Fundamenta Informaticae XX (2002) 1–26
IOS Press

Complexity of the Unique Extension Problem in Default Logic

Xishun Zhao[*,†,C]
Institute of Logic and Cognition
Zhongshan University
Guangzhou, 510275, P. R. China
lcxszhao@hotmail.com

Paolo Liberatore[*]
Department of Computer and System Sciences
University of Rome "La Sapienza"
Via Salaria 113, 00198, Rome, Italy
liberato@dis.uniroma1.it

Abstract. In this paper we analyze the problem of checking whether a default theory has a single extension. This problem is important for at least three reasons. First, if a theory has a single extension, nonmonotonic inference can be reduced to entailment in propositional logic (which is computationally easier) using the set of consequences of the generating defaults. Second, a theory with many extensions is typically weak i.e., it has few consequences; this indicates that the theory is of little use, and that new information has to be added to it, either as new formulae, or as preferences over defaults. Third, some applications require as few extensions as possible (e.g. diagnosis).

We study the complexity of checking whether a default theory has a single extension. We consider the combination of several restrictions of default logics: seminormal, normal, disjunction-free, unary, ordered. Complexity varies from the first to the third level of the polynomial hierarchy. The problem of checking whether a theory has a given number of extensions is also discussed.

Keywords: default logic, complexity, unique extension existence problem

[*]The authors thank Thomas Eiter and the anonymous reviewers for their valuable comments and suggestions for improving the paper.
[†]The first author was partially supported by the Chinese Natural Science Fund project 10161005, and the Chinese Social Science Fund project 02BZX046.
[C]Corresponding author

1. Introduction

One of the most prominent formal approaches to nonmonotonic reasoning is default logic [19]. It is different from standard (propositional) logic because of default rules, which model human reasoning made in prototypical situations when complete information is lacking. Informally, a default rule states a rule of inference (given a fact, we can conclude some other fact) that can only be applied if a given premise cannot be proved to be false. Since the premise is assumed to be true whenever it is not known, a default rule models a reasoning step in which a conclusion is drawn by default (in absence of a contradictory information). A default theory is composed of a propositional part, plus a set of default rules. Formally, it is a pair (D, W) in which W is a set of propositional formulas called initial knowledge while D is a set of default rules.

Default logic provides a powerful tool for knowledge representation and reasoning, as it allows for formalizing rules that are easy to state informally, but would require large knowledge bases in propositional logics. However, there is a price to be paid to gain such advantage. As it is the case for any nonmonotonic reasoning formalism, inference in default logic is computationally hard to do. As proved by Gottlob [7], checking whether a fact is implied by a default theory is at the second level of the polynomial hierarchy, and is therefore harder than propositional inference, which is only at the first level.

Another drawback of default logic is that some theories do not have extensions; if this is the case, no information can be derived from them. This problem motivates the introduction of restricted forms of the default rules: for example, normal defaults always generate extensions. Another solution is the introduction of new semantics: for instance, Przymusinska and Przymusinski [18] introduced the notion of stationary extensions and proved that a stationary extension always exists. Many other variants of the original semantics exist. All of them are harder than propositional logic.

Complexity is due to two causes: first, reasoning requires propositional inference, which is known to be intractable; second, a default theory may have exponentially many extensions, and all of them have to be taken into account in the process of inference. A large number of attempts to lower complexity of default reasoning have been put forward in the literature. Two main directions have been followed: one is to use restricted forms of the propositional part (e.g. Horn); the other one is to use only defaults of a specific form (e.g. normal). However, most restrictions are still intractable.

The large number of extensions can be seens either as a drawback, or as a feature. Since it is due to a large number of conflicting rules, one can either advocate that default reasoning is good as it allows to reason in presence of many conflicts (which is impossible in the standard propositional logic); on the other hand, having too many extensions may be a problem. In any case, knowing the number of extensions of a default theory is key feature of the default theory, for several reasons:

Compilation. If a default theory has a single extension, then it can be translated into propositional logics without changing its consequences (and, therefore, the information it carries). This allows for solving the problem of inference by an algorithm with preprocessing [2, 3, 15]: the preprocessing step is that of translating the default theory; once it is done, queries can be solved in the (easier) propositional calculus.

Expressiviness. A default theory being equivalent to a propositional theory can be seen as implying that we are using a computationally complex formalism (default logic), while the information can be encoded into the propositional logic. In this sense, having one extension is a drawback.

Weakness. If a default theory has many extensions, it usually have few consequences, i.e. it is weak. This may indicate that information has not be encoded incorrectly, or that it is simply deficient. In the latter case, new information has to be added, either in terms of new plain facts (propositional formulae), or in terms of priorities among defaults.

Use of extensions. In some applications of default logic, a small number of extensions is to be preferred. For example, if we encode a model-based diagnosis problem in default logic, then extensions corresponds to possible diagnoses, and a small number of them is clearly to be preferred.

These points show that a default theory having few extensions can be seen as a drawback or as a feature, depending on the point of view. Knowing the number of extensions allows for evaluating the theory, which can be then regarded as good or bad, depending both on the number of extensions and on the point of view.

In this paper we mainly analyze the complexity of the unique extension existence problem, we denote by UEE, that is, the complexity of the problem of determining whether a default theory has exact one extension. We not only study the general case but also analyze various restrictions (normal defaults, etc.)

We begin with the case of prerequisite-free normal defaults in Section 3, and show that UEE is $P^{NP[\log n]}$-complete. When defaults are normal but prerequisite are allowed, complexity goes up to the second level of the polynomial hierarchy, as the UEE is Π_2^P-complete. For semi-normal default theories the UEE problem is in D_2^P, i.e., it can be expressed as the intersection of a Σ_2^P problem and a Π_2^P problem. Since every normal default is semi-normal, the UEE problem for semi-normal default theories is Π_2^P-hard.

In Section 4, we analyze the problem under the strongest restrictions of default logics: disjunction-free, unary, and ordered. While it is foreseeable that the problem is simplified by these assumption, it is somehow surprising to see how much it is. For normal disjunction-free default theories the problem becomes polynomial. For semi-normal disjunction-free theories the problem is in D^P, i.e., it is the intersection of one NP problem and one coNP problem. But whether it is D^P-complete is open. We only show that it is at least as hard as the unique satisfiability problem. Even for unary default theories, the complexity of the UEE problem does not decrease. For ordered theories, since the existence of extensions is guaranteed, the problem becomes a coNP problem. We show that it is coNP-complete. In section 5, we first informally prove that the problem EE(k), deciding if a default has exact k extensions, has the same complexity as the UEE problem, where k is any fixed natural number. We then roughly discuss the reason why the problem UEE has high computational complexity. Our main results are summarized in Table 1.

The complexity analysis reported in this paper extends previous work about complexity of default logic by analyzing the problem of uniqueness of extensions, that have been missing. Indeed, the decision problems analyzed so far are: the existence of extensions; credulous reasoning (deciding whether a formula appears in at least one extension); skeptical reasoning (deciding whether a formula appears in all extensions); model checking (deciding whether an interpretation is a model of all extensions); and extension checking (deciding whether a set of defaults represents an extensions).

The first three problems have been analyzed by Gottlob [7], model checking has been analyzed by Liberatore and Schaerf [16] and by Baumgarden and Gottlob [1], and extension checking has been studied by Rosati [20].

All these works have shown that default logic is usually at the second level of the polynomial hierarchy, and is in general harder than the underlying monotonic reasoning. This leads many researchers to

	NM	semi-NM	NM $\vee$-free	semi-NM $\vee$-free	unary	O-unary
member of	Π_2^P	D_2^P	P	D^P	D^P	coNP
hard for	Π_2^P	Π_2^P	–	Unique-SAT	Unique-SAT	coNP

"NM" stands for "normal"; "$\vee$-free" stands for "disjunction-free"

"O-unary" is for "ordered unary"

Table 1. Complexity of the unique extension existence problem

studying the complexity of default reasoning in special cases [11, 21, 4, 24, 23]. For example, Kautz and Selman [11] analyzed disjunction-free default theories. Although the underlying monotonic inference is, in this very restricted case, very easy, default reasoning is still intractable. Kautz and Selman [11] also introduced some subclasses of disjunction-free default theories, such as unary theories, ordered theories and ordered unary theories, etc. However, default reasoning for most of them remains intractable. Our work is related to that of actually finding the "right" extension, which has been considered by Lang and Marquis [14].

2. Preliminaries

2.1. Default Logic

First we recall some definitions and results on Reiter's default logic. In this paper we only consider finite propositional default theories in which the initial knowledge is consistent. A default is a rule of the form

$$\frac{\varphi : \psi}{\theta},$$

where φ, ψ, θ are propositional formulas. φ is called the prerequisite of the default, ψ is called its justification, and θ is the consequence. Given a default d, we write $p(d)$ for the prerequisite of d, $j(d)$ for its justification, and $c(d)$ for its consequence. Given a set D of defaults, define

$$p(D) = \{p(d)|d \in D\},$$
$$j(D) = \{j(d)|d \in D\},$$
$$c(D) = \{c(d)|d \in D\}.$$

A default is normal if its justification and consequence are the same. A semi-normal default is of the form

$$\frac{\varphi : \theta \wedge \psi}{\theta}.$$

A (normal, semi-normal, respectively) default theory is a pair (D, W) where D is a set of (normal, semi-normal) defaults and W is a consistent set of formulas.

Let (D, W) be a default theory, S be a set of formulas. Then $\Gamma(S)$ is the smallest theory such that

1. $W \subseteq \Gamma(S)$,
2. $\Gamma(S)$ is closed under propositional deduction,
3. For any $d \in D$, if $p(d) \in \Gamma(S)$ and $S \not\vdash \neg j(d)$ then $c(d) \in \Gamma(S)$.

We say a theory E is an extension of (D, W) if and only if $E = \Gamma(E)$.

2.2. Monotonic Rule Systems

A monotonic rule system (R, W) consists of a finite set

$$R = \left\{ \frac{\alpha_1}{\gamma_1}, \cdots, \frac{\alpha_n}{\gamma_n} \right\}$$

of monotonic propositional inference rules and a set W of propositional formulas. Intuitively, a monotonic rule $\frac{\alpha}{\beta}$ means that whenever α is derived, β must be added to the knowledge base. A formula φ is derivable from (R, W), denoted as $(R, W) \vdash \varphi$, if and only if φ can be obtained from W and the axioms of propositional calculus and by a finite number of applications of the modus ponens and some rules in R. The set of all propositional formulas derivable form (R, W) is denoted as $Cn(R, W)$.

Gottlob [8, Lemma 2.2] proved that $\Gamma(S) = Cn(D_0, W)$, where

$$D_0 = \left\{ \frac{\alpha}{\theta} \,\middle|\, \frac{\alpha : \beta}{\theta} \in D, S \not\vdash \neg\beta \right\}.$$

2.3. Equivalent Semantics of Default Logic

Reiter [19] has shown that each normal default theory has at least one extension. However, there exist semi-normal default theories with no extension. Furthermore, normal default logic has semi-monotonicity property, that is, for any two normal default theories (D', W) and (D, W) such that $D' \subseteq D$, every extension E' of (D', W) can be extended to an extension of (D, W). In addition, we also need to recall the proof theory for normal default theories.

Let (D, W) be a normal default theory, φ a propositional formula. We say that a sequence $\delta = (d_1, d_2, \cdots, d_m)$ of defaults from D is a default proof of φ if

1. d_1 is prerequisite-free or $W \vdash p(d_1)$.
2. For each i, $1 \leq i \leq m$, $W \cup c(\{d_1, \cdots, d_{i-1}\}) \vdash p(d_i)$
3. $W \cup c(\delta)$ is consistent and $W \cup c(\delta) \vdash \varphi$

It has been proved that for normal default logic a formula φ has a default proof if and only if it appears in an extension. Usually only minimal default proofs are of interest, in other words, default proofs from which we cannot delete any defaults without losing the properties of a default proof of the required formula.

The following definitions will be used for characterizing the extensions of normal default theories.

Definition 2.1. Let (D, W) be a normal default theory, S a consistent set of formulas such that $W \subseteq S$. Define S_i, D_i by induction as follows. First, $S_0 = S$ and $D_0 = \emptyset$. For any $i > 0$:

$$\begin{array}{rcl} D_{i+1} & = & \{d \in D \mid S_i \vdash p(d), S_i \not\vdash \neg j(d)\} \\ S_{i+1} & = & \begin{cases} S_i \cup c(D_{i+1}) & \text{if it is consistent} \\ S_i & \text{otherwise} \end{cases} \end{array}$$

Intuitively, D_{i+1} is the set of defaults that are applicable in S_i, that is, the defaults whose preconditions are implied by S_i, and whose justifications are consistent with S_i. The set S_{i+1} is the result of applying all of them, that is, adding all their consequences to S_i. Note that D_{i+1} can very well contain conflicting defaults, which results in an inconsistency. In this case, S_{i+1} is equal to S_i. In words, the sequence of S_i's is obtained by applying all applicable defaults until some conflicting defaults are found. This is important, as conflicting defaults may lead to multiple extensions. More precisely, they always lead to multiple extensions if all defaults are normal.

Let us define $\Lambda(S)$ to be the deductive closure of the sequence of S_i, and $\Delta(S)$ be the defaults that are applied in the process of deriving it:

$$\begin{array}{rcl} \Lambda(S) & = & Cn\left(\bigcup_{i \geq 0} S_i\right) \\ \Delta(S) & = & \{d \in D \mid p(d) \in \Lambda(S) \text{ and } c(d) \in \Lambda(S)\} \end{array}$$

Clearly, $\Lambda(S)$ is the deductive closure of $S \cup c(\Delta(S))$.

2.4. Complexity Classes

Next we give a brief review of the relevant notions of complexity theory. Recall that the classes Δ_k^P, Σ_k^P, and Π_k^P of the polynomial hierarchy [10] are defined as follows.

$$\Delta_0^P = \Sigma_0^P = \Pi_0^P = P$$

and for $k \geq 0$,

$$\Delta_{k+1}^P = P^{\Sigma_k^P}, \Sigma_{k+1}^P = NP^{\Sigma_k^P}, \Pi_{k+1}^P = \text{co-}\Sigma_{k+1}^P.$$

In particular, $NP = \Sigma_1^P$, co-$NP = \Pi_1^P$, and $\Delta_2^P = P^{NP}$. Thus Δ_2^P is the class of all problems that are solvable in polynomial time on a deterministic Turing machine with polynomially many calls to an NP oracle. D_k^P is the class of problems which can be described as the intersection of one Σ_k^P problem and one Π_k^P problem. In the literature, D^P is used instead of D_1^P.

The notion of completeness we employ is based on many-one polynomial time transformations. The problem of determining if a formula appears in at least one extension of a normal default theory is Σ_2^P-complete. A well known Π_2^P-complete problem (see e.g. [12]) is deciding the validity of a quantified Boolean formula of the form $\forall p_1 \cdots \forall p_n \exists q_1 \cdots \exists q_m E$, where E is a Boolean formula over variables $\{p_1, \cdots, p_n, q_1, \cdots, q_m\}$. This quantified formula is valid if and only if every truth assignment v to the variables $p_1, \cdots, p_n$ can be extended to the variables $q_1, \cdots, q_m$ so that E is true under v.

The following problem MSA-odd is complete for Δ_2^P [13]: given a satisfiable propositional formula F on variables $p_1, \cdots, p_n$, decide whether the lexicographically maximum truth assignment satisfying F is odd, that is, assigns the truth value 1 to p_n. The lexicographically maximum truth assignment v is recursively defined as

$$v(p_1) = \begin{cases} 1, & \text{if } F(1, p_2, \cdots, p_n) \text{ is satisfiable} \\ 0, & \text{otherwise} \end{cases}$$

$$v(p_{i+1}) = \begin{cases} 1, & \text{if } F(v(p_1), \cdots, v(p_i), 1, p_{i+2}, \cdots, p_n) \text{ is satisfiable} \\ 0, & \text{otherwise} \end{cases}$$

The class $P^{NP[\log n]}$, consists of all problems solvable in polynomial time with $O(\log n)$ queries to an NP oracle [22]. It is known that $P^{NP[\log n]}$ coincides with the class of all problems solvable with parallel n queries to NP oracles. The problem SAT^n_{odd} is complete for this class, where SAT^n_{odd} is the problem of determining whether the number of satisfiable formulas among n CNF-formulas is odd.

The uniquely satisfiability problem (Unique-SAT, for short) is the problem of determining whether a propositional formula has exactly one satisfying truth assignment. It has been proved that Unique-SAT is in D^P and coNP-hard. However, the D^P-completeness of Unique-SAT is still open [12].

3. Complexity of UEE, General Case

3.1. Prerequisite-free Normal Default Theories

We begin the computational analysis of uniqueness of extensions with one of the simplest cases, that of defaults that have no preconditions and are normal. When a default has no precondition, its applicability only depends on the set of justifications. In turns, the justifications coincide with the consequences, which makes defaults of this kind very intuitive: $\frac{:\beta}{\beta}$ means that β should be taken for true whenever possible.

The set of formulae $\Gamma(W)$ has a special role, in this case. When no default has precondition, Γ applies all defaults that are individually applicable. In other words, for each default, its justifications are checked for consistency with W. No default is however applied during this process. Instead, all defaults that are not in contradiction with W are applied only at the end.

The set $\Gamma(W)$ can therefore be inconsistent. Indeed, whenever the application of one default contradicts the justification of another default and vice versa, applying both defaults at the same time makes the result inconsistent. On the other hand, applying one default only leads to an extension. This case is important for us, as it is the only case in which a prerequisite-free normal default theory can have more than one extension.

Lemma 3.1. A normal default theory (D, W) has a single extension, if $\Gamma(W)$ is consistent.

Proof:
Suppose $\Gamma(W)$ is consistent. Gottlob [8, Theorem 4.1] proved that $\Gamma(W) = \Gamma^2(W)$. Then $\Gamma(W)$ is an extension of (D, W). On the other hand, $\Gamma(W)$ is the fixed point of $\Gamma^{2i}(W)$. That is, $\Gamma(W)$ is the smallest stationary extension. Suppose E is an extension of (D, W). E is also a stationary extension (note that every extension is stationary). Thus, $\Gamma(W) \subseteq E$. Since any two different extensions are non-including, it follows that $E = \Gamma(W)$. Consequently, $\Gamma(W)$ is the unique extension of (D, W). □

The converse only holds if all defaults are free of prerequisite.

Lemma 3.2. If a prerequisite-free normal default theory (D, W) has a single extension then $\Gamma(W)$ is consistent.

Proof:
Suppose that $\Gamma(W)$ is inconsistent. We shall prove that the default theory does not have a single extension. Notice that we have assumed that W is consistent. Since $\Gamma(W)$ is defined to be the smallest theory satisfying three conditions of the operator Γ, inconsistency means that no consistent theory satisfies these three conditions at the same time. In particular, no consistent and deductively closed theory W' can imply W, and contain β whenever W is consistent with β. Let D' be as follows:

$$D' = \{d \mid W \wedge j(d) \text{ is consistent } \}$$

This is the set of defaults that are applicable in W, as they are exactly those defaults whose justifications are consistent with W. By assumption, the set of all their consequences are inconsistent with W. However, any single default in it can be applied without leading to inconsistency. This means that the default theory has at least two extensions: take the first by applying one default at time from the first one on; then, take the second by first applying one of the remaining default first. The second one cannot include all defaults of the first one, as the first one is by assumption maximal. □

As a simple consequence, we have the following corollary.

Corollary 3.1. A prerequisite-free normal default theory has a single extension if and only if $\Gamma(W)$ is consistent. Moreover, $\Gamma(W)$ is the unique extension of (D, W).

This corollary allows not only for determining whether a prerequisite-free normal default theory has a single extension, but it also tells the extension itself: indeed, if $\Gamma(W)$ is consistent, it is the only extension of the theory. From a computational point of view, checking its consistency is $P^{NP[\log n]}$-complete.

Theorem 3.1. The problem of determining whether $\Gamma(W)$ is an extension of a prerequisite-free normal default theory (D, W) is $P^{NP[\log n]}$-complete.

Proof:
In order to check consistency of $\Gamma(W)$, we first check all justifications of defaults for consistency with W; we then add to W all their consequences. This algorithm shows that the problem can be solved with an NP-tree: in the leaves we have $|D|$ independent consistency tests; in the root, we have a final check that depends on the result of the previous tests. As Gottlob [9] have shown, this linear number of calls to an NP-oracles can be replaced with a logarithmic number of them, thus showing that the problem is in $P^{NP[\log n]}$.

To prove the hardness we employ a method in [8]. We will present a polynomial time reduction from SAT^n_{odd}, the problem of determining whether the number of satisfiable formulas among n formulas in CNF is odd. Let $F_1, \cdots, F_n$ be n formulas in CNF. Without loss of generality we assume that formulas $F_1, \cdots, F_n$ are mutually disjoint in their propositional variables. We also assume that at least one F_i

is satisfiable (otherwise, we consider formulas p, q, $F_1, \cdots, F_n$). Now we construct a default theory (D, W) as follows. Let $f_1, \cdots, f_n$ be new variables not occurring in any F_i. Define

$$D = \left\{ \frac{: f_1}{f_1}, \cdots, \frac{: f_n}{f_n}, \frac{: f_1 \oplus \cdots \oplus f_n}{f_1 \oplus \cdots \oplus f_n} \right\}$$

and let

$$W = \{f_1 \to F_1, \cdots, f_n \to F_n\},$$

where $\oplus$ stands for exclusive or. Note that $\frac{\top}{f_i}$ is in D_0 if and only if F_i is satisfiable and that if F_i is unsatisfiable then $W \vdash \neg f_i$. Therefore, for each f_i, either $f_i \in \Gamma(W)$ or $\neg f_i \in \Gamma(W)$. In addition, $\frac{\top}{f_1 \oplus \cdots \oplus f_n} \in D_0$ since $W \cup \{f_1 \oplus \cdots \oplus f_n\}$ is consistent by our assumption that some F_i is satisfiable. Now it is easy to see that $Cn(D_0, W)$ is consistent if and only if the number of satisfiable formulas F_i is odd. □

Note that the above proof also proves $P^{NP[\log n]}$-hardness in the case in which there are preconditions. However, the uniqueness of extensions in this case does not necessarily imply the consistency of $\Gamma(W)$. Indeed, the presence of preconditions may result in all defaults initially applicable consistent, but inconsistency appears later. The following default theory shows this case:

$$\left(\emptyset, \left\{ \frac{: a}{a}, \frac{a : \neg a}{\neg a} \right\} \right)$$

Since W is consistent with a, we have $a \in \Gamma(W)$ by condition 3 of the definition of Γ. This implies, on the other hand, that the precondition of the second default is in $\Gamma(W)$, while its justification is still consistent with W. This implies, still by condition 3, that $\neg a \in \Gamma(W)$, thus proving that $\Gamma(W)$ is inconsistent.

On the other hand, this theory has only one extension, namely $Cn(a)$. This being an extension is obvious. It is easy to prove that the theory has no other extensions: $\emptyset$ is easy to rule out as $\Gamma(W) = \Gamma(\emptyset)$ and is larger than $Cn(a)$; the only other possibility is $Cn(\neg a)$, which cannot be an extension as $\Gamma(Cn(\neg a)) = \emptyset$: this holds as $\emptyset$ verifies the three conditions of being an extension.

This example shows why the consistency of $\Gamma(W)$ is not equivalent to the uniqueness of extensions. $\Gamma(W)$ is the smallest set that contains all consequences of applicable defaults, but preconditions are checked against $\Gamma(W)$, while justifications are checked against W. This results in a being in $\Gamma(W)$, while $\neg a$ is still consistent with W, thus making the second default applicable, while it should be not as a is $\Gamma(W)$ (intuitively, the second default should be irrelevant, as its precondition is inconsistent with its justification.)

3.2. Normal Default Theories

In this section, we consider the problem of checking whether a default theory has a single extension, when defaults are normal but have preconditions. As shown in the previous section, the consistency of $\Gamma(W)$ is only a sufficient condition to ensure uniqueness of extensions, but is not necessary.

We show two conditions that implies the multiplicity of extensions. It will later be proved that they are, together, necessary and sufficient. We first need an easy corollary of Lemma 3.1.

Corollary 3.2. For any normal default theory (D, W), let $R_D = \left\{ \frac{p(d)}{c(d)} \middle| d \in D \right\}$. If $Cn(R_D, W)$ is consistent then (D, W) has a unique extension.

The first possible cause of multiple extensions can be explained as follows. Suppose that we start from $S_0 = W$ and compute the sequence of S_i's, but end up with a theory S_i in which two or more defaults are individually applicable (their justifications are individually consistent with S_i) but cannot be applied together (the union of their justifications is not consistent with S_i.) In this case, $S_{i+1} = S_i = \Lambda(W)$, and there are defaults in D_{i+1} that are applicable in $\Lambda(W)$ but are not applied (their consequences are not in $\Lambda(W)$.)

In this case, the sequence S_i does not change from this point on. On the other hand, the process of applying defaults can be continued by applying only a subset of the applicable defaults. When defaults are normal, each choice will lead to an extension. We have therefore found a first condition that leads to multiple extensions.

Lemma 3.3. Let (D, W) be a normal default theory. Suppose there is a default $d \in D$ such that $p(d) \in \Lambda(W)$ but $c(d) \notin \Lambda(W)$ and $\neg c(d) \notin \Lambda(W)$. Then (D, W) has at least two extensions.

Proof:
Define D^* as follows.

$$D^* = \{d \in D \mid p(d) \in \Lambda(W), c(d) \notin \Lambda(W) \text{ and } \neg c(d) \notin \Lambda(W)\}.$$

By assumption, D^* is non-empty. From the definition of $\Lambda(W)$, we see that $\Lambda(W) \cup c(D^*)$ is inconsistent. Now we pick two subsets D_1^* and D_2^* so that $\Lambda(W) \cup c(D_1^*) \cup c(D_2^*)$ is inconsistent while both $\Lambda(W) \cup c(D_1^*)$ and $\Lambda(W) \cup c(D_2^*)$ are consistent. By Corollary 3.2, $(\Delta(W) \cup D_i^*, W)$ has exactly one extension, the deductive closure of $\Lambda(W) \cup c(D_i^*)$, $i = 0, 1$. By the semi-monotonicity, (D, W) has two extensions E_1 and E_2 such that $\Lambda(W) \cup c(D_i^*) \subseteq E_i$, $i = 1, 2$. Since $\Lambda(W) \cup c(D_1^*) \cup c(D_2^*)$ is inconsistent, E_1 and E_2 are different. □

The condition of the above lemma, however, does not cover all possible theories generating multiple extensions. The problem is that, when we move from S_i to S_{i+1}, we are applying all defaults that are applicable in S_i at the same time. However, a default d can be such that $p(d)$ is made true while $c(d)$ is made false at the same time, but these are consequences of different defaults. For example, let us consider (D, W), where $W = \emptyset$, and $D = \{d_1, d_2, d_3\}$, where:

$$d_1 = \frac{: a}{a}, \;\; d_2 = \frac{: b}{b}, \;\; d_3 = \frac{a : \neg b}{\neg b}$$

Clearly, $D_1 = \{d_1, d_2\}$, and $S_1 = Cn(\{a, b\})$. In other words, both d_1 and d_2 are applicable in W, and do not generate inconsistency. This makes d_3 not applicable at all in S_1. On the other hand, once d_1 is applied, d_3 can be applied as well, generating a different extension $Cn(\{a, \neg b\})$.

In general, the problem is that a default d may be left out from D_i because its justification are false when its preconditions are made true. In particular, if the defaults that make the justification false are different from those making the preconditions true, then d can be used to generate an extension. The following lemma formalizes this condition.

Lemma 3.4. Let (D, W) be a normal default theory , $d \in D$ a default such that $\Lambda(W) \vdash \neg c(d)$. Suppose there is a subset $D' \subseteq \Delta(W)$ such that $p(d)$ belongs to an extension of (D', W) (see Corollary 3.2 and Definition 2.1) while $\neg c(d)$ does not. Then (D, W) has at least two extensions.

Proof:
It is easy to see that $c(d)$ belongs to one extension of $(D' \cup \{d\}, W)$. By semi-monotonicity, there is an extension E' of (D, W) such that $c(d) \in E'$. On the other hand, there is an extension E such that $\Lambda(W) \subseteq E$. Since $\neg c(d) \in \Lambda(W)$, E and E' are different. □

The conditions of the two lemmas above can be proved to be the only two cases in which a default theory has more than one extension. The following lemma, indeed, proves that the uniqueness of extension is equivalent to the falsity of one of the above conditions.

Lemma 3.5. Let (D, W) be a normal default theory. Then (D, W) has only one extension if and only if the following conditions hold.

1. There is no $d \in D$ such that $p(d) \in \Lambda(W)$ but $c(d) \notin \Lambda(W)$ and $\neg c(d) \notin \Lambda(W)$.
2. For every default $d \in D$ such that $\Lambda(W) \vdash \neg c(d)$, there is no $D' \subseteq \Delta(W)$ such that $p(d)$ appears in an extension of (D', W) while $\neg c(d)$ is not.

Proof:
The *only if* part directly follows from Lemma 3.3-3.4. For the *if* part we suppose conditions 1 and 2 hold. First we show that, for any default, if it appears in one default proof then it is in $\Delta(W)$. Suppose otherwise, then there is a default proof in which some default is not in $\Delta(W)$. Pick such a default proof $(d_1, \cdots, d_m)$ with the minimum length. Then $d_i \in \Delta(W)$ for each $1 \leq i < m$. Then $p(d_m) \in \Lambda(W)$. By condition 1, we have $\neg c(d_m) \in \Lambda(W)$. However, by using condition 2, we have $W \cup c(\{d_1, \cdots, d_{m-1}\}) \vdash \neg c(d_m)$. That is, $(d_1, \cdots, d_m)$ is not a default proof, a contradiction. Consequently, every extension is a subset of $\Lambda(W)$. Since (D, W) has at least one extension, it follows that $\Lambda(W)$ is the unique extension of (D, W). □

Lemma 3.5 shows that the uniqueness check can be done by first computing $\Lambda(W)$, and then checking whether conditions 1 and 2 in Lemma 3.5 hold. $\Lambda(W)$ can be computed with a polynomial number of queries to an NP oracle. Condition 1 can also be verified with a polynomial number of queries. However, checking Condition 2 cannot be done in the same way, as it requires checking all subsets $D' \subseteq \Delta(W)$, and these are exponentially many. We can actually show that the problem cannot be simplified (unless the polynomial hierarchy collapses): the unique extension existence problem is Π_2^P-complete, and cannot therefore be solved with a polynomial number of queries to an NP oracle.

We first introduce some auxiliary problems. The first one is related to the second condition of the lemma above; it consists in checking whether a consistent subset of a monotonic rule system implies a formula.

Problem P1

Instance: A monotonic rule system (R, W) such that $Cn(R, W)$ is inconsistent and a formula φ.

Query: Is there $R' \subseteq R$ such that $(R', W) \vdash \varphi$ and $Cn(R', W)$ is consistent?

Lemma 3.6. Problem P1 is Σ_2^P-complete.

Proof:
The following procedure to solve P1 shows it is in Σ_2^P. Guess a subset $R' \subseteq R$. Then check whether $(R', W) \vdash \varphi$ and $Cn(R', W)$ is consistent. If this condition holds then return yes. It is easy to see that this is a non-deterministic Turing machine with polynomially many calls to the oracle of consistency and inference checking. Thus problem P1 is in Σ_2^P.

Hardness is proved by reduction from credulous default reasoning (determining if a formula occurs in at least one extension). Given a normal default theory (D, W) and a formula φ, we consider a new atom p, and define R as follows.

$$R = \left\{ \frac{\alpha}{\beta} \,\middle|\, \frac{\alpha : \beta}{\beta} \in D \right\} \cup \left\{ \frac{W}{p}, \frac{W}{\neg p} \right\}$$

The set R and the formula φ are a valid instance of P1, as $Cn(R, W)$ is inconsistent: this is due to the last two rules, which enforce both p and $\neg p$ to be consequences. By putting only one of the two rules in $R' \subset R$ we remove this source of inconsistency, but R' may still be inconsistent due to the other rules, that corresponding to the defaults of the original theory.

We now prove that this is a reduction from credulous default reasoning to P1: φ belongs to at least one extension of (D, W) if and only if there exists $R' \subseteq R$ such that $(R', W) \vdash \varphi$ and $Cn(R', W)$ is consistent.

Let us first prove the direction from the left to the right. Suppose φ appears in one extension. Then φ has a default proof, say, $(d_1, \cdots, d_n)$. Let

$$R' = \left\{ \frac{p(d_1)}{c(d_1)}, \cdots, \frac{p(d_n)}{c(d_n)} \right\}.$$

Clearly, $(R', W) \vdash \varphi$ and $Cn(R', W)$ is consistent.

To prove the converse, suppose there is $R' \subseteq R$ such that $(R', W) \vdash \varphi$ and $Cn(R', W)$ is consistent. The last two rules in R are not helpful for the proof of φ; therefore, if either one is in R', it can be removed without making φ not implied. We therefore assume that none of them is in R'. By replacing each rule in R' with the original default, we obtain a default proof of φ. Thus, φ appears in at least one extension of (D, W). □

The second auxiliary problem we consider is also related to the second condition of uniqueness. Indeed, P2 is based on the existence of a subset of rules that implies a formula but does not imply another one.

Problem P2

Instance: A monotonic rule system (R, W) such that $Cn(R, W)$ is consistent, two formulas φ, ψ such that $(R, W) \vdash \psi$.

Query: Does there exist $R' \subseteq R$ such that $(R', W) \vdash \varphi$ but $(R', W) \not\vdash \psi$?

This problem has the same complexity of the previous one, that is, it is Σ_2^P-complete. In particular, hardness is proved by reduction from P1.

Lemma 3.7. Problem P2 is Σ_2^P-complete.

Proof:
Membership is proved in the same way which has been done for P1 in Lemma 3.6. Hardness is proved by reduction from P1. Let (R, W) and a formula φ be an instance of P1. We build an instance of P2 as follows. Let t be a new variable. Define W_t and R_t as follows.

$$\begin{array}{rcl} W_t & = & \{F \vee t \mid F \in W\} \\ R_t & = & \left\{ \dfrac{\alpha \vee t}{\beta \vee t} \middle| \dfrac{\alpha}{\beta} \in R \right\}. \end{array}$$

$Cn(R_t, W_t)$ is consistent thanks to the variable t that is disjoined to all involved formulae. This is why t has been introduced, indeed. However, since $Cn(R, W)$ is inconsistent, $(R_t, W_t) \vdash t$. The lemma follows from the following statement.

> There exists $R' \subseteq R$ such that $(R', W) \vdash \varphi$ and $Cn(R', W)$ is consistent if and only if there is $R'' \subseteq R_t$ such that $(R'', W_t) \vdash \varphi \vee t$ and $(R'', W_t) \not\vdash t$.

Informally, this statement is true as deleting rules from R to remove inconsistency corresponds to deleting rules from R_t to make t not implied any more.

First we show the direction from the left to the right. Suppose $R' \subseteq R$, $(R', W) \vdash \varphi$ and $Cn(R', W)$ is consistent. It is easy to see that $(R'_t, W_t) \vdash \varphi \vee t$ and $(R'_t, W_t) \not\vdash t$. Conversely, suppose $R'' \subseteq R_t$, $(R'', W_t) \vdash \varphi \vee t$ and $(R'', W_t) \not\vdash t$. Let R' be obtained from R'' by dropping t. Since $(R'', W_t) \not\vdash t$, $Cn(R'', W_t \cup \{\neg t\})$ is consistent. This, together with $(R'', W_t) \vdash \varphi \vee t$, implies that $Cn(R', W)$ is consistent and $(R', W) \vdash \varphi$. □

Now we show that the unique extension existence problem for normal default theories is Π_2^P-complete. Hardness is shown by reduction from P2.

Theorem 3.2. The problem of determining whether a normal default theory has a unique extension is Π_2^P-complete.

Proof:
We show that the problem of determining if a normal default theory (D, W) has at least two different extensions is in Σ_2^P. We employ the following procedure to determine if (D, W) has at least two extensions. First, guess two permutations of defaults in D. For each permutation compute the corresponding extension [17]. If we get two different extensions then return yes. To compute an extension from a permutation costs polynomially many calls to the oracle of consistency and inference checking. Thus, the problem is in Σ_2^P.

Hardness is proved by reduction from problem P2. Given a monotonic rule system (R, W) such that $Cn(R, W)$ is consistent, and given two formulas φ, ψ such that $(R, W) \vdash \psi$. Define

$$D = \left\{ \frac{\alpha : \beta}{\beta} \middle| \frac{\alpha}{\beta} \in R \right\} \cup \left\{ \frac{: \psi}{\psi}, \frac{\varphi : \neg\psi}{\neg\psi} \right\}.$$

We only need to prove the following statement.

> There exists $R' \subseteq R$ such that $(R', W) \vdash \varphi$ and $(R', W) \not\vdash \psi$ if and only if (D, W) has at least two extensions.

The last two defaults, if they are at some point both applicable, generate two different extensions, one implying ψ and one implying $\neg\psi$. Applicability of both of them is only possible if, applying the other defaults, we can make φ true while ψ is not implied (i.e. $\neg\psi$ is consistent). This corresponds to finding a subset of rules that implies φ but not ψ.

Let us formally prove the claim. We first show the direction from the left to the right. Suppose $R' \subseteq R$, $(R', W) \vdash \varphi$ and $(R', W) \not\vdash \psi$. Let D' be obtained from R' by replacing each monotonic rule $\frac{\alpha}{\beta}$ with $\frac{\alpha:\beta}{\beta}$. From the assumption, we know that $\neg\psi$ appears in an extension of $(D' \cup \{\frac{\varphi:\neg\psi}{\neg\psi}\}, W)$. By the semi-monotonicity, $\neg\psi$ must occur in an extension of (D, W). It is easy to see ψ appears in some extension of (D, W). Therefore (D, W) has at least two extensions.

Now we prove the other direction. Suppose (D, W) has at least two extensions. Let

$$D_R = \left\{ \frac{\alpha : \beta}{\beta} \middle| \frac{\alpha}{\beta} \in R \right\}.$$

Since $Cn(R, W)$ is consistent, then by Corollary 3.1, (D_R, W) has only one extension which contains ψ because $(R, W) \vdash \psi$. Therefore, (D, W) must have an extension such that $\neg\psi \in E$ and $\neg\psi$ is obtained by applying $\frac{\varphi:\neg\psi}{\neg\psi}$. As a consequence, there is a default proof containing $\frac{\varphi:\neg\psi}{\neg\psi}$. From such a default proof we can easily obtain a subset $R' \subseteq R$ such that $(R', W) \vdash \varphi$ but $(R', W) \not\vdash \psi$. □

As we have seen in the previous section, $\Gamma(W)$ has a special role, when there are no preconditions. Indeed, it is an extension if and only if it is the unique extension of the theory. Preconditions, however, make it a necessary condition only: if $\Gamma(W)$ is an extension, it is unique, but the converse does not necessarily hold.

Another necessary condition for uniqueness is $\Lambda(W)$ being an extension. We can prove that this problem is actually easier than checking uniqueness. This is important, as we can solve this problem as a preliminary step of checking uniqueness: if $\Lambda(W)$ is not an extension, then we know that the default theory does not have a single extension without solving the Π_2^p-hard problem of checking uniqueness.

Theorem 3.3. For a normal default theory (D, W), the problem of determining whether $\Lambda(W)$ is an extension is Δ_2^P-complete.

Proof:
A finite base of $\Lambda(W)$ can be computed with polynomially many calls to an NP-oracle. It is not hard to see that $\Lambda(W)$ is an extension if and only if condition 1 in Lemma 3.5 holds. To verify if $\Lambda(W)$ satisfies condition 1 also needs polynomially many calls to the oracle. Thus, the problem is in Δ_2^P.

Hardness is proved by reduction from the problem MSA-odd: the definition of this problem is in Section 2; it is Δ_2^P-hard [13]. Let φ be an instance of this problem; let $p_1, \ldots, p_n$ be its variables. For each variable p_i we introduce a new variable t_i. We also need an additional variable t_0. Let $W = \{\varphi, t_0\}$, and D be the set of the following defaults.

$$D = \bigcup_{1 \leq i \leq n} \left\{ \frac{t_{i-1} : p_i}{p_i}, \frac{t_{i-1} \wedge p_i : t_i}{t_i}, \frac{t_{i-1} \wedge \neg p_i : t_i}{t_i} \right\} \cup \left\{ \frac{\neg p_n : a}{a}, \frac{\neg p_n : \neg a}{\neg a} \right\}$$

The idea is as follows. When computing $\Lambda(W)$, at the first step the only applicable defaults are those corresponding to $i = 1$. Indeed, t_0 is the only variabile t_i that is in W. In particular, the first default is only applicable if φ is consistent with p_1. This means that p_1 is added to the set we are building if and only if it is consistent with φ. In the other case, $\varphi \models \neg p_1$. Either way, either p_1 or $\neg p_1$ is now implied. This result in making t_1 be implied as well, which makes the defaults with $i = 2$ applicable.

Therefore, defaults are applied in order: first the ones with $i = 1$, then the ones with $i = 2$, etc. At each step, p_i is added to the extension if and only it is consistent with it. Therefore, at the end of this process we have the lexicographically maximal model of φ. The presence of p_n blocks the application of the last two defaults and we end up with $\Lambda(W)$ being the only extension of the theory. On the other hand, if the lexicographically maximal model contains $\neg p_n$, both the last two defaults are individually applicable, but cannot be applied together. Therefore, their conclusions are not in $\Lambda(W)$, while applying either one generates an extension.

This informally shows that $\Lambda(W)$ is an extension if and only if p_n is implied by the lexicographically maximal model of φ. Formal proof is omitted as it is easy to derive.

□

3.3. Semi-Normal Default Theories

The problem of checking whether a default theory has a single extension can be expressed as checking whether it has at least one extension, and it has no more than one of them. The condition (checking existence of extensions) has already been studied: Gottlob [7] and Stillman [21] proved it to be in Σ_2^P. It is easy to see that checking whether a default theory has more than one extension is in Σ_2^P as well. Since what we need is the converse of this problem, we conclude that UEE can be expressed as the intersection of a problem in Σ_2^P with a problem in Π_2^P. This holds for all default theories.

Theorem 3.2 implies that UEE is Π_2^P-hard for normal default theories, and this result extends to the general case. It therefore remains a little gap between the lower and the upper bound we have. However, it is possible to prove that UEE is at least as hard as another uniqueness problem on QBFs.

We consider quantified Boolean formulas that are valid whenever there exists an unique truth assignment on a part of the variables that makes the resulting formula valid. Such formulas are written as follows.

$$\Phi = \exists! y_1 \cdots \exists! y_m \forall x_1 \cdots \forall x_n \varphi.$$

This quantified Boolean formula is defined valid if there exists exactly one truth assignment v to the propositional variables $y_1, \cdots, y_m$ such that, however v is extended to the propositional variables

$x_1, \cdots, x_n$, formula φ is true. To check whether Φ is true, we first check that $\exists y_1 \cdots \exists y_m \forall x_1 \cdots \forall x_n \varphi$ is true. This is a Σ_2^P problem. Then we check that there is at most one truth assignment v to $y_1, \cdots, y_m$ such that v makes φ to be tautology. This is a Π_2^P problem. Therefore, the problem of deciding the validity of a quantified Boolean formula with the above form can be described as the intersection of a problem in Σ_2^P and one in Π_2^P. The next lemma shows that it is Π_2^P-hard.

Lemma 3.8. The problem of deciding the validity of a quantified Boolean formula with the form $\exists! y_1 \cdots \exists! y_m \forall x_1 \cdots \forall x$ is Π_2^P-hard.

Proof:
This lemma directly follows from the fact that $\exists y_1 \cdots \exists y_m \forall x_1 \cdots \forall x_n \varphi$ is false if and only if $\exists! y_0 \exists! y_1 \cdots \exists! y_m \forall x_1 \cdots \forall x$ $\neg y_0) \vee (y_0 \wedge y_1 \wedge \cdots \wedge y_m)$ is true, where y_0 is a new variable. □

We now prove that this problem can be reduced to UEE in polynomial time.

Theorem 3.4. The UEE problem for semi-normal default theories is at least as hard as the problem of deciding the validity of a quantified formula of the form $\exists! y_1 \cdots \exists! y_m \forall x_1 \cdots \forall x_n \varphi$.

Proof:
Let Φ be the quantified Boolean formula $\exists! y_1 \cdots \exists! y_m \forall x_1 \cdots \forall x_n \varphi$, and let $t_1, \cdots, t_m$ and t be new variables. Define $W = \emptyset$ and D to be the set consisting of the following defaults

(I) The defaults

$$\frac{\neg\varphi : \neg t_1 \wedge \cdots \wedge \neg t_n \wedge t}{t}, \quad \frac{\neg\varphi : \neg t_1 \wedge \cdots \wedge \neg t_n \wedge \neg t}{\neg t},$$

(II) For each $i = 1, \cdots, m$, the defaults

$$\frac{: y_i \wedge \varphi}{y_i}, \frac{: \neg y_i \wedge \varphi}{\neg y_i}$$

(III) For each $i = 1, \cdots, m$, the defaults

$$\frac{y_i : t_i}{t_i}, \quad \frac{\neg y_i : t_i}{t_i}$$

(IV) Finally, the default

$$\frac{t_1 \wedge \cdots \wedge t_n : \neg\varphi}{\neg\varphi}$$

The QBF formula must be valid if and only if the default theory has a single extension. Therefore, whenever the QBF formula is not valid, the corresponding default theory is made having zero or more than one extension. In particular, if φ is contradictory, the defaults in group (I) generate two extensions.

The default in the group (II) are defined in such a way either one for each i has to be applied, if φ is not contraditory. Intuitively, each time we apply a default we make an arbitrary choice on setting y_i to true or false. Whenever such an assignment is done, one of the default in group (III) is applicable, which results in t_i being set whenever y_i is assigned a value.

Up to this point, the values of y_i are set arbitrarily, while t_i is set whenever y_i is assigned a value. What is still left is a way to guarantee that this is an extension if and only if φ is valid in this partial assignment, regardless of the value of the variables x_i's. This would imply that the QBF is valid if and only if the default theory has exactly one extension. Default (IV) is defined to this purpose: if φ is valid, the last default is not applicable; no other default is applicable, and the theory we have built is already an extension. If φ is instead not valid, then $\neg\varphi$ is consistent, and the last default is applicable. This leads to contradicting the justification of all defaults in group (II). Therefore, this is not an extension.

Formally, we will prove the following statement.

Φ is valid if and only if (D, W) has exactly one extension.

Before proving the statement, we show a related claim. Let v be a truth assignment to the propositional variables $y_1, \cdots, y_m$ that satisfies φ however it is extended to $x_1, \ldots, x_n$. Let L_v be the set of literals (variables and their negations) satisfied by v:

$$L_v := \{l \mid l \in \{y_i, \neg y_i\} \text{ for some } i \text{ and } v(l) = 1\}.$$

Let $E_v = Cn(L_v \cup \{t_1, \ldots, t_m\})$. We prove that E is an extension of (D, W). First, by assumption φ is implied from L_v, hence $\varphi \in E_v$. The defaults in group (II) make each literal in L_v to be in $\Gamma(E_v)$. Moreover, thanks to the defaults in group (III), each t_i is also in $\Gamma(E_v)$. Thus $E_v \subseteq \Gamma(E_v)$. By Gottlob's Lemma 2.2 in [8] it is easy to see that $\Gamma(E_v) \subseteq E_v$. Hence $E_v = \Gamma(E_v)$. That is to say E_v is an extension of (D, W).

Suppose E is an extension of (D, W) such that $\varphi \in E$. It is easy to see that, for each $i = 1, \cdots, m$, either y_i or $\neg y_i$ is in E, and that all $t_i, (i = 1, \cdots, m)$ are in E. Define the truth assignment v_E to $y_1, \cdots, y_m$ as follows. For each $i = 1, \cdots, m$,

$$v_E(y_i) = \begin{cases} 1, & \text{if } y_i \in E \\ 0, & \text{otherwise} \end{cases}$$

Note that $\varphi \notin W$ and there is no default whose consequence is φ. Thus, φ must be provable from $E \cap \{y_1, \neg y_1, \cdots, y_m, \neg y_m\}$. As a result, however v_E is extended to $x_1, \cdots, x_n$, φ is true under v_E.

Now we come the proof of the statement.

$(\Rightarrow)$. Suppose Φ is valid. Let v be the unique truth assignment to $y_1, \ldots, y_m$ such that, in whatever way v is extended, φ is true under v. We already proved that E_v is an extension of (D, W). Let E be an arbitrary extension of (D, W). We claim $\neg\varphi \notin E$. Suppose, by contrary, $\neg\varphi \in E$. Then no t_i is in E. This can be proved as follows. Suppose t_i is in E; then, either y_i or $\neg y_i$ is in E. Hence one of defaults in group (II) is applicable. On the other hand, $\neg\varphi$ blocks the use of defaults in group (II), a contradiction. Thus, the final default is not applicable and hence $\neg\varphi$ must be a tautology. However, φ is satisfiable by assumption, a contradiction. Therefore, $\neg\varphi \notin E$. Then, for each i, either y_i or $\neg y_i$ is in E, hence each t_i is in E. Then we get that φ is in E (otherwise, $\neg\varphi \in E$ by the last default). Consequently, v_E is also a truth assignment satisfying φ regardless the truth values of x_i, $i = 1, \cdots, n$. By the uniqueness of v, we get $v = v_E$. Then $E_v \subseteq E$, and hence $E = E_v$. That is to say, (D, W) has exactly one extension.

$(\Leftarrow)$. Suppose (D, W) has exactly one extension, say E. If $\neg\varphi \in E$, then $\neg\varphi$ would be a tautology (see the above paragraph). Therefore, because of the defaults in group (I) (D, W) would have two extensions, one containing t and the other containing $\neg t$. Thus $\neg\varphi$ is not in E. Then, for each i, either y_i or $\neg y_i$ is in E, hence each t_i is in E. Suppose $\varphi \notin E$. Then by the final default $\neg\varphi \in E$, a contradiction.

Thus $\varphi \in E$. Then, v_E satisfies φ regardless the truth values of x_i, $i = 1, \cdots, n$. Suppose v is another truth assignment with the same property as v_E. Then, E_v is an extension of (D, W) different from E, contradicting the uniqueness of E. Thus Φ is valid.

Therefore, we prove the theorem. □

4. Disjunction-Free Default Theories

A default d is disjunction-free if no formula in it contains the disjunction connective $\vee$, that is, it is in the form:

$$d = \frac{a_1 \wedge \cdots \wedge a_k : b_1 \wedge \cdots \wedge b_m \wedge c_1 \wedge \cdots \wedge c_n}{b_1 \wedge \cdots \wedge b_m},$$

where a_i, b_i, c_i are literals. In this case, we consider $p(d), j(d), c(d)$ as sets of literals. Normality and semi-normality are defined as before. A disjunction-free default theory is a default theory (D, W) in which W is a set of literals and D is a set of disjunction-free defaults. For a normal default theory (D, W), we identify $\Lambda(W)$ with the set of literals in $\Lambda(W)$. It has been shown that disjunction-free default reasoning is still intractable [11].

Suppose (D, W) is a disjunction-free default theory. Let D' be obtained from D by deleting all the defaults in which there is a literal x such that $\neg x \in W$. It is easy to see that (D, W) and (D', W) have the same extensions. We can therefore assume, without loss of generality, that the literals of the defaults in D are all consistent with W.

4.1. Normal Disjunction-Free Default Theories

We have shown, in a previous section, that the unique extension problem for normal default theories is Π_2^P-complete. In this section we will show that the problem is polynomial for disjunction-free normal default theories.

Before technically proving this claim, we give some intuitions about it. First of all, inference is polynomial when all involved formulas are conjunctions of literals. This clearly makes checking applicability of defaults polynomial. In the general case, Lemma 3.5 shows when a default theory has exactly one extension: when $\Lambda(W)$ is an extension and no default is applicable after a subset of $\Delta(W)$ has been applied, while it is not in $\Lambda(W)$.

The first condition is clearly polynomial: start with W, and then apply all defaults that are applicable in it, until either contradiction arises or no other default is applicable. Since applicability is polynomial, the whole process is polynomial.

One of the outcomes of the evaluation of $\Lambda(W)$ is the set of defaults that have been applied so far $\Delta(W)$. This set is needed for checking the second condition of uniqueness. In particular, we have to consider each possible subset of it, and check which defaults are applicable. In general, there are exponentially many such subsets. In this case, however, we only need to consider a specific subset.

Definition 4.1. Let (D, W) be a normal disjunction-free default theory. Let $d \in D$ be such that $\neg a \in \Lambda(W)$ for some $a \in c(d)$. Define:

$$\Delta(W,d) = \Delta(W)\setminus \left\{d' \in \Delta(W) \mid \text{ there exists } x \in c(d) \text{ such that } \neg x \in p(d') \cup c(d')\right\}.$$

That is to say, for any default $d^* \in \Delta(W)$, $d^* \in \Delta(W,d)$ if and only if every literal in d^* is consistent with $c(d)$.

The point is that, while checking for subsets of $\Delta(W)$ that makes d applicable, we only have to consider the defaults that do not involve (as a precondition or as a consequence) a literal whose opposite is in d. We can therefore restrict our attemption to the defaults in $\Delta(W,d)$.

A further simplification of the problem is due to the fact that the defaults in $\Delta(W)$ do not conflict to each other, as this implies the same for the defaults in $\Delta(W,d)$. Therefore, application of them corresponds to generating the unique extension of $(\Delta(W,d),W)$, in which only a part of the defaults in $\Delta(W,d)$ can be applied. Moreover, by construction, the justification of d is necessarily consistent with the unique extension of $(\Delta(W,d),W)$.

Lemma 4.1. Let (D,W) be a normal disjunction-free default theory. (D,W) has only one extension if and only if the following conditions hold.

1'. There is no default $d \in D$ such that $p(d) \subseteq \Lambda(W)$ but $c(d) \not\subseteq \Lambda(W)$ and $\neg x \notin \Lambda(W)$ for all $x \in c(d)$.

2'. For any default $d \in D$ such that $\neg x \in \Lambda(W)$ for some $x \in c(d)$, $p(d)$ is not included in the unique extension of $(\Delta(W,d),W)$.

Proof:
Condition 1 in Lemma 3.5 and condition 1' are in fact the same. We only have to prove that 2 and 2' are the same as well. We first prove the *only if* part. Suppose (D,W) has exactly one extension. Then, by Lemma 3.5, Condition 2 in Lemma 3.5 hold. By the construction of $\Delta(W,d)$, the unique extension of $(\Delta(W,d),W)$ is consistent with $c(d)$. Thus, by Condition 2 we know that $p(d)$ is not included in the unique extension.

To prove the converse, we assume that Condition 2' holds, and prove that Condition 2 holds as well. Suppose, by contrary, that there is a default $d \in D$ such that $\neg x \in \Lambda(W)$ for some $x \in c(d)$ and there exists $D' \subseteq \Delta(W)$ such that $p(d)$ is included in the unique extension of (D',W), while $c(d)$ is consistent with the extension. Then, $D' \subseteq \Delta(W,d)$. Hence $p(d)$ is included in the extension of $(\Delta(W,d),W)$, contradicting the assumption that Condition 2' holds. This completes the proof. □

For a disjuction-free normal default theory (D,W), computing $\Lambda(W)$ needs polynomial time since consistency checking here becomes trivial. To inspect Condition 1' also needs only polynomial time for the same reason. For any default $d \in D$, $(\Delta(W,d),W)$ has exactly one extension, which can be computed in polynomial time. It follows that condition 2' can be checked in polynomial time.

Theorem 4.1. The UEE problem for normal disjunction-free default theories can be solved in polynomial time.

4.2. Semi-Normal Disjunction-Free Default Theories

Part of the problem of uniqueness is checking whether a default theory has at least one extension. This problem is already known to be NP-complete [11, 21, 5]. The other part of the problem is that of checking whether the default theory has two or more extensions or not. This problem is NP-hard. Since UEE can be solved by checking whether a default theory has at least one extensions, and *it has not* two or more extensions, it can be expressed as the intersection of a problem in NP and a problem in coNP. Therefore, it is in D^P when no disjunction is allowed. Proving hardness is however difficult. We show that the problem is at least as hard as the problem of checking whether a propositional formula has exactly one extension.

Theorem 4.2. For disjunction-free semi-normal default theories, the UEE problem is at least as hard as Unique-SAT.

Proof:
Let $\varphi = \alpha_1 \wedge \cdots \wedge \alpha_m$ be a 3-CNF formula with $\alpha_i = (a_i \vee b_i \vee c_i)$. We introduce a new variable t, and a new variable t_i for each clause α_i. Define $W = \emptyset$ and D to be the set consisting of the following groups of defaults.

(I) For each variable $p \in var(\varphi)$, the defaults

$$\frac{: p \wedge t}{p}, \quad \frac{: \neg p \wedge t}{\neg p}.$$

(II) For each $i = 1, 2, \cdots, m$, the defaults

$$\frac{a_i : t_i}{t_i}, \quad \frac{b_i : t_i}{t_i}, \quad \frac{c_i : t_i}{t_i}.$$

(III) For each $i = 1, 2, \cdots, m$, the defaults

$$\frac{\neg a_i \wedge \neg b_i \wedge \neg c_i : \neg t_i}{\neg t_i}$$

(IV) The default

$$\frac{t_1 \wedge t_2 \wedge \cdots \wedge t_m : t}{t}$$

(V) For each $i = 1, \cdots, m$, defaults

$$\frac{\neg t_i : \neg t}{\neg t}.$$

Informally, using defaults in (I), we can get a truth assignment. Since we want this theory to have a single extension if and only if φ has a single model, the other defaults must generate a single extension for each model of the formula, and no extension if the truth assignment does not satisfies φ.

Defaults in (II) and (IV) are defined in such a way that, if the truth assigment satisfies φ, then $t, t_1, \ldots, t_m$ are derived. Defaults in (III) are instead aimed at making all t_i false if the truth assigment falsifies φ. It is easy to see that exactly one extension is obtained in the first case (no other default is applicable). In the second case, the last default is applicable, generating $\neg t$ which contradicts the justification of the defaults in group (I), and this is therefore not an extension.

Formally, we prove the theorem by showing that the following statement is true.

φ is in Unique-SAT if and only if $(D, \emptyset)$ has exactly one extension

Before proving the statement we do some preparation.

1. Let v be a truth assignment such that $v(\varphi) = 1$. Define $E_v = Cn(L_v \cup \{t_1, \cdots, t_m, t\})$: by construction, E_v is an extension of $(D, \emptyset)$.

2. Suppose E is an extension of $(D, \emptyset)$; we prove that $\neg t \notin E$. Suppose otherwise; then $\neg t$ must be obtained from one default in group (V), say $\frac{\neg t_i : \neg t}{\neg t}$. Thus $\neg t_i$ is in E. Since E is an extension, $\neg t_i$ must be obtained by applying the default in group (III). Hence $\neg a_i, \neg b_i, \neg c_i$ belong to E. To obtain $\neg a_i, \neg b_i, \neg c_i$ we have to apply some defaults in group (I). On the other hand, the fact that $\neg t \in E$ would block the use of defaults in group (I), a contradiction. Thus $\neg t \notin E$.

3. We show $t \in E$, if E is an extension of $(D, \emptyset)$. Suppose $t \notin E$. Since $\neg t \notin E$ we have that no $\neg t_i$ belongs to E (otherwise from defaults in group (V) we would get $\neg t$). Since neither t nor $\neg t$ is in E, we get that at least one t_i is not in E (otherwise, t would be in E by the default in group (IV)). Now we know that there is some i such that neither t_i nor $\neg t_i$ is in E. From the defaults in group (III), we get $\{\neg a_i, \neg b_i, \neg c_i\} \not\subseteq E$. Without loosing any generality, we assume that $\neg a_i \notin E$. Then the default $\frac{: a_i \wedge t}{a_i}$ in group (I) is applicable. Hence $a_i \in E$. Since $\neg t_i \notin E$ the default $\frac{a_i : t_i}{t_i}$ is applicable, we have $t_i \in E$. However, we have previously proved that $t_i \notin E$, a contradiction. Consequently, $t \in E$. Then t must be obtained by using the default in group (IV). Thus each t_i appears in E. Further, $\{a_i, b_i, c_i\} \cap E$ is non-empty. As a result, $\varphi \in E$ is satisfiable.

Let E be an extension of (D, W). We define the truth assignment v_E as follows. For each variable $p \in var(\varphi)$, $v_E(p) = 1$ if $p \in E$, $v_E(p) = 0$ otherwise. Clearly, $v_E(\varphi) = 1$. We now prove the statement of the theorem.

$(\Rightarrow)$. Suppose v is the unique truth assignment satisfying φ. By the above argument, we know E_v is an extension of $(D, \emptyset)$. Suppose E is another extension of $(D, \emptyset)$. By the above discussion we know that $\{t_1, \cdots, t_m, t\} \subseteq E_v \cap E$. Thus E and E_v must be different at some variable $p \in var(\varphi)$. That means v and v_E will be different, contradicts the uniqueness of v. Thus $(D, \emptyset)$ has exactly one extension.

$(\Leftarrow)$. Suppose E is the unique extension of $(D, \emptyset)$. Then v_E is a truth assignment satisfying φ. Suppose v is another truth assignment of φ. Then E_v is an extension of $(D, \emptyset)$. Since v_E and v are different, E and E_v are different by the definition of v_E and E_v. This contradicts the uniqueness of E. Consequently, φ is uniquely satisfiable.

Hence, We prove the theorem. □

4.3. Unary defaults and Ordered Theories

In this section, we analyze the problem of uniqueness of extensions for the case of unary defaults and ordered theories. We show that UEE problem for unary default theories is at least as hard as Unique-SAT. A default d is termed unary if it is of the form

$$d = \frac{p : q}{q}, \text{ or } \frac{p : q \wedge \neg r}{q}, \text{ or } \frac{p : \neg q}{\neg q},$$

where p, q, r are propositional variables.

Theorem 4.3. The problem of determining whether a unary default theory has exact one extension is at least as hard as Unique-SAT.

Proof:
We will employ the reduction defined in Definition 5.3 of [11] to prove the hardness (but we use different notations). Let $\varphi = \alpha_1 \wedge \cdots \wedge \alpha_n$ be a 3CNF formula with $\alpha_i = (a_i \vee b_i \vee c_i)$. Let π be the function that maps each positive literal to itself, and maps each negative literal to a new variable. For each clause α_i we introduce two new propositional variables f_i, g_i. Let f and z be two other new variables. Let D be made up of defaults in the following groups.

(A) For each variable p, the defaults:

$$\frac{:p}{p}, \ \frac{:\neg p}{\neg p};$$

(B) For each variable p, the defaults:

$$\frac{p : \neg\pi(\neg p)}{\neg\pi(\neg p)}, \ \frac{: \pi(\neg p) \wedge \neg p}{\pi(\neg p)};$$

(C) For each clause α_i, the defaults:

$$\frac{\pi(\neg a_i) : g_i \wedge \neg\pi(b_i)}{g_i}, \ \frac{g_i : f_i \wedge \neg\pi(c_i)}{f_i}, \ \frac{f_i : f \wedge \neg z}{f};$$

(D) The single default:

$$\frac{f : z}{z}.$$

Kautz and Selman [11] pointed out that φ is satisfiable if and only if $(D, \emptyset)$ has an extension. In fact, from different satisfying truth assignments of φ we can construct different extensions of $(D, \emptyset)$ and vice versa. Thus, φ is uniquely satisfiable if and only if $(D, \emptyset)$ has exact one extension. □

Let us now consider ordered disjunction-free default theories. Since they always have extensions [11, 6], the UEE problem reduces to checking whether they do not have two or more extensions. As a result, the problem is in coNP. However, the problem remains intractable even for ordered unary default theories.

Definition 4.2. ([6, 11]) Suppose (D, W) is a disjunction-free theory, let Lit be the set of all literals occurring in the theory. Define $\ll$ and $\underline{\ll}$ to be the smallest relation over $Lit \times Lit$ such that

1. $\underline{\ll}$ is reflexive,
2. $\underline{\ll}$ is a superset of $\ll$,
3. $\underline{\ll}$ and $\ll$ are transitive,
4. $\ll$ is transitive through $\underline{\ll}$; that is, for literals $x, y, z \in Lit$:

$$[(x \ll y \wedge y \underline{\ll} z) \vee (x \underline{\ll} y \wedge y \ll z)] \rightarrow x \ll z,$$

5. for every $d \in D$, and every $a \in p(d), b \in c(d)$, and $c \in j(d) - c(d)$:

$$a \underline{\ll} b, \ \neg c \ll b.$$

Then (D, W) is said to be ordered if and only if there is no literal $x \in Lit$ such that $x \ll x$

Theorem 4.4. The problem of determining whether an ordered unary default theories has exactly one extension is coNP-complete.

Proof:
We define a reduction from a 3CNF formula φ containing a negative clause (i.e., every literal in the clause is negative) to an ordered unary theory. Pick an additional new variable t. Let D be made up of defaults in groups (B), (C) of Theorem 4.3 and defaults in the following groups (A'), (E)

(A') For each variable p, the rules:

$$\frac{:p}{p}, \quad \frac{t:\neg p}{\neg p};$$

(E) The single rule:

$$\frac{: t \wedge \neg f}{t}.$$

We first check that $(D, \emptyset)$ is ordered. For $\neg p$, because $\neg p$ does not occur in the prerequisite of any defaults and because p does not occur in the justification of any default except for $\frac{:p}{p}$, we know that there is no other literal L such that $\neg p \ll L$ or $\neg p \underline{\ll} L$. Hence it is impossible that $\neg p \ll \neg p$. For $\neg\pi(\neg p)$, it does not occur as prerequisite in any default and $\pi(\neg p)$ does not occur as justification in any default. Therefore, $\neg\pi(\neg p) \not\ll \neg\pi(\neg p)$. For any other literal L, from the default in (A'),(B), (C),(E), we can see that in each maximal path (with respect to $\ll$ and $\underline{\ll}$) from L, any two literals are different. Hence, $L \not\ll L$. Consequently, $(D, \emptyset)$ is ordered.

Suppose φ unsatisfiable. Clearly,

$$\{p, \neg\pi(\neg p) \mid p \text{ occurs in } \varphi\} \cup \{f_i, g_i \mid \alpha_i \text{ is a negative clause }\} \cup \{f\}$$

is an extension. Suppose $(D, \emptyset)$ has another extension E. If $t \in E$, then $f \notin E$. Then E could determine a satisfying truth assignment v, contradicting the unsatisfiability. Thus, $t \notin E$. It follows that $f \in E$ (otherwise, t could be obtained by applying the default in (E)). Then, there must be a variable p such that $p \notin E$. If $\neg p \notin E$ then $\frac{:p}{p}$ would be applicable. Thus $\neg p \in E$. Then $\neg p$ must be obtained by applying $\frac{t:\neg p}{\neg p}$. Hence $t \in E$. That means $f \notin E$. This would imply the satisfiability of φ. Therefore, $(D, \emptyset)$ has exactly one extension.

Suppose $(D, \emptyset)$ has exactly one extension. Since φ contains a negative clause, it follows that

$$\{p, \neg\pi(\neg p) \mid p \text{ occurs in } \varphi\} \cup \{f_i, g_i \mid \alpha_i \text{ is a negative clause }\} \cup \{f\}$$

is an extension of $(D, \emptyset)$ (from the defaults constructed from a negative clause we get f, thus the default in (E) is not applicable). Suppose φ is satisfiable. Let v be a satisfying truth assignment of φ. From v we can construct an extension which contains t. Then $(D, \emptyset)$ has at least two extensions, a contradiction. Thus φ is unsatisfiable.

□

5. Conclusions

In this paper we mainly analyzed the complexity of the unique extension existence problem. As explained in the Introduction, the problem of checking whether a default theory has a small number of extensions is also of interest. For example, for some fixed natural number $k \geq 1$, one may ask whether a default theory has at most k extensions. We denote this problem by EE(k), and discuss its complexity.

For normal default theories, membership to Π_2^P is easy to prove: the opposite problem of determining whether a normal default theory has more than k extensions is in Σ_2^P: guess $k+1$ permutations of defaults, for each permutation construct an extension; if we get $k+1$ different extensions then return yes. Thus, EE(k) is in Π_2^P for normal default theories.

For semi-normal default theories, the problem might get harder, since we also have to check the existence of extensions. For normal disjunction-free default theories, EE(k) can still be solved in polynomial time. The proof is similar to that in section 4.1; being long but tedious, it is omitted. For semi-normal disjunction-free default theories, EE(k) is still in D^P for any fixed k.

Summarizing, EE(k) for any fixed $k > 1$ does not get harder than UEE with respect to polynomial transformations. On the other hand, EE(k) for fixed $k > 1$ can not get easier than UEE. This can be seen as follows (for simplicity, we take $k = 3$ as example). For any default theory (D, W), pick new variables p, q. Define

$$D' = D \cup \left\{ \frac{:p}{p}, \frac{:\neg p}{\neg p}, \frac{p:q}{q}, \frac{p:\neg q}{\neg q} \right\}.$$

It is easy to see that (D, W) has exactly one extension if and only if (D', W) has exactly three extensions. The case of $k = 3$ easily generalizes to any fixed $k > 1$: all the hardness results for UEE hold for EE(k) as well.

Let us now discuss the significance of the results. One might think of a default theory with a small number of extensions as a rational system. Our results show that the problem of deciding whether a default theory is rational is also very hard. While it is clear that it cannot be any easier than default resoning, an additional source of complexity is that of checking conflict accessibility. As we have seen in the previous section, multiple extensions are due to the presence of defaults that can be applied alone, but not together (conflicting defaults.) Nevertheless, such conflicts may not be accessible, in the sense that the conditions to apply such defaults are never met while actually computing extensions. In particular, a default theory may contain many conflicts, but most of them are never encountered when computing extensions. Im this case, the default theory has only a small number of extensions. Let us consider the following example.

$$D = \left\{ \frac{:p}{p}, \frac{p:r}{r}, \frac{r:\neg p}{\neg p} \right\}.$$

The default theory $(D, \emptyset)$ contains a conflict between the first and the third defaults of D, as their consequences contradict each other. However, the last default is never applicable, hence the default theory has only one extension. Let us consider another example.

$$D' = \left\{ \frac{:p}{p}, \frac{p:r}{r}, \frac{r:\neg p}{\neg p}, \frac{:t}{t}, \frac{t:r}{r} \right\}.$$

The default theory $(D', \emptyset)$ also contains a conflict. Since both the two conflicting defaults $\frac{:p}{p}$ and $\frac{r:\neg p}{\neg p}$ are accessible, $(D', \emptyset)$ has two extensions. Testing whether a conflict is accessible is what makes the UEE problem hard. The restricted cases in which the test of accessibility of conflicts is easy, the complexity of the UEE problem decreases. For example, the problem of determining whether $\Gamma(W)$ is an extension of a normal default theory (D, W) is $P^{NP[\log n]}$-complete, since conflict accessibility is easy in this case. In the normal disjuction-free case, the test of accessibility can be done efficiently, hence the UEE problem is solvable in polynomial time.

Some problems that have been left open by this work. The UEE problem for semi-normal default theories is in D^P, but it is open whether it is D^P-complete or not. We have shown that the problem is at least as hard as the unique satisfiability problem which, is in D^P and has been shown to be coNP-hard. Even for very simple (e.g., unary) theories, the problem is still at least as hard as Unique-SAT.

The problem simplifies when the default theory is very simple: for ordered theories, since the existence of extensions is guaranteed, the UEE problem is coNP-complete. For ordered unary theories the problem is still coNP-complete.

References

[1] Baumgartner, R., Gottlob, G.: On the Complexity of Model Checking for Propositional Default Logics: New Results and Tractable Cases, *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI'99)*, 1999.

[2] Cadoli, M., Donini, F. M.: A Survey on Knowledge Compilation, *AI Communications—The European Journal on Artificial Intelligence*, **10**, 1997, 137–150.

[3] Cadoli, M., Donini, F. M., Liberatore, P., Schaerf, M.: *Preprocessing of intractable problems*, Technical Report DIS 24-97, Dipartimento di Informatica e Sistemistica, Università di Roma "La Sapienza", November 1997, To appear in *Information and Computation*.

[4] Cadoli, M., Donini, F. M., Schaerf, M.: Is Intractability of Non-monotonic Reasoning a Real Drawback?, *Artificial Intelligence*, **88**(1–2), 1996, 215–251.

[5] Cadoli, M., Schaerf, M.: A Survey of Complexity results for Non-monotonic Logics, *Journal of Logic Programming*, **17**, 1993, 127–160.

[6] Etherington, D. V.: *Reasoning with incomplete information*, Morgan Kaufmann, Los Altos, Los Altos, CA, 1987.

[7] Gottlob, G.: Complexity Results for Nonmonotonic Logics, *Journal of Logic and Computation*, **2**, 1992, 397–425.

[8] Gottlob, G.: The Complexity of Default Reasoning under the Stationary Fixed Point Semantics, *Information and Computation*, **121**, 1995, 81–92.

[9] Gottlob, G.: NP trees and Carnap's modal logic, *Journal of the ACM*, **42**(2), 1995, 421–457.

[10] Johnson, D. S.: A Catalog of Complexity Classes, in: *Handbook of Theoretical Computer Science* (J. van Leeuwen, Ed.), vol. A, chapter 2, Elsevier Science Publishers (North-Holland), Amsterdam, 1990, 67–161.

[11] Kautz, H. A., Selman, B.: Hard Problems for Simple Default Logics, *Artificial Intelligence*, **49**, 1991, 243–279.

[12] Kleine Büning, H., Lettmann, T.: *Propositional Logic: Deduction and Algorithm*, Cambridge University Press, 1999.

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

[14] Lang, J., Marquis, P.: In search of the right extension, *Proceedings of the Eighth International Conference on Principles of Knowledge Representation and Reasoning (KR 2002)*, 2002.

[15] Liberatore, P.: Monotonic Reductions, Representative Equivalence, and Compilation of Intractable Problems, *Journal of the ACM*, **48**(6), 2001, 1091–1125.

[16] Liberatore, P., Schaerf, M.: The complexity of model checking for propositional default logics, *Proceedings of the Thirteenth European Conference on Artificial Intelligence (ECAI'98)*, 1998.

[17] Marek, W., Truszczyński, M.: *Nonmonotonic Logics*, Springer, 1993.

[18] Przymusinka, H., Przymusinski, T.: Stationary Default Extensions, *Fundamenta Informaticae*, **21**, 1994, 67–87.

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

[20] Rosati, R.: Model Checking for Nonmonotonic Logics, *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI'99)*, 1999.

[21] Stillman, J.: It's Not My Default: The Complexity of Membership Problems in Restricted Propositional Default Logics, *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI'90)*, 1990.

[22] Wadner, K.: Bounded Query Classes, *SIAM Journal on Computing*, **19**, 1990, 833–846.

[23] Zhao, X., Ding, D.: Complexity Results for 2CNF Default Theories, *Fundamenta Informaticae*, **45**, 2001, 393–404.

[24] Zhao, X., Ding, D., Kleine Büning, H.: Complexity Results for Restricted Credulous Default Reasoning, *AI Communications—The European Journal on Artificial Intelligence*, **13**, 2000, 249–257.