On the Complexity of Propositional Knowledge Base Revision, Updates, and Counterfactuals

Preliminary Report*

Thomas Eiter and Georg Gottlob

{eiter,gottlob}@vexpert.dbai.tuwien.ac.at

Christian Doppler-Laboratory for Expert Systems
Information Systems Department, Vienna University of Technology
Paniglgasse 16, A-1040 Wien, Austria

Abstract

We study the complexity of several recently proposed methods for updating or revising propositional knowledge bases. In particular, we derive complexity results for the following problem: given a knowledge base T, an update p, and a formula q, decide whether q is derivable from $T \circ p$, the updated (or revised) knowledge base. This problem amounts to evaluating the counterfactual $p > q$ over T. Besides the general case, also subcases are considered, in particular where T is a conjunction of Horn clauses, or where the size of p is bounded by a constant.

1 Introduction

Modeling databases as logical theories is a well-known approach to account for incomplete information and integrity constraints. Similarly, in many areas of AI logical theories capture what is known about the world. Since things may change, the problem of updating logical theories (or knowledge bases) is thus a basic issue considered in the database as well as in the AI community [17, 52].

A number of approaches for changing knowledge bases have been proposed recently. This paper deals with the following methods: Ginsberg's "Possible Worlds" approach [21], introduced by Fagin, Ullman, and Vardi [15, 14] and further studied by Nebel [41, 42], the Cross-Product approach [15], the WIDTIO approach [21, 47, 52], Dalal's revision [7, 8], the approach by Forbus [16], Borgida's method [4], Winslett's "Possible Models" approach [51], Satoh's method [49], and Weber's approach [50].

All these methods adhere to the principle of *Minimality of Change*, which states that the knowledge base should change as small as possible if new information is incorporated, and each of them is suited for particular applications; the current belief is that here is no general method that will "do the right thing" under all circumstances [52].

The computational complexity of the different approaches was posed as an important issue in [29, 52, 42, 24]. Although formula implication is known to be intractable for each method, the precise complexity of this problem remained open. Furthermore, it was not clear under which restrictions tractability can be gained.

In this paper we shed light on these issues. The "implication problem" is as follows: given a knowledge base T, an update p, and a formula q, decide whether q is derivable from the updated (or revised) knowledge base $T \circ p$. This problem exactly corresponds to evaluating a counterfactual, which is a conditional statement of the form "if p were true, then q would hold", where p is known or assumed to be false. By the *Ramsey Test* [17], such a counterfactual is true in the knowledge base T exactly if $T \circ p$ logically implies q.

We consider the implication problem in the propositional case. Besides the general case, the Horn case (all formulas are Horn conjunctions) and the case where the size of p is bounded by a constant is considered.

For each method and each combination of restrictions,

*This paper is an overview of [11], which contains detailed proofs of all complexity results.

we locate the implication problem at the appropriate level of the polynomial hierarchy. It turns out that without restrictions, it resides for all approaches (except for Dalal's) at the second level, and it is thus much harder than classical implication in propositional logic. In the Horn case, it resides at the first level in each case and is thus still intractable. The same applies to the case where the size of p is bounded by a constant. Tractability is obtained only if we consider both restrictions simultaneously.

Our results extend and complement very recent results. Nebel [42] showed that under the possible worlds approach, the implication problem is $\mathbf{\Pi}_2^P$-complete and hence at the second level of the polynomial hierarchy. His proof method, however, does not carry over to most other approaches that we discuss. Nebel also shows that in Dalal's setting, the problem is both $\mathbf{NP}$-hard and co-$\mathbf{NP}$-hard, but in $\mathbf{P}^{\mathbf{NP}}$. We resolve the issue of the exact complexity in this case showing that the problem is complete for the class $\mathbf{P}^{\mathbf{NP}[O(\log n)]}$.

Winslett [52] gives algorithms to compute $T \circ p$ for several approaches and derives complexity bounds therefor. For the WIDTIO approach, she shows that to decide if for a formula $f \in T$ it holds that $f \notin T \circ p$ is in $\mathbf{\Sigma}_2^P$; it is posed as an open issue if this problem is $\mathbf{\Sigma}_2^P$-complete [52, p.33]. A corollary to our results affirmatively resolves this open question, even if p and f are literals.

The rest of this extended abstract is organized as follows. In Section 2 we introduce notation and relevant concepts of complexity theory. In Section 3 we define the different change operators and relate them to counterfactuals. Section 4 presents an overview and a thorough discussion of our results. Proofs for some results are provided in Section 5. Finally, in Section 6 we review closely related results.

2 Preliminaries

Throughout this paper, the language L of propositional logic with finite variable set V is considered. A knowledge base is a finite set of propositional formulae in L. T always denotes a knowledge base and p, q are a propositional formulae. As usual, we will identify T with the conjunction of its formulae. If $S = \{T_1, \ldots, T_m\}$ where $T_1, \ldots, T_n$ are knowledge bases, S is identified with the formula $T_1 \vee \cdots \vee T_n$.

$V(T)$ denotes the variables of V that occur in T. Any truth value assignment to V is a model, which we identify with its set of true variables. A model M is a model of T iff T is satisfied by M (also written $M \models T$). The set of all models of T is denoted by $\mathbf{M}(T)$. Every $x \in V$ is an *atom*, and T is called *atomic* if every formula in T is an atom. Any atom or negated atom is a *literal*, and T is a *literal base* if it contains only literals. A *clause* is a disjunction $l_1 \vee \cdots \vee l_n$ of literals.

Given a knowledge base T and a formula p, $T \circ p$ denotes the change of T by incorporating p into the knowledge base applying the change operator $\circ$.

We assume that the reader is familiar with the basic concepts of the theory of $\mathbf{NP}$-completeness [18]. The problems we discuss reside in low levels of the *polynomial hierarchy* (PH). Recall that the classes Δ_k^P, $\mathbf{\Sigma}_k^P$, and $\mathbf{\Pi}_k^P$ of PH are defined as follows:

$$\Delta_0^P = \mathbf{\Sigma}_0^P = \mathbf{\Pi}_0^P = \mathbf{P},$$

and for all $k \geq 0$,

$$\Delta_{k+1}^P = \mathbf{P}^{\Sigma_k^P}, \mathbf{\Sigma}_{k+1}^P = \mathbf{NP}^{\Sigma_k^P}, \mathbf{\Pi}_{k+1}^P = \text{co-}\mathbf{\Sigma}_{k+1}^P.$$

A problem is at the k-th level of the PH if it is complete for Δ_{k+1}^P under Turing reductions (that is, in Δ_{k+1}^P and $\mathbf{\Sigma}_k^P$-hard or $\mathbf{\Pi}_k^P$-hard).

More recently, $\mathbf{P}^{\mathbf{NP}}$ has been refined to take into account how many $\mathbf{NP}$ oracle calls are needed to solve a problem [43, 32, 28, 27]. $\mathbf{P}^{\mathbf{NP}[f(n)]}$ are the problems decidable in polynomial time with $\leq f(n)$ calls to an $\mathbf{NP}$ oracle, where n is the input size [27]. $\mathbf{P}^{\mathbf{NP}[k]}$ for integer k and $\mathbf{P}^{\mathbf{NP}[O(\log n)]}$ are the classes of problems polynomially decidable with at most k and $O(\log n)$ $\mathbf{NP}$ oracle calls, respectively. Similarly, $\mathbf{P}^{\Sigma_2^P[O(\log n)]}$ contains all problems decidable in polynomial time with at most $O(\log n)$ calls to an $\mathbf{\Sigma}_2^P$ oracle.

3 Knowledge change and counterfactuals

The methods to incorporate p into T differ noticeably from one another, if they are compared according to the following three orthogonal classification criteria:

Irrelevance of Syntax.

In some approaches, p is simply added to T. If the knowledge base becomes inconsistent, then some formulas are retracted before adding p such that consistency is preserved. In this way, a formula is either retracted as a whole or left in T. These approaches, referred to as *formula-based changes*,

possible worlds	$T \circ_G p = \{T' \cup \{p\} : T' \in W(p,T)\}$
cross product	$T \circ_C p = \{p_1 \vee \cdots \vee p_m : p_i \in T_i, 1 \leq i \leq m\}$, where $\{T_1, \ldots, T_m\} = T \circ_G p$
WIDTIO	$T \circ_{Wid} p = \bigcap_{W \in W(p,T)} W \cup \{p\}$

Table 1: Formula-based revision operators

heavily depend on the syntax of T. In contrast to this, *model-based changes* refer only to the models of T in respect of the principle of *Irrelevance of Syntax* [8]. Model-based change is favored by many researchers, but formula-base change also has its uses [30].

Revision versus Update.

Different change operators are appropriate if p consists of an *update*, which reflects a change in the world, or of a *revision*, that is a correction of the description of the (unchanged) world [31, 51, 52, 29]. Rationality postulates for update semantics different from postulates for revision semantics are developed in [29]. An important difference is that revision operates "global" on the models of T while update works "pointwise". That is, in case of a revision the models of the changed knowledge base are models as close to T as possible, where the measure of closeness depends on all models of T. In case of an update, however, each model of T must be changed separately under minimal change, and closeness depends only on that model. The result of an update is the union of all the separate model changes.

Measure of Closeness.

All change approaches need a notion of "closeness". In the formula-based approaches, closeness is measured between sets of formulas, while in the model-based approaches closeness is determined by model difference (that is, difference between sets of atoms). In both approaches, the *inclusion-based* measure occurs, which defines closeness by set difference ordered by inclusion, while the *cardinality-based* measure, which defines closeness by the cardinality of set difference, has been considered merely for model-based change.

3.1 Formula-based operators

The formula-based approaches we consider are based on "possible worlds" [21], which are obtained by the inclusion-based measure. Assume T is consistent. If $T \wedge p$ is inconsistent, a minimal set S of formulas is retracted from T such that $T - S \cup \{p\}$ is consistent. Since S will not be unique in general, several "possible worlds" $T - S$ for p will exist. The family

$$W(p,T) = \{T' \subseteq T : T' \not\models \neg p,\ \ T' \subset U \subseteq T \Rightarrow U \models \neg p\}$$

contains all that possible worlds [21]. Then, the formula-based operators can be defined as in Table 1.

The "Possible Worlds" (or Set-Of-Theories) approach is named according to [21]; it has been introduced in [15, 14] and was reconsidered in [41, 42]. A sentence q is true after change with p iff it is true in all knowledge bases of $T \circ_G p$, that is, if $T \circ_G p \models q$ holds. The cross product approach [15] and the WIDTIO (When In Doubt Throw It Out [52]) approach [21, 47, 41] have been proposed to combine all knowledge bases in $T \circ_G p$ into a single one. It is not difficult to verify that $T \circ_C p \equiv T \circ_G p$; note that the WIDTIO approach retracts in the worst case all knowledge in T.

Example 3.1

Let $T = \{a, b, a \wedge b \Rightarrow c\}$ *and* $p = \neg c$. *Then,*

$$W(p,T) = \{\{a \wedge b \Rightarrow c, b\}, \{a \wedge b \Rightarrow c, a\}, \{a, b\}\},$$

hence

$$T \circ_G p = \{\{a \wedge b \Rightarrow c, b, \neg c\}, \{a \wedge b \Rightarrow c, b, \neg c\}, \{a, b, \neg c\}\}.$$

Under WIDTIO semantics, we have $T \circ_{Wid} p = \{\neg c\}$.

3.2 Model-based operators

The considered model-based change operators rely heavily on the notion of model difference defined by

operator $\circ$	$\mathbf{M}(T \circ p)$ if $T \wedge p$ is inconsistent	$\mathbf{M}(T \circ p)$ if $T \wedge p$ is consistent
Dalal [7, 8] $\circ_D$	$\{M \in \mathbf{M}(p) : \exists M' \in \mathbf{M}(T) \text{ s.t. } \lvert M \triangle M' \rvert = \lvert\triangle\rvert^{min}(T,p)\}$	$T \wedge p$
Forbus [16] $\circ_F$	$\bigcup_{M \in \mathbf{M}(T)} \{M' \in \mathbf{M}(p) : \lvert M' \triangle M\rvert = \lvert\triangle\rvert_M^{min}(p)\}$	
Borgida [4] $\circ_B$	$\bigcup_{M \in \mathbf{M}(T)} \{M' \in \mathbf{M}(p) : M \triangle M' \in \triangle_M^{min}(p)\}$	$T \wedge p$
Winslett [51] $\circ_{Win}$	$\bigcup_{M \in \mathbf{M}(T)} \{M' \in \mathbf{M}(p) : M \triangle M' \in \triangle_M^{min}(p)\}$	
Satoh [49] $\circ_S$	$\{M \in \mathbf{M}(p) : \exists M' \in \mathbf{M}(T) \text{ s.t. } M \triangle M' \in \triangle^{min}(p,T)\}$	$T \wedge p$
Weber [50] $\circ_{Web}$	$\{M \in \mathbf{M}(p) : \exists M' \in \mathbf{M}(T) \text{ s.t. } M - \Omega_{T,p} = M' - \Omega_{T,p}\}$	$T \wedge p$

Table 2: Model-based revision operators

$M \triangle M$, that is the set of propositional symbols in which two models M, M' differ. Let M be a model and p, q be propositional formulas. Then, define

$$\begin{aligned}
\triangle^{min}(p,q) &= \min_{\subseteq}(\{M \triangle M' : M \in \mathbf{M}(p), \\
&\qquad\qquad M' \in \mathbf{M}(q)\}) \\
\triangle_M^{min}(p) &= \min_{\subseteq}(\{M \triangle M' : M' \in \mathbf{M}(p)\}) \\
|\triangle|^{min}(p,q) &= \min\{|D| : D \in \triangle^{min}(p,q)\} \\
|\triangle|_M^{min}(S) &= \min\{|D| : D \in \triangle_M^{min}(p)\} \\
\Omega_{T,p} &= \{x : \exists D \in \triangle^{min}(p,q) \text{ s.t. } x \in D\}
\end{aligned}$$

($\min_{\subseteq}$ selects the minimal elements under set inclusion).

Table 2 describes the considered model-based change operators, where T as well as p is assumed to be consistent. If T or p is inconsistent, then $\mathbf{M}(T \circ_D p) = \mathbf{M}(p)$, $\mathbf{M}(T \circ_{Web} p) = \mathbf{M}(T)$; all other operators work in that case as in column one. Note that $\circ_F, \circ_B, \circ_{Win}$ operate pointwise and that all other operators work global. $\circ_B$ works pointwise only if $T \wedge p$ is inconsistent, however, and gobal otherwise. Pointwise operation in case where $T \wedge p$ is consistent, as performed by $\circ_{Win}$, can be useful, however.

Example 3.2 *Consider the classic example*

$$T = \{a \equiv \neg b\}, \quad p = a$$

appearing in [51]. Since $T \wedge p$ is consistent, $T \circ_B p \equiv a \wedge \neg b$, but $T \circ_{Win} p \equiv a$. If a means 'a book is on the floor' and b means 'a magazine is on the floor', T states that either the book or the magazine is on the floor. If a robot is ordered to put the book on the floor, it is reasonable to accept only that a holds thereafter; why should $a \wedge \neg b$ be accepted, i.e. the magazine is not on the floor ?

3.3 Counterfactuals

A *counterfactual* is a conditional statement of the form "If p, then q", where p is known or expected to be false [33, 21]. The epistemic semantics for counterfactuals of our concern is based on the *Ramsey Test*, which informally means "Accept a sentence of the form 'If A, then C' in a state of belief K if and only if the minimal change of K needed to accept A also requires accepting C." [17, p. 147]. Extensive studies of counterfactuals of this meaning are [21, 17].

Given T and change operator $\circ$, define that "If p, then q" (denoted by $p > q$) is true iff $T \circ p \models q$. The problem of our study is the following problem COUNTERFACTUAL (CF): Given T and formulae p, q, is $p > q$ true in T ?

Example 3.3

Let r denote "It is raining", c, "It is cold", h, "I stay at home", and g, "Gorbachev is president of the United States", and assume

$$T = \{r, \neg c, \neg g, r \wedge c \Rightarrow h\}.$$

Consider $c >_G h$, which means, "If it would be could, I would stay at home". We get

$$W(T, c) = \{\{r, \neg g, r \wedge c \Rightarrow h\}\},$$

and hence

$$T \circ_G c = \{\{r, \neg g, r \wedge c \Rightarrow h, c\}\}.$$

Clearly $T \circ_G c \models h$ holds, and therefore $c >_G h$ is true over T. On the other hand, $c >_G g$, that is "If it would be cold, Gorbachev would be president of the United States", is false over T as $T \circ_G c \not\models g$ holds. Note that the "classical" implications $c \Rightarrow h$ and $c \Rightarrow g$ are both valid in T.

We remark that counterfactuals are distinct from hypothetical reasoning studied in [2, 3]. Roughly speaking, Bonner's hypothetical queries are based, in our terms, on adding nonderivable formulas to T; the salient point is that these formulas are consistent (or not expected to be inconsistent) with T, however.

4 Overview and discussion of results

The main results of this paper are compactly presented in Table 3. Different methods according to the different change semantics for evaluating a counterfactual $p > q$ over a propositional knowledge base correspond to different rows in this table. Recall that a counterfactual $p > q$ is true over a knowledge base T iff $T \circ p \models q$, where "$\circ$" is the chosen change operator; in other terms, evaluating a counterfactual $p > q$ is equivalent to testing whether q could be inferred after a knowledge base change p.

The table contains four columns. The first two columns account for the case where T is a general propositional knowledge base, while the last two columns reflect the Horn case. In the Horn case we assume that p,q and all formulae in T are conjunctions of Horn clauses. Columns one and three describe those cases in which no bound is imposed on the size of the knowledge base change (or antecedent) p, columns two and four account for those cases where the size of p is bounded by a constant. This restriction has been argued for as a reasonable postulate since it will certainly be met by human-generated updates [52, p.33].

Let us discuss the results. The first thing to notice is that in the general case with no restriction on p, evaluating $p > q$ is at the second level of the PH for all approaches except Dalal's. More precisely, for most change operators the problem is $\mathbf{\Pi}_2^P$-complete and for two (WIDTIO and Weber's) it is very close to $\mathbf{\Pi}_2^P$. Under the currently believed hypothesis that the PH does not collapse[26], this means that counterfactual inference is harder than classic inference. On the other hand, the second level is still a very low level of PH, and, in a certain sense, much closer to the first level (that is, **NP**) than to **PSPACE**.

In more practical terms, the $\mathbf{\Pi}_2^P$-hardness results state that with current methods, if we use any subroutine for solving **NP**-complete problems, we would need (in the worst case) an exponential number of calls to this subroutine for deciding whether the counterfactual $p > q$ holds.

Loosely speaking, the $\mathbf{\Pi}_2^P$-hardness is mostly due to the interaction of two different sources of complexity. The first source is the complexity of inference in propositional logic (co-**NP**-hard) and the second source is the large number of possible choices to make a knowledge base (or a model) T consistent with a proposition p performing a minimal change.

The only exception in the first column of our table is Dalal's approach, which is complete for the class $\mathbf{P}^{\mathbf{NP}[O(\log n)]}$. This means that $p > q$ is decidable with $O(\log n)$ **NP** oracle calls, where n is the size of T plus p plus q. The deep reason for this relative easiness is that in Dalal's approach the two sources of complexity are not intermingled and can be separated from one another.

Interestingly, the syntax-based approaches remain $\mathbf{\Pi}_2^P$-hard even if the size of p is bounded by a constant k. This insensitivity towards the size of the knowledge change p can be intuitively explained as follows: even a small change may be have drastic effects by forcing some originally unessential part of a formula in T to become true. For instance, if T contains two formulas $F_1 = \neg u \vee (p_1 \wedge p_2 \wedge \ldots \wedge p_n)$, $F_2 = \neg u$ and if p is equal to u, then p forces the originally completely irrelevant conjunction $(p_1 \wedge p_2 \wedge \ldots \wedge p_n)$ to become true.

On the other hand, the complexity of the model-based approaches decreases when the size of p is bounded by a constant. The intuitive reason for this is that a small change p can affect only few literals; since in this case there are only few alternatives for changing a given model, the second source of complexity disappears.

In the Horn case, the complexity of all approaches except Forbus's is at the first level of the PH, more precisely, co-**NP**-equivalent. Intuitively, this lowering of complexity is due to the fact that propositional inference for propositional Horn clauses is tractable, and hence one source of complexity is eliminated. The cardinality based approaches of Dalal and Forbus are the only ones whose complexity does not decrease under the Horn restriction. This is because their cardinality-based measure of closeness between models has a detrimental effect on the benign properties of Horn clauses.

Let us finally consider the case where both restrictions,

COUNTERFACTUAL $p > q$ true in T?	GENERAL CASE		HORN FORMULAS	
	arbitrary p	$\|p\| \leq k$	arbitrary p	$\|p\| \leq k$
Ginsberg : $p >_G q$	$\mathbf{\Pi}_2^P$-complete		co-$\mathbf{NP}$-complete	
Cross-Product : $p >_C q$	$\mathbf{\Pi}_2^P$-complete		co-$\mathbf{NP}$-complete	
WIDTIO : $p >_{Wid} q$	$\mathbf{\Pi}_2^P$-hard, in $\mathbf{P}^{\mathbf{\Sigma}_2^P[O(\log n)]}$		co-$\mathbf{NP}$-hard, in $\mathbf{P}^{\mathbf{NP}[O(\log n)]}$	
Dalal : $p >_D q$	$\mathbf{P}^{\mathbf{NP}[O(\log n)]}$-complete	$\mathbf{NP}$- and co-$\mathbf{NP}$-hard, in $\mathbf{P}^{\mathbf{NP}[O(1)]}$	$\mathbf{P}^{\mathbf{NP}[O(\log n)]}$-complete	$\mathbf{P}$
Forbus : $p >_F q$	$\mathbf{\Pi}_2^P$-complete	co-$\mathbf{NP}$-complete	$\mathbf{\Pi}_2^P$-complete	$\mathbf{P}$
Borgida : $p >_B q$	$\mathbf{\Pi}_2^P$-complete	$\mathbf{NP}$- and co-$\mathbf{NP}$-hard, in $\mathbf{P}^{\mathbf{NP}[2]}$	co-$\mathbf{NP}$-complete	$\mathbf{P}$
Winslett : $p >_{Win} q$	$\mathbf{\Pi}_2^P$-complete	co-$\mathbf{NP}$-complete	co-$\mathbf{NP}$-complete	$\mathbf{P}$
Satoh : $p >_S q$	$\mathbf{\Pi}_2^P$-complete	$\mathbf{NP}$- and co-$\mathbf{NP}$-hard, in $\mathbf{P}^{\mathbf{NP}[O(1)]}$	co-$\mathbf{NP}$-complete	$\mathbf{P}$
Weber : $p >_{Web} q$	$\mathbf{\Pi}_2^P$-hard, in $\mathbf{P}^{\mathbf{\Sigma}_2^P[O(\log n)]}$	$\mathbf{NP}$- and co-$\mathbf{NP}$-hard, in $\mathbf{P}^{\mathbf{NP}[O(1)]}$	co-$\mathbf{NP}$-hard, in $\mathbf{P}^{\mathbf{NP}[O(\log n)]}$	$\mathbf{P}$

Table 3: Complexity results for propositional counterfactual evaluation

Horn formulas and bounded p, are combined (see the fourth column of our table). As in the general case, bounding the size of p does not improve the complexity of the formula-based methods. The model-based methods, however, all become polynomial under these restrictions and allow counterfactual evaluation even in linear time in this case.

Note that the theorems in the full paper often state yet more acute results than the ones shown in Table 3. For instance, most hardness and completeness results hold even if T is literal base, i.e. T consists of a set (or conjunction) of atoms or negated atoms.

All results presented in Table 3 are new except the $\mathbf{\Pi}_2^P$-completeness of Ginsberg's approach for the general case. The latter was recently shown by Nebel in [42]. Note that in Nebel's proof it is assumed that the given knowledge base T is inconsistent, which is not reasonable in the context of counterfactuals [21]. Our proof is completely different and does not make this inconsistency assumption. Moreover, Nebel's proof method is unsuited for the model-based approaches, which require a different proof technique.

An important feature we do not deal explicitly with are *protected* formulas in T. The set $I \subseteq T$ of protected formulas of T can be viewed as integrity constraints that must not be violated by any change. It is thus required that $T \circ p$ satisfies I, i.e. $T \circ p \models I$. If p is not subject to restrictions, however, this can always be enforced by "adding" I to the change p, i.e. computing $T \circ (p \wedge I)$. Therefore, protected formulas do not alter our complexity results in this case. Moreover, it is quite easy to see that protected formulas leave the worst case complexity unchanged if $|p|$ is bounded by a constant. For these reasons, we do not treat protected formulas as a special case in this paper.

Nebel has also studied in [42] a restriction for the syntax based methods that we do not consider here: disambiguating the knowledge base change by attaching different priorities to the sentences in T. Clearly, such a policy eliminates the second source of complexity, and therefore the complexity level will be lowered by one in general. Note that this prioritization of sentences cannot be applied in the model-based cases. Here we could use a total order on the propositional variables (similar in some sense to the partial order used in prioritized circumscription [34]) indicating which variable should change its value first if an inconsistency has to be removed. We leave a detailed study of such restrictions for the future.

Another interesting issue are compound counterfactuals and sequences of knowledge change. A counterfactual of the form $p > (q > r)$ corresponds to the question "if p would hold, would q counterfactually imply r ?". Some results of the authors indicate that as long as we use nested sequences of counterfactuals, the problem remains $\mathbf{\Pi}_2^P$-complete for most approaches. However, if we mix counterfactual implication ($>$) with classical connectives (such as negation) then the problem becomes **PSPACE**-complete for most of the operators. The latter has also been observed by Mendelzon and Grahne for Winslett's operator $\circ_{Win}$ [25]. Finally, we remark that this result can be sharpened to the Horn case.

5 Proof sketches

In this section, we give proof sketches for some of the results in Table 3. Proofs for all results can be found in the full paper [11].

The following technical lemma captures relationships between various change semantics that are helpful to derive several complexity results with few proofs.

Lemma 5.1 *Let T be a consistent literal base, and let p, q be propositional formulas.*

(1) *$T \circ_G p \equiv T \circ_S p$ and $p >_G q$ is true in T iff $p >_S q$ is true in T.*

(2) *If $V(p) \subseteq V(T)$, then $T \circ_S p \equiv T \circ_B p \equiv T \circ_{Win} p$, and the validity of $p >_S q$, $p >_B q$, and $p >_{Win} q$ in T is equivalent.*

(3) *If $V(p) \subseteq V(T)$, then $T \circ_D p \equiv T \circ_F p$, and $p >_D q$ is true in T iff $p >_F q$ is true in T.*

(4) *If p is consistent, then $T \circ_{Wid} p \equiv T \circ_{Web} p$, and $p >_{Wid} q$ is true in T iff $p >_{Web} q$ is true in T.*

(5) *Let $q \in T$, and let p be consistent. Then the validity of $p >_G q$, $p >_{Wid} q$, $p >_S q$, and $p >_{Web} q$ in T is equivalent, and $q \in T \circ_{Wid}$ iff $p >_G q$ is true in T.*

Exploiting these relationships, $\mathbf{\Pi}_2^P$-hardness of CF under a number of change semantics can be shown with a single transformation.

Lemma 5.2 *Let T be an atomic knowledge base. Deciding if $p > q$ is true in T is $\mathbf{\Pi}_2^P$-hard for $> \in \{>_G, >_C, >_{Wid}, >_B, >_{Win}, >_S, >_{Web}\}$, even if q is an atom with $q \in T$.*

Proof: (sketch) Recall that the most prominent $\mathbf{\Pi}_2^P$-complete problem is deciding validity of a formula from $\mathrm{QBF}_{2,\forall}$, the set of quantified Boolean formulas of the form $\forall \underline{x} \exists \underline{y} E(\underline{x}, \underline{y})$, where $\underline{x} = \langle x_1, \ldots, x_m \rangle, \underline{y} = \langle y_1, \ldots, y_n \rangle$ are tuples of Boolean variables and $E(\underline{x}, \underline{y})$ is a Boolean formula on $x_1, \ldots, x_m, y_1, \ldots, y_n$.

Given an instance $F = \forall \underline{x} \exists \underline{y} E(\underline{x}, \underline{y})$ of $\mathrm{QBF}_{2,\forall}$, we construct in polynomial time an atomic knowledge base T and formulae p, q such that q is an atom from T and $p >_G q$ is true in T iff F is true; since $V(p) \subseteq V(T)$ holds, the claim will be easy to prove with Lemma 5.1.

Let $\underline{x} = \langle x_1, \ldots, x_m \rangle$ and $\underline{y} = \langle y_1, \ldots, y_n \rangle$. Let $z_1, \ldots, z_m$ and r be additional propositional variables. We define T, p, and q as follows:

$$
\begin{aligned}
T &= \{x_1, \ldots, x_m, z_1, \ldots, z_m, y_1, \ldots y_n, r\} \qquad (1)\\
p &= [(x_1 \equiv \neg z_1) \wedge \ldots \wedge (x_m \equiv \neg z_m) \wedge \\
&\quad (r \Rightarrow E(\underline{x}, \underline{y})) \wedge ((y_1 \vee \ldots \vee y_n) \Rightarrow r)] \\
q &= r
\end{aligned}
$$

Note that T and p are consistent, but $T \wedge p$ is inconsistent. The construction of T, p, q is obviously feasible in time polynomial in the size of F.

Let us say a few words about the spirit behind this transformation. The possible worlds for p correspond to truth assignments to $\underline{x}$ and $\underline{y}$ such that $(r \Rightarrow E(\underline{x}, \underline{y})) \wedge ((y_1 \vee \cdots \vee y_n) \Rightarrow r)$ is satisfied. For every truth assignment to $\underline{x}$ there exists a candidate for a possible world in which r and all y_i are false. This candidate can be eliminated if there is a world consistent with p that has the same truth assignment to the x_i

and determines a truth assignment to the y_i such that $E(\underline{x}, \underline{y})$ is satisfied; in this case, r can be added to this world without introducing inconsistency with p. The z_i serve to assure that this candidate can be eliminated only by such a world and not by a world in which the x_i have a different truth value assignment. Thus exactly if all candidates are eliminated, which means that r appears in all possible worlds, the formula F is true.

It is shown in the full paper that $p >_G q$ is true in T iff the formula F is true.

Thus the claim is true for $> \in \{>_G, >_C\}$ and by (1) in Lemma 5.1 also for $>_S$. Since $V(p) \subseteq V(T)$ and $q \in T$, it is by (2) and (5) in Lemma 5.1 further true for $> \in \{>_B, >_{Win}, >_{Wid}, >_{Web}\}$.

□

Remark: $\mathbf{\Pi}_2^P$-hardness of CF_F does not follow from construction (1) (this would imply by (3) in Lemma 5.1 $\mathbf{\Pi}_2^P$-hardness of CF_D, which is in Δ_2^P); this result requires an extra proof.

Theorem 5.3 *Deciding if a counterfactual $p > q$ is true in a knowledge base T is $\mathbf{\Pi}_2^P$-complete for $> \in \{>_G, >_C, >_{Win}, >_B, >_S\}$, even if T is atomic and q is an atom with $q \in T$.*

Proof: By Lemma 5.2 it remains to show membership in $\mathbf{\Pi}_2^P$. We outline how to disprove nondeterministically $p > q$ in polynomial time with an **NP** oracle, without treatment of special cases (e.g., T, p inconsistent).

- $>_G, >_C$: Guess $T' \subseteq T$. Call for every $f \in T - T'$ the oracle whether $T' \cup \{f, p\}$ is inconsistent. Call the oracle to check that $T' \wedge p \not\models q$.
- $>_{Win}$: Guess $M \in \mathbf{M}(T), D \subseteq V(T \wedge p)$. Test if M modified by D is a model of $p \wedge \neg q$. Call the oracle whether there is no $M' \in \mathbf{M}(p)$ such that $M \triangle M' \subset D$.
- $>_B$: Test by oracle call whether $T \wedge p$ is consistent; if this is true, check $T \wedge p \not\models q$ by oracle call, else proceed as for $>_{Win}$.
- $>_S$: Guess $D \subseteq V(T \wedge p)$. Call the oracle whether there exist $M \in \mathbf{M}(T)$, $M' \in \mathbf{M}(p \wedge \neg q)$ such that $M \triangle M' = D$. Call the oracle whether there exist no $M \in \mathbf{M}(T), M' \in \mathbf{M}(p)$ such that $M \triangle M' \subset D$.

□

CF_D is an interesting problem since it is by our knowledge the first logical inference problem known to be complete for the class $\mathbf{P}^{\mathbf{NP}[O(\log n)]}$.

Theorem 5.4 *Deciding if $p >_D q$ is true in T is $\mathbf{P}^{\mathbf{NP}[O(\log n)]}$-complete, even if T is atomic.*

Proof: (sketch)

We show first membership in $\mathbf{P}^{\mathbf{NP}[O(\log n)]}$ and then completeness for this class.

To prove membership in $\mathbf{P}^{\mathbf{NP}[O(\log n)]}$, we sketch an algorithm that decides $p >_D q$ in polynomial time with $O(\log m)$ calls to an NP oracle, where $m = |V(T \wedge p)|$. If T or p is inconsistent, then $p >_D q$ equals $p \models q$, which is decided by an oracle call. Otherwise, proceed in two steps. First, the distance $d = |\triangle|^{min}(T, p)$ between T and p is determined. d can be computed by binary search with $O(\log m)$ oracle calls (to decide $d \leq k$? for given k is obviously in **NP**). The second step is to call the oracle if there exist $M, M' \subseteq V(T \wedge p \wedge q)$ such that $|M \triangle M'| = d$ and $M \in \mathbf{M}(T)$ and $M' \in \mathbf{M}(p \wedge \neg q)$; such M, M' do not exist iff $p >_D q$ is true in T.

For the proof of $\mathbf{P}^{\mathbf{NP}[O(\log n)]}$-hardness, we refer to the problem UOCSAT: Given a set $C = \{C_1, \ldots, C_m\}$ of clauses, decide if all truth assignments that satisfy a maximum number of clauses in C always satisfy the same set of clauses in C. This problem is proved $\mathbf{P}^{\mathbf{NP}[O(\log n)]}$-complete in [28]. It can be transformed into CF_D in polynomial time as follows. Let $c_1, \ldots, c_m$ be new variables that do not appear in C. Define

$$\begin{aligned} T_1 &= \{c_1, \ldots, c_m\} \\ p_1 &= (C_1 \vee \neg c_1) \wedge \cdots \wedge (C_m \vee \neg c_m) \end{aligned} \tag{2}$$

It is not difficult to show that a model M is in $\mathbf{M}(T_1 \circ_D p_1)$ iff the truth assignment by M to the variables in $V(C)$ satisfies as many clauses in C as simultaneously satisfiable (C_i is satisfied by M iff $M \models c_i$). Now let T_1' and p_1' be the knowledge base resp. formula that results if all of its variables are primed, i.e. $T_1' = \{c_1', \ldots, c_m'\}$, and $p_1' = [(C_1' \vee \neg c_1') \wedge \cdots \wedge (C_m' \vee \neg c_m')]$. ($C_i'$ is the modified clause C_i.) Define

$$\begin{aligned} T &= T_1 \cup T_1', \quad p = p_1 \wedge p_1', \\ q &= [(c_1 \equiv c_1') \wedge \cdots \wedge (c_m \equiv c_m')]. \end{aligned} \tag{3}$$

$T \circ_D p$ "combines" each model of $T_1 \circ_D p_1$ on $V(T_1 \wedge p_1)$ with each model of $T_1' \circ_D p_1'$ on $V(T_1' \wedge p_1')$. Then, $p >_D q$ is true in T iff all truth assignments that satisfy a maximum number of clauses in C always satisfy the same clauses in C. □

Theorem 5.5 *It is $\mathbf{\Pi}_2^P$-complete for $> \in \{>_G, >_C\}$ and $\mathbf{\Pi}_2^P$-hard for $>_{Wid}$ to decide if $p > q$ is true in a complete knowledge base T, even if p, q are atoms and $q \in T$.*

Proof: Let s be a new propositional variable. Define T', p', q' from T, p, q in (1) by $T' = T \cup \{\neg s, \neg s \vee p\}$, $p' = s$, $q = r$. It is immediately verified that for each variable $x \in V = T \cup \{s\}$, $T \models x$ or $T \models \neg x$, and that $p' >_G q'$ in T' iff $p >_G q$ in T. Thus by (5) in Lemma 5.1 for T', p', q' it follows that $p' >_G q'$ iff $\neg r \in T' \circ_{Wid} p'$ iff $p' >_{Wid} q'$.

□

Remark. By this theorem, the open problem in [52, p.33], if it is $\mathbf{\Sigma}_2^P$-complete to decide if, under WIDTIO semantics, updating T with p, where the length of p is bounded by a constant, removes $f \in T$ from T is immediately answered positively.

Theorem 5.6 *Let T be a set of Horn formulas and let p, q be Horn formulas, where $|p| \leq k$ for k a constant. Deciding validity of $p > q$ in T is possible in linear time for $> \in \{>_D, >_F, >_B, >_{Win}, >_S\}$.*

Proof: (sketch) Denote by $\mathcal{M}(T)$ the models of T on variables $V(p)$ and by $\mathcal{M}(p)$ the models of p on $V(p)$. The key observation is that under the constraint $|p| \leq k$ both $|\mathcal{M}(T)|$ and $|\mathcal{M}(p)|$ is bounded by the constant 2^k; hence (e.g. by the algorithms in [10]), $\mathcal{M}(T)$ and $\mathcal{M}(p)$ are computable in linear time. As an immediate consequence, $\triangle^{min}(p, T)$ is also computable in linear time (note: $|\triangle^{min}(p, T)| < 2^k$).

$T \circ_S p$ is representable as a disjunction $D = D_1 \vee \cdots \vee D_l$ of Horn formulas, which is efficiently constructible in linear time from T, p, $\triangle^{min}(p, T)$ as follows: for all $M \in \mathcal{M}(T), M' \in \mathcal{M}(p)$ such that $M \triangle M' \in \triangle^{min}(p, T)$, set up a disjunct $f(M') \wedge T[M]$, where $f(M')$ is a conjunction of all literals true in M' that have their variable in $V(p)$, and where $T[M]$ is a Horn formula obtained as follows: First, the truth values of the variables $V(p)$ in M' are substituted into T, and then the formula resulting from this step is simplified accordingly (remove tautologic clauses, false literals etc.) This is certainly possible in linear time. Given D, $p >_S q$ is false in T iff there is some disjunct D_i in D such that $D_i \not\models q$; Since l is bounded by a constant, $p >_S q$ is evaluable in T in linear time.

Concerning Winslett's change operation, for each $M \in \mathbf{M}(T)$, the models $C(M, p)$ of p that are "closest" to M, i.e. $C(M, p) = \triangle_M^{min}(p)$, are computable in linear time, and from this, as for $T \circ_S p$, a disjunction of Horn formulas with no more than a constant number of disjuncts that represents $\mathbf{M}(T \circ_{Win} p)$ is constructible in linear time. Therefore, evaluation of $p >_{Win} q$ is possible in linear time; it is immediately verified that the same holds for $p >_B q$, too. In Dalal's and Forbus's approach, $\triangle^{min}(p, T)$ resp. $C(M, p)$ just has to be "pruned" by eliminating all but the smallest sets in $\triangle^{min}(p, T,)$ resp. the models in $C(M, p)$ closest to M, and then the same procedure as for $>_S$ resp. $>_{Win}$ applies. Since pruning is easily performed in constant time, linear time complexity of testing $p >_D q$ resp. $p >_F q$ is clear.

□

6 Related Results

Note that theory revision, theory update and counterfactual reasoning are particular forms of *nonmonotonic reasoning*. Complexity results for other, yet more common forms of nonmonotonic reasoning in the propositional context have recently been derived by the authors.

The nonmonotonic inference techniques of *closed world reasoning* are widely used in Artificial Intelligence, Database Theory, and Logic Programming [1]. Starting with the (naive) Closed World Assumption (*CWA*) introduced by Reiter [45], several formalizations of closed world reasoning have been developed.

In [12], besides the *CWA*, the following well-known approaches to closed world reasoning are considered: the Generalized Closed World Assumption (*GCWA*) by Minker [39], the Extended Generalized Closed World Assumption (*EGCWA*) by Yahya and Henschen [54], the Careful Closed World Assumption (*CCWA*) by Gelfond and Przymusinska [20], and the Extended Closed World Assumption (*ECWA*) by Gelfond, Przymusinska, and Przymusinski [19]. Circumscription was introduced by McCarthy in [36]. It is known that circumscription coincides with the *ECWA* in the case of propositional logic [19].

In [12] we deal with the complexity of closed world reasoning applied to general propositional theories and to unrestricted clause-sets. In particular, we study the deduction problem, i.e., given formulae F and G, does G follow from F under a certain closed-world inference rule. We show that the deduction problems under the *GCWA*, the *EGCWA*, the *CCWA*, the *ECWA* are all at the second level of PH. In particular, we show that the deduction problem under both the *EGCWA* and the *ECWA* (= propositional circumscription) is $\mathbf{\Pi}_2^P$-complete. We show that this holds also under the restriction that F is in clause form with at most three literals per clause and G is a literal. This is proved to be a limit case for $\mathbf{\Pi}_2^P$-completeness, since if F has at most 2 literals per clause, the deduction problem is co-**NP**-complete for both *EGCWA* and *ECWA* (co-**NP**-hardness is shown in [6]).

For the *GCWA* and the *CCWA*, we provide upper bounds by showing that the inference problem can be solved with $O(\log m)$ calls to a $\mathbf{\Sigma}_2^P$-oracle, where m is the number of propositional variables in the formula F. We also show that testing whether the closure under *CWA* of a formula F is consistent (a co-**NP**-hard problem, but most likely not in co-**NP**) can be done with $O(\log m)$ calls to an **NP**-oracle, where m is the number of propositional variables in F.

The results in [12] can be interpreted as follows: If the polynomial hierarchy does not collapse, then, for all closed world reasoning principles that we consider, the deduction problem is strictly harder than the deduction problem in classical propositional logic. In particular, depending on the chosen closed world principle except for the naive *CWA*, the deduction problem is either complete for $\mathbf{\Pi}_2^P$ or only "mildly" harder than $\mathbf{\Pi}_2^P$. The deduction problem under the *CWA* is complete for Δ_2^P under Turing reductions, and thus not much harder than the deduction problem in the classical propositional logic. It is harder, however, since we show that this problem is neither in **NP** nor in co-**NP** (more precisely, not in the class D^P [43]) unless the polynomial hierarchy collapses.

The complexity of different reasoning tasks in *non-monotonic propositional logics* is investigated in [23]. The systems of logics considered are Reiter's default logic [46], McDermott's and Doyle's nonmonotonic logic [38, 37], Moore's autoepistemic logic [40], and Marek's and Truszczyński's nonmonotonic logic N [35]. All these logics have in common that the semantics of an initially given set of premises is explained through a corresponding set of fixed points (also called extensions or expansions). Given a set of premises, the main reasoning tasks in these logics are

1. testing for existence of a fixed point
2. deciding whether a given formula belongs to at least one fixed point (brave reasoning)
3. deciding whether a given formula belongs to all fixed points (cautious reasoning).

It is shown that for all mentioned logics, the first and the second problem are complete for the class $\mathbf{\Sigma}_2^P$ of the polynomial hierarchy, while the third problem is complete for the complementary class $\mathbf{\Pi}_2^P$.

The proofs of the complexity results for default logic and autoepistemic logic (especially the hardness-parts) are surprisingly simple. To illustrate this, we state completeness proofs for brave and cautious reasoning in propositional default logic.[1]

Theorem 6.1

a.) *(Brave reasoning) Deciding whether a formula ϕ is an element of some extension of a propositional default theory $\langle W, D\rangle$ is $\mathbf{\Sigma}_2^P$-complete. The problem is $\mathbf{\Sigma}_2^P$-complete even if $\langle W, D\rangle$ is a normal default theory.*

b.) *(Cautious reasoning) Deciding whether a formula ϕ belongs to all extensions of a propositional default theory $\langle W, D\rangle$ is $\mathbf{\Pi}_2^P$-complete. The problem is $\mathbf{\Pi}_2^P$-complete even if $\langle W, D\rangle$ is a normal default theory.*

Proof:

Hardness. We prove hardness-results for normal default theories in both cases. For both proofs, let $F = \exists\underline{x}\forall\underline{y}E(\underline{x},\underline{y})$ be a quantified Boolean formula, where $\underline{x} = \langle x_1,\ldots,x_m\rangle, \underline{y} = \langle y_1,\ldots,y_n\rangle$ are tuples of (distinct) Boolean variables and $E(\underline{x},\underline{y})$ is a Boolean formula on $x_1,\ldots,x_m, y_1,\ldots,y_n$. Deciding validity of such a formula is well-known $\mathbf{\Sigma}_2^P$-hard.

a.) (brave reasoning). Transform F into the default theory $\langle\emptyset, D_1\rangle$ with

$$D_1 = \left\{\frac{\top : Mx_1}{x_1}, \frac{\top : M\neg x_1}{\neg x_1}, \ldots, \frac{\top : Mx_m}{x_m}, \frac{\top : M\neg x_m}{\neg x_m}\right\}.$$

The transformation is obviously feasible in polynomial time. Note that the 2^m possible truth value assignments to the variables $x_1,\ldots,x_m$ are in a one-to-one correspondence to the extensions of $\langle\emptyset, D_1\rangle$. Indeed, each extension, must either contain x_i or $\neg x_i$, $1 \le i \le m$, and each possible choice for the x_i's corresponds to a

[1]For the background on default logic, terminology, and notation see [46].

valid extension. It follows that F is valid iff there exists an extension of $\langle \emptyset, D_1 \rangle$ containing E. This proves $\mathbf{\Sigma}_2^P$-hardness.

b.) (cautious reasoning). Translate F to the default theory $\langle \emptyset, D_2 \rangle$,

$$D_2 = \left\{ \frac{\top : Mx_1}{x_1}, \frac{\top : M\neg x_1}{\neg x_1}, \ldots, \frac{\top : Mx_m}{x_m}, \frac{\top : M\neg x_m}{\neg x_m}, \frac{\top : M\neg E}{\neg E} \right\}.$$

Clearly, F is not valid iff $\neg E$ belongs to each extension of $\langle \emptyset, D_2 \rangle$. This proves $\mathbf{\Pi}_2^P$-hardness.

Membership. Let us denote the prerequisite of a default d by $p(d)$, the set of justifications by $j(d)$, and the conclusion by $c(d)$. The set of consequents $\{c(d) : d \in H\}$ of a set H of defaults is denoted by $Cs(H)$. The deductive closure of a set of formulas V is written $Cn(V)$. By results of Zhang and Marek [55, proof of Theorem 6], a propositional default theory $\Delta = \langle D, W \rangle$ has an extension containing ϕ iff there exists a set $D^* \subseteq D$ partitioned into a finite number of strata $D_0^*, \ldots, D_n^*$, such that

1. $D_0^* = \{d : d \in D^* \textit{ where } p(d) \in Cn(W)\}$;

2. for $1 \leq i \leq n$,
$D_i^* = \{d : d \in D^* - \bigcup_{j=0}^{i-1} D_j^*,$
$p(d) \in Cn(W \cup Cs(\bigcup_{j=0}^{i-1} D_j^*))\}$;

3. $\forall d \in D^* \; \forall \beta \in j(d) \; \neg\beta \notin Cn(W \cup Cs(D^*))$;

4. for each $d \in D$, if $p(d) \in Cn(W \cup Cs(D^*))$, and $\forall \beta \in j(d) \; \neg\beta \notin Cn(W \cup Cs(D^*))$, then $d \in D^*$.

5. $\phi \in Cn(W \cup Cs(D^*))$.

Thus, for showing that ϕ belongs to an extension of $\langle D, W \rangle$, it is sufficient to guess an appropriate set $D^* \subseteq D$ (the *generating defaults*) together with a stratification and check conditions 1-5. The checking phase consists of a polynomial number of classical inference tests. Hence, brave reasoning is in $\mathbf{NP}^{\mathbf{NP}} = \mathbf{\Sigma}_2^P$. Let us turn to cautious reasoning. The complementary problem of cautious reasoning consists in deciding whether there exists an extension of $\langle D, W \rangle$ *not* containing ϕ. It is not hard to see that this is the case iff there exists a set $D^* \subseteq D$ satisfying 1-4 but not 5. Thus, the complement of cautious reasoning is in $\mathbf{\Sigma}_2^P$, hence cautious reasoning is in $\mathbf{\Pi}_2^P$.

□

The reader is referred to [23] for more details and for results on other nonmonotonic logics. Similar complexity results on abductive reasoning are given in [13], while in [48] it is shown that different reasoning tasks of De Kleer's assumption-based truth management system (ATMS) [9] are complete for $\mathbf{\Pi}_2^P$.

In summary, it appears that the complexity of almost all approaches to nonmonotonic propositional reasoning is located at the second level of the polynomial hierarchy. As a practical consequence, most formalizations of nonmonotonic reasoning can be polynomially translated into one another, and existing algorithms or proof procedures for one approach can be fruitfully applied to the other.

For instance, $\mathbf{\Pi}_2^P$-completeness of Winslett's method and cautious default reasoning guarantees that the problem of evaluating a counterfactual $p >_{Win} q$ can be polynomially transformed into an equivalent problem of brave reasoning in default logic. Thus, one can take profit of existing sophisticated procedures for default reasoning and of implemented default theorem-provers to solve counterfactuals. Similarly, by the established complexity results, one can polynomially translate most decision problems related to theory update or revision under various approaches to the problem of inferencing under (propositional) circumscription. Practical methods for inferencing under circumscription have been developed by [5, 22, 44]. Note that in the more general context of first-order logic, Winslett already pointed out the translatability of theory-update problems to circumscription [53].

References

[1] K. Apt. Introduction to Logical Programming. Technical report, University of Texas, Department of Computer Science, 1987.

[2] A. J. Bonner. A Logic for Hypothetical Reasoning. In *Proceedings AAAI-88*, pages 480–484, 1988.

[3] A. J. Bonner. Hypothetical Datalog: Negation and Linear Recursion. In *Proceedings PODS-91*, pages 286–300, 1991.

[4] A. Borgida. Language Features for Flexible Handling of Exceptions in Information Systems. *ACM Trans. on Database Syst.*, 10:563–603, 1985.

[5] G. Bossu and P. Siegel. Saturation, Nonmonotonic Reasoning and the Closed World Assumption. *Artificial Intelligence*, 25:13–63, 1985.

[6] M. Cadoli and M. Lenzerini. The Complexity of Closed World Reasoning and Circumscription. In *Proceedings AAAI-90*, pages 550–555, 1990.

[7] M. Dalal. Investigations into a Theory of Knowledge Base Revision: Preliminary Report. In *Proceedings AAAI-88*, pages 475–479, 1988.

[8] M. Dalal. Updates in Propositional Databases. Technical Report DCS-TR-222, Rutgers University, Dept. of Computer Science, Feb. 1988.

[9] J. de Kleer. An Assumption-Based Truth Maintenance System. *Artificial Intelligence*, 28:127–162, 1986.

[10] W. Dowling and J. H. Gallier. Linear-time Algorithms for Testing the Satisfiability of Propositional Horn Theories. *Journal of Logic Programming*, 3:267–284, 1984.

[11] T. Eiter and G. Gottlob. On the Complexity of Propositional Knowledge Base Revision, Updates, and Counterfactuals. Technical Report CD/TR 91-23, Christian Doppler Laboratory for Expert Systems, Vienna University of Technology, Austria, July 1991.

[12] T. Eiter and G. Gottlob. Propositional Circumscription and Extended Closed World Reasoning are $\mathbf{\Pi}_2^P$-complete. Technical Report CD/TR 91-20, Christian Doppler Laboratory for Expert Systems, Vienna University of Technology, Austria, May 1991. To appear in *Theoretical Computer Science.*

[13] T. Eiter and G. Gottlob. Higher Level Complexity Results for Abduction. Technical report, Information Systems Dept., TU Vienna, 1992. forthcoming.

[14] R. Fagin, G. Kuper, J. Ullman, and M. Y. Vardi. Updating Logical Databases. In P. Kanellakis and F. Preparata, editors, *Advances in Computing Research*, volume 3. JAI Press, 1986.

[15] R. Fagin, J. D. Ullman, and M. Y. Vardi. On the Semantics of Updates in Databases. In *Proceedings PODS-83*, pages 352–365, 1983.

[16] K. D. Forbus. Introducing Actions into Qualitative Simulation. In *Proceedings IJCAI-89*, pages 1273–1278, 1989.

[17] P. Gärdenfors. *Knowledge in Flux.* Bradford Books, MIT Press, 1988.

[18] M. Garey and D. S. Johnson. *Computers and Intractability – A Guide to the Theory of NP-Completeness.* W. H. Freeman, New York, 1979.

[19] M. Gelfond, H. Przymusinksa, and T. Przymusinski. On the Relationship Between Circumscription and Negation as Failure. *Artificial Intelligence*, 38:75–94, 1989.

[20] M. Gelfond and H. Przymusinska. Negation as Failure: Careful Closure Procedure. *Artificial Intelligence*, 30:273–287, 1986.

[21] M. L. Ginsberg. Counterfactuals. *Artificial Intelligence*, 30:35–79, 1986.

[22] M. L. Ginsberg. A Circumscriptive Theorem Prover. *Artificial Intelligence*, 39:209–230, 1989.

[23] G. Gottlob. Complexity Results for Nonmonotonic Logics. Technical Report CD/TR 91-24, Christian Doppler Laboratory for Expert Systems, Vienna University of Technology, Austria, August 1991. To appear in *Journal of Logic and Computation.*

[24] G. Grahne. Updates and Counterfactuals. In *Proceedings KR-91*, pages 269–276, 1991.

[25] G. Grahne and A. O. Mendelzon. Updates and Subjunctive Queries. Technical Report KRR-TR-91-4, University of Toronto, CS Dept., July 1991. Also Proceedings Workshop on Nonstandard Queries and Answers, Toulouse, July 1-3, 1991, pp. 33-52.

[26] J. Hartmanis. New Developments in Structural Complexity Theory. *Theoretical Computer Science*, 71:79–93, 1990.

[27] D. S. Johnson. A Catalog of Complexity Classes. volume A of *Handbook of Theoretical Computer Science*, chapter 2. Elsevier Science Publishers B.V. (North-Holland), 1990.

[28] J. Kadin. $P^{NP[O(\log n)]}$ and Sparse Turing-Complete Sets for NP. *Journal of Computer and System Sciences*, 39:282–298, 1989.

[29] H. Katsuno and A. O. Mendelzon. On the Difference between Updating a Knowledge Base and Revising it. In *Proceedings KR-91*, pages 387 – 395, 1991.

[30] H. Katsuno and A. O. Mendelzon. Propositional Knowledge Base Revision and Minimal Change. *Artificial Intelligence*, 52:253–294, 1991.

[31] A. Keller and M. Winslett-Wilkinson. On the Use of an Extended Relational Model to Handle Changing Incomplete Information. *IEEE Transactions on Software Engineering*, SE:11(7):620–633, 1985.

[32] M. Krentel. The Complexity of Optimization Problems. *Journal of Computer and System Sciences*, 36:490–509, 1988.

[33] D. Lewis. *Counterfactuals*. Harvard University Press, Cambridge, MA, 1973.

[34] V. Lifschitz. Closed-World Databases and Circumscription. *Artificial Intelligence*, 27:229–235, 1985.

[35] W. Marek and M. Truszczyńsky. Modal logic for default reasoning. *Annals of Mathematics and Artificial Intelligence*, 1:275–302, 1990.

[36] J. McCarthy. Circumscription – A Form of Non-Monotonic Reasoning. *Artificial Intelligence*, 13:27–39, 1980.

[37] D. McDermott. Non-Monotonic Logic II: Non-monotonic Modal Theories. *Journal of the ACM*, 29:33–57, 1982.

[38] D. McDermott and J. Doyle. Non-Monotonic Logic I. *Artificial Intelligence*, 13:41–72, 1980.

[39] J. Minker. On Indefinite Data Bases and the Closed World Assumption. In *Proceedings of the 6^{th} Conference on Automated Deduction (CADE)*, pages 292–308, 1982.

[40] R. Moore. Semantical Considerations on Non-monotonic Logics. *Artificial Intelligence*, 25:75–94, 1985.

[41] B. Nebel. A Knowledge Level Analysis of Belief Revision. In *Proceedings KR-89*, pages 301–311, 1989.

[42] B. Nebel. Belief Revision and Default Reasoning: Syntax-Based Approaches. In *Proceedings KR-91*, pages 417–428, 1991.

[43] C. H. Papadimitriou and M. Yannakakis. The Complexity of Facets (And Some Facets of Complexity). *Journal of Computer and System Sciences*, 28:244–259, 1984.

[44] T. Przymusinski. An Algorithm to Compute Circumscription. *Artificial Intelligence*, 38:49–73, 1989.

[45] R. Reiter. On Closed-World Databases. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, pages 55–76. Plenum Press, New York, 1978.

[46] R. Reiter. A Logic for Default Reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[47] R. Reiter. A Theory of Diagnosis From First Principles. *Artificial Intelligence*, 32:57–95, 1987.

[48] V. Rutenburg. Complexity Classification in Truth Maintenance Systems. In *Proceedings STACS-91*, number 480 in Lecture Notes in Computer Science, pages 272–283, 1991.

[49] K. Satoh. Nonmonotonic Reasoning by Minimal Belief Revision. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, pages 455–462, 1988.

[50] A. Weber. Updating Propositional Formulas. In *Proceedings of the First Conference on Expert Database Systems*, pages 487–500, 1986.

[51] M. Winslett. Reasoning about Action Using a Possible Models Approach. In *Proceedings AAAI-88*, pages 89–93, 1988.

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

[53] M. Winslett. Circumscriptive Semantics for Updating Databases. *Annals of Mathematics and Artificial Intelligence*, 3:429–450, 1991.

[54] A. Yahya and L. Henschen. Deduction in Non-Horn Databases. *Journal of Automated Reasoning*, 1(2):141–160, 1985.

[55] A. Zhang and W. Marek. On the Classification and Existence of Structures in Default Logic. *Fundamenta Informaticae*, 13(4):485–499, 1990.