hard problems. In general, a problem which is nu-comp-C-complete for a class C can be regarded as the "toughest" problem in C, even after arbitrary preprocessing of the fixed part. On the other hand, a problem in nu-comp-C is a problem that, after preprocessing of the fixed part, becomes a problem in C (i.e., it is "compilable to C").

We end this section by proving the nu-comp-Σ_2^p-completeness of a problem that will be used in several proofs. We define $\epsilon\exists\forall$-QBF as follows:

$$\{\langle x, y\rangle \mid x = \epsilon \text{ and } y \text{ is a true quantified boolean formula of the form } \exists\forall F\}$$

Proposition 2 *$\epsilon\exists\forall$-QBF is* nu-comp-Σ_2^p*-complete. It remains* nu-comp-Σ_2^p*-complete when F is of the form $\neg\Pi$, where Π is a set of 3CNF clauses.*

3.2 Compactness of belief revision formalisms

In this section, adapted from [CDLS96a], we show how to use our compilability classes to compare the compactness of belief revision formalisms. The result on compilability allows for comparing the compactness of knowledge representation formalisms [CDLS96a]. In particular, we want to know how compact are the various formalisms in representing sets of models via the revision (update) of an original knowledge base T with a new formula P.

Let F_1 and F_2 be two belief revision formalisms representing sets of models. We prove in this section that any pair of formulae $\langle T, P\rangle$ in F_1 can be reduced, via a poly-size reduction, to a new pair of formulae $\langle T', P'\rangle$ in F_2 satisfying model-preservation if and only if the compilability class of the problem of *model checking* (first input: $\langle T, P\rangle$, second input: interpretation M) in F_1 is higher than or equal to the compilability class of the problem of model checking in F_2.

In order to simplify the presentation we introduce some definitions.

Definition 5 (Model hardness/completeness) *Let F be a belief revision formalism. If the problem of model checking for F belongs to the compilability class* nu-comp-C*, we say that F is in* model-C*. Similarly, if model checking is* nu-comp-C*-complete (hard), we say that F is* model-C*-complete (hard).*

These definitions implicitly define the model hierarchy (model-C). We can now formally establish the connection between compactness of representations and compilability classes. In the following theorems, the complexity classes C, C_1, C_2 belong to the polynomial hierarchy [Sto76]. In Theorem 4 we assume that the polynomial hierarchy does not collapse.

Theorem 3 (Thm. 3 of [CDLS96a]) *Let F_1 and F_2 be two belief revision formalisms. If F_1 is* model-C*-complete and F_2 is* model-C*-hard, then there exists a poly-size reduction $f : F_1 \mapsto F_2$ satisfying model preservation.*

Theorem 4 (Thm. 4 of [CDLS96a]) *Let F_1 and F_2 be two belief revision formalisms. If the polynomial hierarchy does not collapse, F_1 is* model-C_1*-hard, F_2 is in* model-C_2*, and $C_2 \subset C_1$, then there is no poly-size reduction $f : F_1 \mapsto F_2$ satisfying model preservation.*

The above theorems show that the hierarchy of classes model-C exactly characterizes the space efficiency of a formalism in representing sets of models. In fact, two formalisms at the same level in the model hierarchy can be reduced one into the other via a poly-size reduction (Theorem 3), while there is no poly-size reduction from a formalism (F_1) higher up in the hierarchy into one (F_2) in a lower class (Theorem 4). In the latter case we say that F_1 is *more space-efficient* than F_2.

4 Relative Compactness

In this section we prove the results on the ability of the belief revision and update formalisms to compactly represent knowledge. In particular, we show that the various formalisms can be divided into four classes of equivalence w.r.t. their capability of compactly representing sets of models. While for most of the formalisms the level of compactness is directly related to the complexity of model checking, for others the two results are not correlated. In this section we use some special notation:

Notation 1 Given a set X of propositional variables, we denote with Π_X the set of all possible clauses of length 3 over the alphabet X. $\diamond$

We can now prove our theorems.

Theorem 5 *Ginsberg's belief revision formalism is* model-coNP*-complete.*

Proof. We show the thesis by providing a nucomp reduction from the problem MMC (minimal model checking) that is proved to be nu-comp-coNP-complete in [CDLS96b]. Given an instance $\langle KB, I\rangle$ of MMC we construct T, P and M as follows: $T = \{\neg x | x \in X\}$, $P = KB$ and $M = I$. This a nu-comp-reduction, and by [LS97, Theorem 25], we have that I is a minimal model of KB if and only if $M \models T *_G P$. As a consequence, Ginsberg's belief revision formalism is model-coNP-hard. Since model checking is in coNP we have that this formalism is model-coNP-complete. $\square$

Theorem 6 *Winslett and Borgida's update formalisms are* model-Σ_2^p*-complete.*

Proof. We will prove the nu-comp-Σ_2^p-hardness of model checking for Winslett's operator, hence showing model-Σ_2^p-hardness. The proof for Borgida's revision is similar. Let Π be a set of clauses of three literals over the alphabet $X \cup Y$. We use $\Pi_{X\cup Y}$ to denote the set of all the clauses of three literals over the alphabet $X \cup Y$. We prove that $\exists X \forall Y. \neg\Pi$ if and only if $M \models T *_W P$, where

$$\begin{aligned}
T &= (X = Z) \wedge Y \wedge w \wedge r \\
P &= (\neg X \wedge \neg Y \wedge Z \wedge r \wedge \neg w) \vee \Big[(X = Z) \wedge \big(\bigwedge_{\gamma_i \in \Pi_{X\cup Y}} (c_i \rightarrow \gamma_i) = w = r\big)\Big] \\
M &= \{c_i \mid \gamma_i \in \Pi\} \cup Z \cup \{r\}
\end{aligned}$$

where C is a set of new variables, one to one with the clauses of $\Pi_{X \cup Y}$, and Z is a set of new variables, one to one with X. Notice that the above one is a nu-comp reduction. Given the set of clauses Π and a subset X_1 of X, there is exactly one model of T, namely

$$I = \{c_i \mid \gamma_i \in \Pi\} \cup X_1 \cup \{z_i \mid x_i \in X_1\} \cup Y \cup \{w, r\}$$

By definition, $\mathcal{M}(T *_W P) = \cup_{I \in \mathcal{M}(T)} \mathcal{F}(I) *_W P$. We prove that

1. $M \not\models \mathcal{F}(J) *_W P$, if $J \cap (C \cup X) \neq I \cap (C \cup X)$.
2. $M \models \mathcal{F}(I) *_W P$, if and only if for each subset Y_1 of Y the interpretation $X_1 \cup Y_1$ is not a model of Π.

The first proposition is easy to prove. Indeed, consider the model N defined as

$$N = (J \cap (C \cup X)) \cup [M \cap (Y \cup Z \cup \{w, r\})]$$

This is a model of P. Consider the symmetric difference between J and M:

$$J \Delta M = [(J \cap (C \cup X)) \Delta (M \cap (C \cup X))] \cup (J \Delta N)$$

Since $J \cap (C \cup X)$ is different by hypothesis from $M \cap (C \cup X)$, we have that N is strictly closer to J than M. As a result, $M \not\models \mathcal{F}(J) *_W P$.

Let now prove the second point, that is, $M \models \mathcal{F}(I) *_W P$ if and only if it is not possible to extend X_1 to form a model of Π.

Suppose Π unsatisfiable. Consider the set $\mathcal{F}(I) *_W P$. Since Π is false for any Y_1, the formula P has two kinds of models (apart from M): the models with $C = C_1$, $X = X_1$ and both w and r false, and the models with a different evaluation of C or X. None of these models can prevent M from being a model of the revision $\mathcal{F}(I) *_W P$. The models with a different C or X are not less distant from I than M; Indeed, let J be that model: if I has both x_i and z_i, but J has none of these literals, then J is not closer of I than M, since the last has z_i but not x_i. The same argument holds if I has none of x_i and z_i but J has both. For the differences on C the fact that J is not closer to I than M is trivial, since I and M agree on C, while I and J do not. The models with the same value of X but a negative value of w and r are more distant from I than M, since M and I agree on w and r while the other models disagree.

Suppose Π satisfiable. We can prove that M is not a model of $\mathcal{F}(I) *_W P$. Consider that for any X_1 we have in P the model $M' = \{c_i \mid \gamma_i \in \Pi\} \cup X_1 \cup Z_1 \cup Y_1 \cup \{w, r\}$ (with $Z_1 = \{z_i \mid x_i \in X_1\}$) where Y_1 is the value of Y that satisfies Π. This model is closer to I than M, thus M is not in the revised base.

As a result, we have that M can be a model only of $\mathcal{F}(I) *_W P$, where $I \cap C = \{c_i \mid \gamma_i \in \Pi\}$. Moreover, M is a model of this partial revision if and only if for any subset Y_1 of Y, it holds that $(I \cap X) \cup Y_1$ is not a model of Π. As a result, M is a model of $T *_W P$ if and only if there exists a subset X_1 of X such that for any Y_1, the interpretation $X_1 \cup Y_1$ is not a model of Π.

The same proof can be used to prove the compactness of Borgida's operator. $\square$

Theorem 7 *Satoh's update formalism is* model-Σ_2^p*-complete.*

Proof. We reduce Winslett's revision to Satoh's using a $\leq_{nu-comp}$ reduction. Let T and P be two formulas, and M be a model over the alphabet X. We prove that $M \models T *_W P$ if and only if $M' \models T' *_S P'$, where

$$\begin{aligned} T' &= T \wedge (X = Y) \wedge (Y \neq Z) \\ P' &= \neg Y \wedge \neg Z \wedge P \\ M' &= M \end{aligned}$$

where Y and Z are set of new variables one-to-one with X. Note that M' is not the same as M, since it is considered over an extended alphabet.

Let us consider the formal definitions of $M \models T *_W P$ and $M' \models T' *_S P'$.

$$\begin{array}{lll} M \models T *_W P & \text{iff} & \exists I \in \mathcal{M}(T) \ \nexists N \in \mathcal{M}(P) \ . \ I \Delta N \subset I \Delta M \\ M' \models T' *_S P' & \text{iff} & \exists I' \in \mathcal{M}(T') \ \nexists J' \in \mathcal{M}(T'), N' \in \mathcal{M}(P') \ . \ J' \Delta N' \subset I' \Delta M' \end{array}$$

Consider the second definition. Assume $y_i \notin J' \Delta N'$ and $y_i \in I' \Delta M'$. Since $y_i \notin N'$ this implies $y_i \notin J'$, thus $z_i \in J'$. As a result, $z_i \in J' \Delta N'$, which implies $z_i \in I' \Delta M'$, thus $z_i \in I'$ and also $y_i \notin I'$ and then $y_i \notin I' \Delta M'$ contradicting the hypothesis. As a result, $J' \Delta N' \subset I' \Delta M'$ implies $J' \cap (Y \cup Z) = I \cap (Y \cap Z)$.

Since the value of each x_i in I' is equal to the value of y_i, and the same for J', we have $I' = J'$. As a result

$$J' \Delta N' \subset I' \Delta M' \quad \text{iff} \quad I' \Delta N' \subset I' \Delta M'$$

The definition of $M' \models T' *_S P'$ can thus be simplified

$$M' \models T' *_S P' \quad \text{iff} \quad \exists I' \in \mathcal{M}(T') \ \nexists N' \in \mathcal{M}(P') \ . \ J' \Delta N' \subset I' \Delta M'$$

This is very close to the definition of Winslett's update. The only point left to prove is the mapping between I and I' and N and N'. This is clearly true, since the models of P and P' are the same, while the models of T' are extensions of the models of T. Formally we have

$$\begin{array}{lll} I' \in \mathcal{M}(T') & \text{iff} & I \cap X \in \mathcal{M}(T) \text{ and } x_i \in I' \text{ iff } y_i \in I' \text{ iff } z_i \notin I' \\ N' \in \mathcal{M}(P') & \text{iff} & N \cap (Y \cup Z) \text{ and } N \cap X \in \mathcal{M}(P) \end{array}$$

This means that our condition $M' \models T' *_S P'$ holds if and only if:

$$
\begin{array}{ll}
 & \exists I' \in \mathcal{M}(T') \; \nexists N' \in \mathcal{M}(P') \; . \; J' \Delta N' \subset I' \Delta M' \\
\text{iff} & \exists I \in \mathcal{M}(T) \; \nexists N \in \mathcal{M}(P) \; . \\
 & (I \cup \{y_i \mid x_i \in I\} \cup \{z_i \mid x_i \notin I\}) \Delta N \subset (I \cup \{y_i \mid x_i \in I\} \cup \{z_i \mid x_i \notin I\}) \Delta M \\
\text{iff} & \exists I \in \mathcal{M}(T) \; \nexists N' \in \mathcal{M}(P) \; . \; I \Delta N \subset I \Delta M \\
\text{iff} & M \models T *_W P
\end{array}
$$

This proves that model checking under Winslett's semantics can be reduced to model checking under Satoh's. Now, since T' and P' depends only on T and P, this is indeed a $\leq_{nu-comp}$ reduction. Since Winslett's revision is model-Σ_2^p-complete, this is a proof of model-Σ_2^p-completeness of Satoh's revision. □

Theorem 8 *Forbus' update formalism is* model-Σ_2^p-*complete.*

Proof. Membership follows from the complexity result of [LS00, Theorem 6]. We will prove the nu-comp-Σ_2^p-hardness of model checking for Forbus's operator, hence showing model-Σ_2^p-hardness. We use a reduction from $\epsilon\exists\forall QBF$. Let Π be a set of propositional clauses over the alphabet $X \cup Y$. We assume, without loss of generality, that $|X| = |Y| = n$. We prove that $\exists X \forall Y . \neg\Pi$ if and only if $M \models T *_F P$, where

$$
\begin{array}{rcl}
T & = & (X = Z) \wedge \neg Y \wedge \neg W \wedge \neg r \\
P & = & (X \wedge \neg Z \wedge \neg Y \wedge \neg W \wedge r) \vee [(X = Z) \wedge (Y \neq W) \wedge \neg r \wedge \bigwedge_{\gamma_i \in \Pi_{X \cup Y}} (c_i \rightarrow \gamma_i)] \\
M & = & X \cup \{r\} \cup \{c_i \mid \gamma_i \in \Pi\}
\end{array}
$$

Notice that T and P do not depend on Π, but only on its alphabet. Let us assume that there exists a subset X_1 of X such that Π is false for each model over $X \cup Y$ extending X_1. Consider the following model of T:

$$I = X_1 \cup Z_1 \cup \{c_i \mid \gamma_i \in \Pi\}$$

where $Z_1 = \{z_i \mid x_i \in X_1\}$. We prove that $M \models \mathcal{F}(I) *_F P$. The distance between I and M is

$$|I \Delta M| = |X \backslash X_1| + |Z_1| + 1 = n + 1$$

Now, consider a model N of $(X = Z) \wedge (Y \neq W) \wedge \neg r \wedge \bigwedge_{\gamma_i \in \Pi_{X \cup Y}} (c_i \rightarrow \gamma_i)$. Since Π is false for each model extending X_1, a model of the latter formula must have a different truth evaluation for X or a different evaluation for C. Thus, N must be a model of the kind:

$$N_1 = X_2 \cup Z_2 \cup Y_2 \cup W_2 \cup \{c_i \mid \gamma_i \in \Pi\}$$

where $W_2 = \{w_i \mid y_i \notin Y_2\}$, provided that Π is true in the model $X_2 \cup Y_2$. Or a model of the kind:

$$N_2 = X_1 \cup Z_1 \cup Y_1 \cup W_1 \cup C_2$$

where Y_1 is any valuation of the set Y and $W_1 = \{w_i \mid y_i \notin Y_1\}$, provided that $X_1 \cup Y_1 \cup C_2 \models \bigwedge_{\gamma_i \in \Pi_{X \cup Y}}(c_i \rightarrow \gamma_i)$. The difference between I and N_1 is

$$|I \Delta N_1| = |X_1 \Delta X_2| + |Z_1 \Delta Z_2| + |Y_2| + |W_2| = |X_1 \Delta X_2| + |Z_1 \Delta Z_2| + n > n + 1$$

This is because X_1 and X_2 differ for at least a literal, and the same holds for Z_1 and Z_2. The difference between I and N_2 is

$$|I \Delta N_2| = |Y_1| + |W_1| + |C_2 \Delta \{c_i \mid \gamma_i \in \Pi\}| = n + |C_2 \Delta \{c_i \mid \gamma_i \in \Pi\}| \geq n + 1$$

The update P has two kind of models, namely M and the models of $(X = Z) \wedge (Y \neq W) \wedge \neg r \wedge \bigwedge_{\gamma_i \in \Pi_{X \cup Y}}(c_i \rightarrow \gamma_i)$. Since the models of the latter (N_1 and N_2) are not closer to I than M, it follows that M is a model of $\mathcal{F}(I) *_F P$, thus it is a model of $T *_F P$.

Now, let us assume that there is no subset X_1 of X such that all the models extending X_1 over $X \cup Y$ falsify Π. Thus, for each X_1 there exists a Y_1 such that $X_1 \cup Y_1 \cup C_p i \models \bigwedge_{\gamma_i \in \Pi_{X \cup Y}}(c_i \rightarrow \gamma_i)$. Consider the *generic* model I of T:

$$I = X_1 \cup Z_1 \cup C_1$$

where $Z_1 = \{z_i \mid x_i \in X_1\}$. We prove that M is not a model of $\mathcal{F}(I) *_F P$ (for any $I \in \mathcal{M}(T)$), thus proving that M is not a model of $T *_F P$. The distance between I and M is

$$|I \Delta M| = |X \backslash X_1| + |Z_1| + |C_1 \Delta \{c_i \mid \gamma_i \in \Pi\}| + 1 = n + 1 + |C_1 \Delta \{c_i \mid \gamma_i \in \Pi\}|$$

Since for each X_1 there is a model $X_1 \cup Y_1 \cup \{c_i \mid \gamma_i \in \Pi\} \models \bigwedge_{\gamma_i \in \Pi_{X \cup Y}}(c_i \rightarrow \gamma_i)$, the formula P contains the following model

$$N = X_1 \cup Z_1 \cup Y_1 \cup W_1 \cup C_1$$

where $W_1 = \{w_i \mid y_i \notin Y_1\}$. Consider the distance between I and N

$$|I \Delta N| = |Y_1| + |W_1| = n$$

As a result, M is not a model of $\mathcal{F}(I) *_F P$, and this holds for any $I \in \mathcal{M}(T)$. As a result, M is not a model of $T *_F P$. $\square$

Theorem 9 *Dalal and Weber's update formalisms are* model-NP-*complete.*

Proof. Membership: Given T and P, we off-line compute $k_{T,P}$ (for Dalal's operator) and Ω (for Weber's operator). Now, given a model M, this is a model of $T * P$ if and only if there exists another model N such that $N \models T$ and $|N \Delta M| = k_{T,P}$ (for Dalal's operator) or $N \models T$ and $N \Delta M \subseteq \Omega$ for Weber's operator. Notice that both checks can be completed with an NP-machine. As a consequence, the compilability level of model checking in nu-comp-NP, and, therefore, the formalisms are in model-NP.

Hardness: We show a $\leq_{nu-comp}$ reduction from ϵSAT. Given a set of propositional clauses Π on an alphabet X and a set Y of new letters in one-to-one correspondence with letters of X, we define:

$$\begin{array}{rcl} T & = & \bigwedge_{\gamma_i \in \Pi_X} (c_i \rightarrow \gamma_i) \wedge (X \not\equiv Y) \\ P & = & (\neg X \wedge \neg Y) \\ M & = & \{c_i \mid \gamma_i \in \Pi\} \end{array}$$

We show that Π is satisfiable iff $M \models T *_D P$ iff $M \models T *_{Web} P$.

Notice that in every model of T exactly n atoms from $X \cup Y$ are true, while in every model of $T *_D P$ all atoms from $X \cup Y$ are false. Hence, recalling the definition of Dalal's revision, $k_{T,P} \geq n$. Moreover, $k_{T,P} = n$ since X is a model of T and $\emptyset$ is a model of P, and the cardinality of their difference is n. Now $\{c_i \mid \gamma_i \in \Pi\}$ is a model of P, hence it is also a model of $T *_D P$ iff there exists a model of T such that $card(\{c_i \mid \gamma_i \in \Pi\} \Delta M) = k_{T,P} = n$.

Assume Π is satisfiable. Let X_1 be a model of Π, $\overline{Y_1} = \{y_i \mid x_i \notin X_1\}$ and $M = \{c_i \mid \gamma_i \in \Pi\} \cup X_1 \cup \overline{Y_1}$. We show that M is a model of T. In fact, M satisfies $X \not\equiv Y)$ by construction of $\overline{Y_1}$, and also M satisfies $\bigwedge_{\gamma_i \in \Pi_X} (c_i \rightarrow \gamma_i)$, because for each clause of $\bigwedge_{\gamma_i \in \Pi_X} (c_i \rightarrow \gamma_i)$, either $c_i \notin \{c_i \mid \gamma_i \in \Pi\}$ or γ_i is satisfied by X_1. Now observe that $card(\{c_i \mid \gamma_i \in \Pi\} \Delta M) = card(X_1 \cup \overline{Y_1}) = n$. Hence, $\{c_i \mid \gamma_i \in \Pi\}$ is a model of $T *_D P$. Since the models of Dalal's revision are also models of Weber's revision, we also have that $\{c_i \mid \gamma_i \in \Pi\}$ is a model of $T *_{Web} P$.

Now, observe that $\mathcal{M}(P) = 2^C$, i.e., every subset of C is a model of P. Since X and Y are models of T, $\mu(X, P) = X$, $\mu(Y, P) = Y$. This implies that both X and Y are in $\delta(T, P) \doteq \bigcup_{M \in \mathcal{M}(T)} \mu(M, P)$. Recall that $\Omega \doteq \cup \delta(T, P)$. Hence, $X \cup Y \subseteq \Omega$. Moreover, since $V(P) \cap C = \emptyset$, $\Omega = X \cup Y$. From the definition of Weber's revision, $\{c_i \mid \gamma_i \in \Pi\}$ is a model of $T *_{Web} P$ iff there exists a model M of T such that $M - \Omega = \{c_i \mid \gamma_i \in \Pi\} - \Omega$. Since $\Omega = X \cup Y$, the last condition is equivalent to $M \cap C = \{c_i \mid \gamma_i \in \Pi\}$.

Assume Π is unsatisfiable and $\{c_i \mid \gamma_i \in \Pi\}$ is a model of $T *_{Web} P$. Then there exists a model M of T such that $M \cap C = \{c_i \mid \gamma_i \in \Pi\}$. We claim that $M \cap X$ is a model of Π. In fact, as for Dalal's revision, M satisfies $\bigwedge_{\gamma_i \in \Pi_X} (c_i \rightarrow \gamma_i)$. Simplifying $\bigwedge_{\gamma_i \in \Pi_X} (c_i \rightarrow \gamma_i)$ with truth values of $M \cap C = \{c_i \mid \gamma_i \in \Pi\}$, we conclude that M satisfies $\bigwedge_{c_i \in \{c_i \mid \gamma_i \in \Pi\}} \gamma_i$, which is exactly formula Π. Since the formula contains only atoms from X, $M \cap X$ satisfies Π, hence Π is satisfiable and contradiction arises. Thus, $\{c_i \mid \gamma_i \in \Pi\}$ is not a model of $T *_{Web} P$. Since the set of models of $T *_{Web} P$ is a superset of the models of $T *_D P$, it follows that $\{c_i \mid \gamma_i \in \Pi\}$ is not a model of $T *_D P$. $\square$

Theorem 10 *Hegner's update formalism is* model-NP*-complete.*

Proof. Membership in model-NP immediately follows from the complexity result. Hardness is shown by a $\leq_{nu-comp}$ reduction from ϵSAT that mimics the proof in [LS00, Theorem 8].

Given a propositional formula Π on an alphabet X, we define:

$$\begin{aligned} T &= (\bigwedge_{\gamma_i \in \Pi_X} (c_i \rightarrow \gamma_i)) \vee (\neg r \wedge \neg X) \\ P &= X \\ M &= \{c_i \mid \gamma_i \in \Pi\} \cup X \cup \{r\} \end{aligned}$$

We show that Π is satisfiable if and only if $M \models T *_H P$.

Assume Π satisfiable. As a consequence, T admits a model N containing $\{c_i \mid \gamma_i \in \Pi\} \cup r$. Hence, the difference between M and N is a subset of X, so M is a model of the revised knowledge base.

Assume Π unsatisfiable and $M \models T *_H P$. By definition, there exists a model $N \models T$ such that $M \Delta N \subseteq X$. Therefore, $\{c_i \mid \gamma_i \in \Pi\} \subseteq N$ and this implies $N \not\models \bigwedge_{\gamma_i \in \Pi_X}(c_i \rightarrow \gamma_i)$. Thus, $N \models (\neg r \wedge \neg X)$. Since $M \Delta N \subseteq X$, we have that $M \models \neg r$, but this contradicts the hypothesis. $\square$

Theorem 11 *The WIDTIO update formalism is in* model-P*.*

Proof. Given T and P, we can compute off-line $T' = T *_{Wid} P$, since, by definition, the size of T' is less than or equal to the size of T. Online, we only need to check whether $M \models T'$, that can be accomplished in polynomial time. As a consequence, model checking is in nu-comp-P and, thus, WIDTIO is in model-P. $\square$

When we consider the capability of belief revision and update formalisms to compactly represent knowledge, we have been able to completely characterize all of them and the formalisms all fall into four equivalence classes. The most compact formalisms are those introduced by Satoh, Winslett, Forbus and Borgida that are complete for the class model-Σ_2^p. Ginsberg's formalism is complete for model-coNP, the formalisms proposed by Dalal, Weber and Hegner are complete for model-NP, while the least compact one is WIDTIO which belongs to model-P. Notice that, as witnessed by the WIDTIO formalism, the compactness of a belief revision formalism is not always relsted to the complexity of model checking or inference.

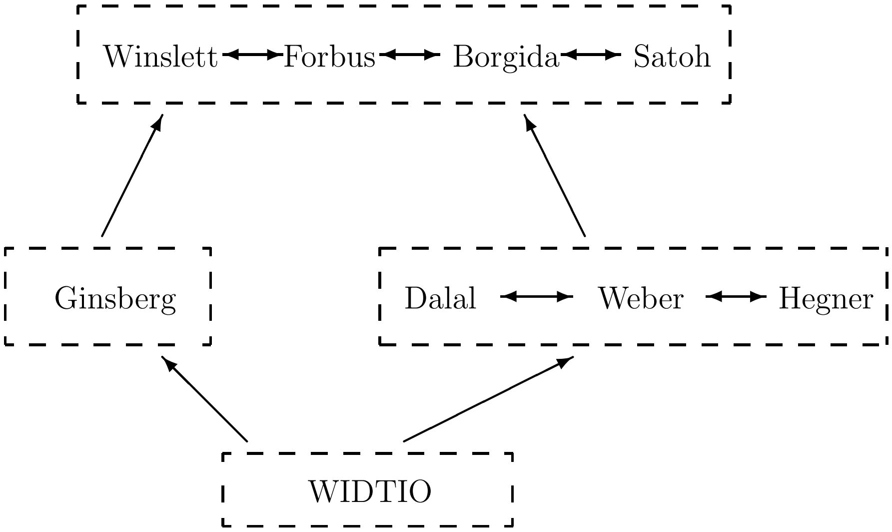

Figure 3: Relative Compactness of Belief Revision and Update Formalisms

5 Conclusions

In this paper we have investigated how compactly the discussed formalisms can represent knowledge.

As it turns out, the various formalisms for belief revision are more compact than classical propositional logic in the representation of sets of models. Furthermore, the complexity of model checking and the compactness are not always related and they are not always related to the complexity of inference.

References

[AGM85] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

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

[CDLS96a] M. Cadoli, F. Donini, P. Liberatore, and M. Schaerf. Comparing space efficiency of propositional knowledge representation formalisms. In *Proceedings of the Fifth International Conference on the Principles of Knowledge Representation and Reasoning (KR'96)*, pages 364–373, 1996.

[CDLS96b] M. Cadoli, F. M. Donini, P. Liberatore, and M. Schaerf. Feasibility and unfeasibility of off-line processing. In *Proceedings of the Fourth Israeli Symposium on Theory of Computing and Systems (ISTCS'96)*, pages 100–109. IEEE Computer Society Press, 1996. URL = `ftp://ftp.dis.uniroma1.it/PUB/AI/papers/cado-etal-96.ps.gz`.

[Dal88] M. Dalal. Investigations into a theory of knowledge base revision: Preliminary report. In *Proceedings of the Seventh National Conference 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 counterfactuals. *Artificial Intelligence*, 57:227–270, 1992.

[For89] K. D. Forbus. Introducing actions into qualitative simulation. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI'89)*, pages 1273–1278, 1989.

[FUV83] R. Fagin, J. D. Ullman, and M. Y. Vardi. On the semantics of updates in databases. In *Proceedings of the Second ACM SIGACT SIGMOD Symposium on Principles of Database Systems (PODS'83)*, pages 352–365, 1983.

[Gär88] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books, MIT Press, Cambridge, MA, 1988.

[Gin86] M. L. Ginsberg. Conterfactuals. *Artificial Intelligence*, 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.

[KM89] H. Katsuno and A. O. Mendelzon. A unified view of propositional knowledge base updates. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI'89)*, pages 1413–1419, 1989.

[KM91] H. Katsuno and A. O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning (KR'91)*, pages 387–394, 1991.

[LS96] P. Liberatore and M. Schaerf. The complexity of model checking for belief revision and update. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI'96)*, pages 556–561, 1996.

[LS97] P. Liberatore and M. Schaerf. Reducing belief revision to circumscription (and viceversa). *Artificial Intelligence*, 93(1–2):261–296, 1997.

[LS00] P. Liberatore and M. Schaerf. Belief revision and update: Complexity of model checking. *Journal of Computer and System Sciences*, 2000. To appear.

[Neb91] B. Nebel. Belief revision and default reasoning: Syntax-based approaches. In *Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning (KR'91)*, pages 417–428, 1991.

[Sat88] K. Satoh. Nonmonotonic reasoning by minimal belief revision. In *Proceedings of the International Conference on Fifth Generation Computer Systems (FGCS'88)*, pages 455–462, 1988.

[Sto76] L. J. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3:1–22, 1976.

[Win89] M. Winslett. Sometimes updates are circumscription. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI'89)*, pages 859–863, 1989.

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