

Prioritizing Default Logic: Abridged Report

Gerhard Brewka¹ and Thomas Eiter²

¹ Institut für Informatik, Abteilung Intelligente Systeme, Universität Leipzig
Augustusplatz 9/10, D-04109 Leipzig, Germany
Email: brewka@informatik.uni-leipzig.de

² Institut für Informatik, Justus-Liebig-Universität Gießen
Arndtstraße 2, D-35392 Gießen, Germany
Email: eiter@informatik.uni-giessen.de

Abstract. A number of prioritized variants of Reiter’s default logic have been described in the literature. In this paper, we introduce two natural principles for preference handling and show that all existing approaches fail to satisfy them. We develop a new approach which does not suffer from these shortcomings. We start with the simplest case, supernormal default theories, where preferences are handled in a straightforward manner. The generalization to prerequisite-free default theories is based on an additional fixed point condition for extensions. The full generalization to arbitrary default theories uses a reduction of default theories to prerequisite-free theories. The reduction can be viewed as dual to the Gelfond/Lifschitz reduction used in logic programming for the definition of answer sets. We finally show how preference information can be represented in the logical language.

1 Introduction

In nonmonotonic reasoning conflicts among defaults are ubiquitous. For instance, more specific rules may be in conflict with more general ones, a problem which has been studied intensively in the context of inheritance networks [21, 26, 27]. When defaults are used for representing design goals in configuration tasks conflicts naturally arise. The same is true in model based diagnosis where defaults are used to represent the assumption that components typically are ok. In legal reasoning conflicts among rules are very common [22] and keep many lawyers busy (and rich).

The standard nonmonotonic formalisms handle such conflicts by generating multiple belief sets. In default logic [23] and autoepistemic logic [18] these sets are called extensions or expansions, respectively. In circumscription [17] the belief sets correspond to different classes of preferred models.

Usually, not all of the belief sets are plausible. We often tend to prefer some of the conflicting rules and are interested in the belief sets generated by the preferred rules only. One way to achieve this is to rerepresent the defaults in such a way that the unwanted belief sets are not generated, for instance by adding new consistency conditions to a default. This approach has the advantage that the logical machinery of the underlying nonmonotonic logic does not have to be changed. On the other hand, rerepresenting

the defaults that way is a very clumsy business. The resulting new defaults tend to be rather complex. Moreover, the addition of new information to the knowledge base may lead to further rerepresentations. In other words, elaboration tolerance is violated.

For this reason we prefer an approach where preferences among defaults can be represented via an explicit preference relation and where the logical machinery is extended accordingly. Indeed, for all major nonmonotonic formalisms, such prioritized versions have been proposed in the past. Among them are prioritized circumscription [12], hierarchic autoepistemic logic [15], prioritized theory revision [2, 19], prioritized logic programming [25, 28], or prioritized abduction [10].

Also several prioritized versions of Reiter's default logic, the logic we are dealing with in this paper, have been described in the literature [16, 4, 1, 9], as well as of defeasible logics beyond default logic [20, 13]. However, as we will show in Section 3, these approaches are not fully satisfactory. It turns out that some of them implicitly recast Reiter's default logic to a logic of graded beliefs, while others do overly enforce the application of rules with high priority, which leads to counterintuitive behavior.

Our approach takes a different perspective, which is dominated by the following two main ideas. The first is that the application of a default rule means to jump to a conclusion, and this conclusion is yet another assumption which has to be used globally in the program for the issue of deciding whether a rule is applicable or not. The second is that the rules must be applied in an order compatible with the priority information. We take this to mean that a rule is applied *unless it is defeated via its assumptions by rules of higher priorities*. This view is new and avoids the unpleasant behavior which is present with the other approaches. Our formalization of these ideas involves a dual of the standard Gelfond-Lifschitz reduction and a certain operator used to check satisfaction of priorities. In order to base our approach on firmer ground we set forth some abstract principles that, as we believe, any formalization of prioritized default logic should satisfy. We demonstrate that our approach satisfies these principles, while other approaches violate them.

The remainder of this paper is organized as follows. The next section recalls the basic definitions of default logic and introduces the notion of a prioritized default theory. Section 3 introduces two basic principles for preference handling, reviews some approaches to prioritized default logic and demonstrates that they fail to satisfy the principles. In Section 4, we then present our approach, by introducing the concept of preferred extensions. We will introduce this notion in a stepwise manner, starting with the simplest case, namely prerequisite-free normal defaults, which are also called supernormal defaults. We then extend our definition to prerequisite-free default theories, showing that an additional fixed point condition is needed to get the definition of preferred extensions right. Finally, we handle arbitrary default theories by reducing them to the prerequisite-free case. The reduction can be viewed as dual to the famous Gelfond/Lifschitz reduction for extended logic programs. In Section 5 we show how preference information can be expressed in the logical language. This makes it possible to reason not only with, but also about preferences among rules. Section 6 discusses related work, and concludes the paper by considering possible extensions and outlining further work.

The work reported here generalizes the approach presented in [7], which covers the fragment of default logic equivalent to extended logic programs, and extends the results.

2 Prioritized Default Theories

We first recall the basic definitions underlying Reiter's default logic. A *default theory* is a pair $\Delta = (D, W)$ of a theory W containing first-order sentences and a set of defaults D . Each default is of the form $a : b_1, \dots, b_n / c$, $n \geq 1$, where a , b_i , and c are first-order formulas. The intuitive meaning of the default is: if a is derived and the b_i are separately consistent with what is derived, then infer c . Formula a is called the *prerequisite*, each b_i a *justification*, and c the *consequent* of the default. For a default d we use $pre(d)$, $just(d)$, and $cons(d)$ to denote the prerequisite, the set of justifications, and the consequent of d , respectively; $\neg just(d)$ denotes $\{\neg a \mid a \in just(d)\}$. As usual, we assume that W and D are in skolemized form and that open defaults, i.e., defaults with free variables, represent the sets of their ground instances over the Herbrand universe [23]; a default theory with open defaults is closed by replacing open defaults with their ground instances. In what follows, we implicitly assume that default theories are closed before extensions etc are considered.

A (closed) default theory generates extensions which represent acceptable sets of beliefs which a reasoner might adopt based on the given default theory (D, W) . Extensions are defined in [23] as fixed points of an operator T_Δ . T_Δ maps an arbitrary set of formulas S to the smallest deductively closed set S' that contains W and satisfies the condition: if $a : b_1, \dots, b_n / c \in D$, $a \in S'$ and $\neg b_i \notin S$ then $c \in S'$. Intuitively, an extension is a set of beliefs containing W such that

1. as many defaults as consistently possible have been applied, and
2. only formulas possessing a noncircular derivation from W using defaults in D are contained.

A default theory Δ may have zero, one or multiple extensions. Default theories possessing at least one extension will be called coherent. We say a default $a : b_1, \dots, b_n / c$ is *defeated* by a set of formulas S , iff $\neg b_i \in S$ for some $i \in \{1, \dots, n\}$.

A default $d = a : b / c$ is called *normal*, if b is logically equivalent to c ; it is called *prerequisite-free*, if a is a logical truth, which is denoted by \top . Defaults which are both prerequisite-free and normal are called *supernormal*. A default theory is called *normal (prerequisite-free, supernormal)*, if all of its defaults are normal (prerequisite-free, supernormal), respectively.

A default d is called *generating* in a set of formulas S , if $pre(d) \in S$ and $\neg just(d) \cap S = \emptyset$; denote by $GD(E, D)$ the set of all defaults from D which are generating in E . It is well-known [23] that every extension of a default theory $\Delta = (D, W)$ is characterized through $GD(D, E)$, i.e.,

$$E = Th(W \cup cons(GD(D, E))) \quad (1)$$

where $cons(D') = \{cons(d) \mid d \in D'\}$ for any set D' . Moreover, if Δ is prerequisite-free, then every E which satisfies (1) is an extension (cf. [16]).

We now introduce the notion of a prioritized default theory. Basically, we extend default theories with a strict partial order $<$ (i.e., $d \not< d$ and $d < d'$, $d' < d''$ implies $d < d''$) on the default rules. A default d will be considered preferred over default d' , whenever $d < d'$ holds.

Definition 1. *A prioritized default theory is a triple $\Delta = (D, W, <)$ where (D, W) is a default theory and $<$ is a strict partial order on D .*

Partially ordered default theories have the advantage that the preference ordering among certain defaults can be left unspecified. This is important because in many cases there is no natural way of assigning preferences. However, the case of arbitrary partial orders can be reduced to particular refinements, namely well-orderings, in a canonical way. Recall that a partial order is a well-ordering, iff every subset of the elements has the least element; observe that any well-ordering is a total ordering.

Definition 2. *A fully prioritized default theory is a prioritized default theory $\Delta = (D, W, <)$ where $<$ is a well-ordering.*

Conclusions of prioritized default theories are defined in terms of preferred extensions, which are a subset of the classical extensions of Δ , i.e., the extensions of (D, W) according to [23]. The definition of preferred extension for fully prioritized default theories will be given in the next section. The general case of arbitrary prioritized default theories can then be reduced to this case as follows.

Definition 3. *Let $\Delta = (D, W, <)$ be a closed prioritized default theory. E is a prioritized extension of Δ iff E is a prioritized extension of a fully prioritized default theory $\Delta' = (D, W, <')$ such that $d < d'$ implies $d <' d'$.*

The preferred extensions of an open prioritized default theory Δ are the preferred extensions of Δ^* obtained by closing Δ . The partial order $<$ is inherited from D to the ground set of instances D^* in the obvious way. We assume here that no conflict arises, i.e., $d < d$ does not result for any $d \in D^*$; otherwise, no preferred extensions are defined.

In the remainder of this paper, we will restrict our discussion to fully prioritized default theories. Unless stated otherwise, all default theories are tacitly assumed to be closed.

3 Problems with Existing Approaches

Different prioritized versions of default logic have been proposed in the literature, e.g. [16, 4, 1, 24, 9]. We will show that all of them fail to satisfy natural principles for preference handling in default logic.

3.1 Principles for priorities

The first principle can be viewed as a meaning postulate for the term “preference” and states what we consider a minimal requirement for preference handling in rule based systems:

Principle I. Let B_1 and B_2 be two extensions of a prioritized default theory Δ generated by the defaults $R \cup \{d_1\}$ and $R \cup \{d_2\}$, where $d_1, d_2 \notin R$, respectively. If d_1 is preferred over d_2 , then B_2 is not a preferred extension of T .

We find it hard to see how the use of the term “preference among rules” could be justified in cases where Principle I is violated.

The second principle is related to relevance. It tries to capture the idea that the decision whether to believe a formula p or not should depend on the priorities of defaults contributing to the derivation of p only, not on the priorities of defaults which become applicable when p is believed:

Principle II. Let E be a preferred extension of a prioritized default theory $\Delta = (D, W, <)$, d a (closed) default such that the prerequisite of d is not in E . Then E is a preferred extension of $\Delta' = (D \cup \{d\}, W, <')$ whenever $<'$ agrees with $<$ on priorities among defaults in D .

Thus, adding a rule which is not applicable in a preferred belief set can never render this belief set non-preferred unless new preference information changes preferences among some of the old rules (e.g. via transitivity). In other words, a belief set is not blamed for not applying rules which are not applicable.

We will see that each of the existing treatments of preferences for default logic, described in [16, 4, 1, 24], violates one of these principles.

3.2 Control of Reiter’s quasi-inductive definition

The first group of proposals [16, 4, 1] uses preferences to control the quasi-inductive definition of extensions [23]: in each step of the generation of extensions the defaults with highest priority whose prerequisites have already been derived are applied. Now what is wrong with this idea? The answer is: the preferred extensions do not take seriously what they believe. It may be the case that a less preferred default is applied although the prerequisite of a conflicting, more preferred default is believed in a preferred extension. As we will see, this can lead to situations where Principle I is violated.

The mentioned approaches differ in technical detail. We do not want to present the exact definitions here. Instead, we will illustrate the difficulties using an example for which all these approaches obtain the same result.

Example 1. Assume we are given the following default theory:

- (1) $a : b/b$
- (2) $\top : \neg b/\neg b$
- (3) $\top : a/a$

Assume further that (1) is preferred over (2) and (2) over (3). This default theory has two classical extensions, namely $E_1 = Th(\{a, b\})$, which is generated by rules (1) and (3), and $E_2 = Th(\{a, \neg b\})$, which is generated by rules (2) and (3). The single preferred extension in the approaches mentioned above is E_2 . The reason is that the prerequisite of (2) is derived before the prerequisite of (1) in the construction of the extension. The approaches thus violate Principle I. ■

The selection of E_2 in the previous example was already observed in [4]. In that paper, the first author tried to defend his approach arguing that there is only weak evidence for the literal a in our example. We revise our view, however, and do not support this argument any longer. After all, default logic is not a logic of graded belief where degrees of evidence should play a role. Default logic models acceptance of belief based on defeasible arguments. Since a is an accepted belief, we believe rule (1) should be applied and E_1 should be the preferred extension in the example.

3.3 Rintanen's approach

An entirely different approach was proposed in [24]. Rintanen uses a total order on (normal) defaults to induce a lexicographic order on extensions.

Call a normal default rule $r = a:b/b$ *applied* in a set of formulas E (denoted $r \in \text{appl}(E)$), if a and b are in E . An extension E is then preferred over extension E' , if and only if there is a default $r \in \text{appl}(E) \setminus \text{appl}(E')$ satisfying the following condition: if r' is preferred over r and $r' \in \text{appl}(E')$, then $r' \in \text{appl}(E)$.

Unfortunately, also this approach leads to counterintuitive results and to a violation of our principles.

Example 2. Consider the following default theory:

- (1) $a : b/b$
- (2) $\top : \neg a/\neg a$
- (3) $\top : a/a$

Again (1) is preferred over (2), and (2) over (3). The default theory has two classical extensions, namely $E_1 = Th(\{\neg a\})$ and $E_2 = Th(\{a, b\})$. Intuitively, since the decision whether to believe a or not depends on (2) and (3) only, and since (2) is preferred over (3), we would expect to conclude $\neg a$, in other words, to prefer E_1 .

However, the approach of Rintanen prefers E_2 . The reason is that in E_2 default (1) is applied. Belief in a is thus accepted on the grounds that this allows us to apply a default of high priority. This is far from being plausible and amounts to wishful thinking. It is also easy to see that Principle II is violated: E_1 clearly is the single preferred extension of rules (2) and (3) in Rintanen's approach. Adding rule (1) which is not applicable in E_1 makes E_1 a non-preferred extension. ■

Since all these approaches suffer from drawbacks, we develop our new approach in the following section.

4 Preferred Extensions

In this section we introduce our new notion of preferred extensions for fully prioritized default theories. As mentioned before, arbitrary prioritized default theories can be reduced to that case in a canonical manner.

We will consider the simplest case first, namely supernormal default theories. We then proceed to prerequisite-free default theories and, finally, to arbitrary default theories.

4.1 Supernormal default theories

Preference handling in prioritized supernormal default theories is rather easy. The obvious idea is to check the applicability of defaults in the order of preference. We first introduce an operator C which, given a fully prioritized prerequisite-free default theory Δ (which is not necessarily supernormal), produces tentative conclusions of Δ .

Call a default d *active* in a set of formulas S , if $pre(d) \in S$, $\neg just(d) \cap S = \emptyset$ and $cons(d) \notin S$ all hold. Intuitively, a default is active in S if it is applicable wrt. S and has not yet been applied.

Definition 4. Let $\Delta = (D, W, <)$ be a fully prioritized prerequisite-free default theory. The operator C is defined as follows: $C(\Delta) = \bigcup_{\alpha \geq 0} E_\alpha$, where $E_0 = Th(W)$, and for every ordinal $\alpha > 0$,

$$E_\alpha = \begin{cases} \underline{E}_\alpha, & \text{if no default from } D \text{ is active in } \underline{E}_\alpha; \\ Th(\underline{E}_\alpha \cup \{cons(d)\}) & \text{otherwise, where} \\ & d = \min_{<} \{d' \in D \mid d' \text{ is active in } \underline{E}_\alpha\}, \end{cases}$$

where $\underline{E}_\alpha = \bigcup_{\beta < \alpha} E_\beta$. (Note that for each successor ordinal $\alpha = \beta + 1$, $\underline{E}_\alpha = E_\beta$.)

In the case of supernormal default theories, the operator C always produces an extension in the sense of Reiter and thus can directly be used to define preferred extensions:

Definition 5. Let $\Delta = (D, W, <)$ be a fully prioritized supernormal default theory. E is the preferred extension of Δ if and only if $E = C(\Delta)$.

It is obvious that there is always exactly one preferred extension. Note that the definition of this extension is fully constructive. It extends the notion of preferred subtheories as developed in [3] to the infinite case.

4.2 Prerequisite-free default theories

Can we simply extend the definition for supernormal defaults to this case? The answer is obviously no. It may be the case that defaults are applied during the construction which are defeated later through the application of defaults of lower priority.

So, can we simply say: if the construction gives us an extension, then that extension is preferred? Unfortunately, the answer is again no.

Example 3. Consider the following default theory:

$$\begin{array}{ll} (1) \top : \neg b/a & (3) \top : a/a \\ (2) \top : \neg a/\neg a & (4) \top : b/b \end{array}$$

Assume $(1) < (2) < (3) < (4)$. Applying operator C to this default theory yields $E = Th(\{a, b\})$. As easily seen, this is a classical extension. Nonetheless, one would certainly not say that this extension preserves priorities. What went wrong? Default (2) is defeated in E by applying a default which is less preferred than (2), namely default (3). In the construction of $C(\Delta)$ this remains unnoticed, since rule (1), although defeated in E , blocks the applicability of (2). In other words, without a special treatment of such cases, a rule (e.g. (3)) may inherit a high preference from a rule with the same consequent (namely (1)), even if that latter rule is not applicable in the extension. ■

To avoid this, we have to impose an additional condition on an extension: in the construction of the tentative conclusions, we have to discard each rule whose consequent is in E , but which is defeated in E . Since we have to take E as the result of the construction into account, this amounts to adding a fixed point condition. What we will do is check whether we arrive at the same set of formulas after eliminating rules which are defeated in E and whose head is in E .

Definition 6. Let $\Delta = (D, W, <)$ be a fully prioritized prerequisite-free default theory. Then, a set E of formulas is a prioritized extension of Δ , if and only if

$$E = C(\Delta^{*E}),$$

where Δ^{*E} is obtained from Δ by deleting all defaults whose consequents are in E and which are defeated in E .

This definition is coherent with the intuition that preferred extensions are distinguished classical extensions. Moreover, as in the case of supernormal theories, the preferred extension (if it exists) is unique.

Proposition 1. Let $\Delta = (D, W, <)$ be a fully prioritized prerequisite-free default theory. Then, every preferred extension E of Δ is a classical extension, and Δ has at most one preferred extension.

Proof. To show the first part, assume that E is a preferred extension of Δ . We show that $E = Th(W \cup cons(GD(D, E)))$ (*) holds; since all defaults are prerequisite-free, this implies that E is a classical extension of Δ (cf. paragraph after Equation (1)).

Since $E = C(\Delta^{*E})$, no default from $D \setminus GD(D, E)$ is applied in the construction of E ; hence, $E \subseteq Th(W \cup cons(GD(D, E)))$ follows. On the other hand, since each E_α is included in E and E does not defeat d , for every $d \in GD(D, E)$, it follows $cons(d) \in C(\Delta^{*E})$; hence, $Th(W \cup cons(GD(D, E))) \subseteq E$ follows. This implies (*), and proves that E is a classical extension of Δ .

For the second part, assume that different preferred extensions $E \neq E'$ exist. We derive a contradiction. Let d be the least default in D such that either (i) $d \in GD(D, E)$ and $cons(d) \notin E'$, or (ii) $d \in GD(D, E')$ and $cons(d) \notin E$. Since $E \neq E'$, d must exist. Consider first the case (i). It follows that $d \in D^{*E'}$; for, otherwise $cons(d) \in E'$ holds, which is a contradiction. The default d must be defeated by E' ; from the definition of $C(\Delta^{*E})$, it follows that d is defeated by $Th(W \cup cons(K))$ for $K = \{d' \in GD(D, E') \mid d' < d\}$. From the minimality of d , it follows that for every $d' \in K$ it holds that $cons(d') \in E$. Hence, $Th(W \cup cons(K)) \subseteq E$, which means E defeats d . This is a contradiction to $d \in GD(D, E)$, however. The case (ii), i.e., $d \in GD(D, E')$ and $cons(d) \notin E$ is analogous. This proves the result. ■

4.3 General default theories

We will now reduce the general case to the prerequisite-free case. The basic idea is the following: in order to check whether an extension E of a fully prioritized default theory Δ is preferred, we evaluate the prerequisites of the default rules according to the

extension E . Evaluating prerequisites means (1) eliminating prerequisites which are contained in the extension E from the corresponding rules, and (2) eliminating rules whose prerequisites are not contained in E . Observe that this operation can be viewed as a dual of the standard Gelfond/Lifschitz reduction from logic programming [11], in which the justifications rather than the prerequisites are used to eliminate and simplify rules.

Finally, we check whether the resulting prerequisite-free theory Δ_E has E as its preferred extension.

Definition 7. Let $\Delta = (D, W, <)$ be a fully prioritized default theory and E a set of formulas. The default theory $\Delta_E = (D_E, W, <_E)$ is obtained from Δ as follows: D_E results from D by

1. eliminating every default $d \in D$ such that $\text{pre}(d) \not\subseteq E$, and
2. replacing $\text{pre}(d)$ by \top in all remaining defaults;

$<_E$ is inherited from $<$ as follows: for any rules d and d' in D_E , $d <_E d'$ holds if and only if $d1 < d1'$ holds for the $<$ -least rules $d1$ and $d1'$ in D which give rise to d and d' (i.e., $d1_E = d$ and $d1'_E = d'$), respectively.

The resulting default theory is clearly prerequisite-free. We thus can define preferred extensions for general default theories as follows:

Definition 8. Let $\Delta = (D, W, <)$ be a fully prioritized default theory. Then, E is a prioritized extension of Δ , if (i) E is a classical extension of Δ , and (ii) E is a prioritized extension of Δ_E .

Let us show that the problematic examples discussed in Section 3 are handled correctly in our approach:

Example 4. Consider the default theory Δ :

- (1) $a : b/b$
- (2) $\top : \neg b/\neg b$
- (3) $\top : a/a$

Again we assume that (1) is preferred over (2) and (2) over (3). This default theory has two classical extensions, namely $E_1 = Th(\{a, b\})$ and $E_2 = Th(\{a, \neg b\})$. Δ_{E_1} consists of the rules

- (1) $\top : b/b$
- (2) $\top : \neg b/\neg b$
- (3) $\top : a/a$

It is not difficult to see that $C(\Delta_{E_1}) = E_1$. Since there are no rules whose head is in E_1 but which are defeated in E_1 , E_1 is a preferred extension. On the contrary, E_2 is not preferred. Note that the dual E_2 -reduct of Δ is the same as the dual E_1 -reduct, and that $C(\cdot)$ applied to the reduct does not reproduce E_2 . ■

Example 5. Consider the following default theory Δ' :

- (1) $a : b/b$
- (2) $\top : \neg a/\neg a$
- (3) $\top : a/a$

Again (1) is preferred over (2), and (2) over (3). The default theory has two Reiter extensions, namely $E_1 = Th(\{-a\})$ and $E_2 = Th(\{a, b\})$. We argued in Section 3 that E_1 should be preferred. Consider Δ'_{E_1} :

- (2) $\top : \neg a/\neg a$
- (3) $\top : a/a$

Clearly, $C(\Delta'_{E_1}) = E_1$, and since again there are no rules defeated in E_1 whose head is in E_1 , we have that E_1 is preferred.

E_2 is not preferred since $C(\Delta'_{E_2}) = Th(\{b, \neg a\})$ which differs from E_2 . ■

The following proposition tells us that all rules which are not generating in a preferred extension must be defeated by some appropriate generating rules, which must have higher priority.

Proposition 2. *Let $\Delta = (W, D, <)$ be a fully prioritized ground default theory, and let E be a classical extension of Δ . Then, E is a preferred extension of Δ , if and only if for each default $d \in D$ such that $pre(d) \in E$ and $cons(d) \notin E$, there exists a set of defaults $K_d \subseteq \{d' \in GD(D, E) \mid d' < d\}$ such that d is defeated in $Th(W \cup cons(K_d))$. ■*

Proof. (\Leftarrow) Suppose that for every default d such that $pre(d) \in E$ and $cons(d) \notin E$ a set $K_d \subseteq \{d' \in GD(D, E) \mid d' < d\}$ exists such that $Th(W \cup cons(K_d))$ defeats d . By (transfinite) induction on the sets E_α , $\alpha \geq 1$, we show that the least active default d_E from D_E^{*E} in \underline{E}_α , provided one exists, stems from some $d \in GD(D, E)$ and $cons(d) \in E_\alpha$ holds.

For $\alpha = 1$, the statement holds. Indeed, the least rule d_E of D_E^{*E} is active. Let d be the least parent of d_E in D , i.e., $d = \min_{<} \{d' \mid d'_E = d_E\}$. Assuming that $d \in D \setminus GD(D, E)$, we obtain $K_d = \emptyset$, and hence d is defeated by $E_0 = Th(W)$. This contradicts that d_E is active in $\underline{E}_1 = E_0$, however. Thus, $d \in GD(D, E)$ holds, and $cons(d) \in E_\alpha$ follows.

Let then $\alpha > 1$ and assume the statement holds for all $1 \leq \beta < \alpha$. Suppose the least default d_E from D_E^{*E} active in \underline{E}_α exists, and that its least parent in D is not in $GD(D, E)$. The induction hypothesis implies that for each $d' \in GD(D, E)$ such that $d' < d$ it holds that $cons(d') \in E_\alpha$. Hence, $Th(W \cup cons(K_d)) \subseteq \underline{E}_\alpha$, which implies that d is defeated by \underline{E}_α . This contradicts that d_E is active. Thus, if d_E exists, then $d \in GD(D, E)$ holds; clearly, $cons(d) \in E_\alpha$. This concludes the induction, from which $C(\Delta_E^{*E}) = Th(W \cup cons(GD(D, E)))$ follows. By Equation (1), it follows $E = C(\Delta_E^{*E})$, which means that E is a preferred extension.

(\Rightarrow) Suppose E is a preferred extension, but some $d \in D$ such that $pre(d) \in E$ and $cons(d) \notin E$ is not defeated by any $Th(W \cup cons(K))$, where $K \subseteq \{d' \in GD(D, E) \mid d' < d\}$. Let d be the least such rule in D . Since E is a preferred extension, for every $d' \in GD(D, E)$ we have $cons(d') \in C(\Delta_E^{*E})$. By the minimality of d , it follows that d_E becomes the least active rule in D_E^{*E} at some step α , and up to

this point, only consequents of active reducts d'_E of $d' \in K$ have been added, i.e., $\underline{E}_\alpha = Th(W \cup cons(K))$ holds. Since \underline{E}_α does not defeat d_E , the rule is applied, which implies $C(\Delta_E^*) \neq E$. Consequently, E is not a preferred extension, which is a contradiction. ■

Exploiting this proposition, we can establish that the principles for a prioritization approach from above are both satisfied by our approach.

Proposition 3. *The approach to preferred extensions satisfies both Principles I and II as described in Section 3.*

Proof. *Principle I.* Let $\Delta = (W, D, <)$ be a prioritized default theory, and let E, E' be classical extensions of Δ such that $GD(D, E) = R \cup \{d\}$ and $GD(D, E') = R \cup \{d'\}$, where $d_1, d_2 \notin R$ and $d < d'$. We have to show that E' is not a preferred extension of Δ .

Towards a contradiction, suppose E' is a preferred extension. Let $\Delta' = (W, D, <')$ be a full prioritization of Δ . Since $d \in GD(D, E)$, it holds that $pre(d) \in Th(W \cup cons(R))$; hence, d survives the dual GL-reduction wrt E' , and d, d' give rise to defaults $d_{E'}, d'_{E'} \in D_{E'}$, respectively.

Since $d \notin GD(D, E')$, it follows that $pre(d) \in E'$ but $cons(d) \notin E'$. Hence, by Proposition 2, it follows that d is defeated by $Th(W \cup cons(K))$ for some $K \subseteq \{d'' \in GD(D, E') \mid d'' < d\}$. It follows that $K \subseteq R$ holds. Since d is defeated by $Th(W \cup cons(K))$ and $Th(W \cup cons(K)) \subseteq E$, it follows that d is defeated by E . This contradicts $d \in GD(D, E)$; satisfaction of Principle I follows.

Principle II. Let E be a preferred extension of $\Delta = (W, D, <)$, and let d be a (closed) default such that $pre(d) \notin E$. We have to show that E is a preferred extension of $\Delta' = (D \cup \{d\}, W, <')$ where $<'$ is compatible with $<$.

Consider the dual reduct of Δ' wrt E , i.e., $\Delta'_E = ((D \cup \{d\})_E, W, <'_E)$. Then, the default d is eliminated in the dual reduct, and we have $\Delta'_E = \Delta_E$. Since E is a preferred extension of Δ_E , it follows immediately that E is a preferred extension of Δ' . Thus, Principle II is satisfied. (Remark: the proof can be easily adapted for an open default d , if closing d does not lead to inconsistency of $<'$.) ■

Thus, our approach satisfies these general benchmarks for a prioritization logic.

On the other hand, a less desirable property of the approach is that in some cases no preferred extension may exist. This is what happens in Example 3; it is easily checked that neither of the two classical extensions $E = Th(\{a, b\})$ and $E' = Th(\{\neg a, b\})$ is a preferred extension.

Another example shows that normal prioritized default theories may have no preferred extension. Thus, the property that supernormal default theories always have extensions is lost if prerequisites are allowed in the defaults.

Example 6. Consider the following defaults:

- (1) $a : \neg b / \neg b$
- (2) $\top : b / b$
- (3) $b : a / a$

This default theory has the unique classical extension $E = Th(\{a, b\})$. However, assuming (1) $<$ (2) $<$ (3), E is not preferred, since preference of (1) requires to conclude $\neg b$. ■

Intuitively, if no preferred extension exists, then the priorities as specified by the user are incompatible with the way in which defaults must be executed to generate an extension. In the preceding example, this is clearly the case. There are different possibilities to react to such an inconsistency, and some of them have been discussed in [7]. There are two main directions for handling such inconsistencies. One direction is to stop on occurrence of such an inconsistency and notify the user that there is an inconsistency in the priorities. The other would be trying to overcome this inconsistency, by reconciling the priority information and the logical entrenchment of default application by relaxing or modifying the priority information in a way such that preferred extensions become possible.

We believe that in general, the first direction is preferable to the second one since the user becomes explicitly aware that there is something wrong with his preferences, which cannot be satisfied. However, we could require that an approach to priorities should be consistent in the sense that if classical extensions exist, then some of them should always be selected by the prioritization method. In this case, a relaxation of our preferred extension approach would be desirable, which selects the preferred extensions if some exist and some classical extensions, according to some rationale, if no preferred extensions exist.

There are different possibilities for generalizing the preferred extensions to such “weakly” preferred extensions. One such possibility is to allow a minimal reordering of the defaults in D , i.e., E becomes a preferred extension after switching as few neighbored defaults in $<$ as possible, cf. [7]. Another approach would be to remove preferences between defaults, e.g., to relax the ordering $<$ as little as possible such that preferred extensions exist. We do not pursue these possibilities any further here. However, we observe some limitations of such weakly preferred extensions.

We call a function χ which selects a subset $\chi(\Delta)$ from the classical extensions of a prioritized default theory D a *consistent preference relaxation* (CPR) of preferred extensions, if $\chi(\Delta)$ selects all and only preferred extensions if preferred extensions exist, and selects some (arbitrary) classical extensions provided some classical extension exists. Then, the following holds.

Proposition 4. *Every consistent preference relaxation χ of preferred extensions must violate both Principle I and Principle II.*

Proof. (Sketch) To show that no CPR χ can satisfy Principle I in general, we consider a prioritized default theory $\Delta = (W, D, <)$ such that Δ has classical extensions but no preferred extensions, and such that χ can not select any of the classical extensions without violating Principle I. Define

$$\begin{aligned} W &= \{b_i \rightarrow (b_{i+1} \wedge c), a_i \leftrightarrow \neg b_i \mid i \geq 0\}, \\ D &= \{\top : a_i/a_i, \top : b_i/b_i \mid i \geq 0\} \cup \{\top : \neg c/\perp\}, \end{aligned}$$

where a_i, b_i and c are propositional atoms, and let $<$ be the well-ordering such that

$$\left. \begin{array}{l} \top : a_i/a_i < \top : b_i/b_i < \top : a_{i+1}/a_{i+1}, \\ \top : a_i/a_i < \top : \neg c/\perp, \top : b_i/b_i < \top : \neg c/\perp \end{array} \right\} i \geq 0,$$

where \perp denotes falsity. It can be seen that the classical extensions of Δ are of the form

$$E^i = Th(W \cup cons(\{\top : a_k/a_k, \top : b_j/b_j \mid 0 \leq k < i, j \geq i\})), \quad i \geq 0.$$

Moreover, for each $i \geq 0$, it holds that $GD^i = GD(D, E^i)$ and $GD^{i+1} = GD(D, E^{i+1})$ are of the form $GD^i = R \cup \{d'\}$ and $GD^{i+1} = R \cup \{d\}$ such that $d < d'$, where $R = GD^i \cap GD^{i+1}$, $d = \top : a_i/a_i$, and $d' = \top : b_i/b_i$.

Hence, if χ satisfies Principle I, then it must not select E^i . Since this holds for all $i \geq 0$, χ cannot select any classical extension. Observe that no preferred extensions exist (cf. Proposition 3). This proves unsatisfiability of Principle I.

That also Principle II is unsatisfiable for any CPR χ is exemplified by the following prioritized default theory $\Delta = (\emptyset, D, \{1 < 2\})$, which is rephrased from [7].

- (1) $\top : \neg b/c$
- (2) $\top : a/b$

The unique classical extension of Δ is $E = Th(\{b\})$, which must be selected by χ . Augment Δ by a default (0) $c : \top/\neg a$, such that $0 < 1$ and $0 < 2$; let Δ' be the resulting default theory. Clearly, $pre(0) \notin E$. However, E cannot be selected by χ , since Δ' has the unique preferred extension $E' = Th(\{\neg a, c\})$. Hence, χ violates Principle II. ■

This result tells us that we have to sacrifice the principles if we want to have a “weakly” preferred extension for each coherent default theory. We take this as additional support for our view that the preferences should be reconsidered in situations where no preferred extension exists.

Observe that the prioritized default theory Δ showing the failure of Principle I is infinite. It turns out that this is essential. In fact, over a finite (closed) Δ , the following CPR χ satisfying Principle I is possible. In the case in which Δ has a preferred extension, χ just returns that collection. In the case in which no preferred extension exists, fix a well-ordering $<'$ compatible with $<$ in $\Delta = (D, W, <)$, and define a relation \prec on the classical extensions of Δ by $E \prec E'$ iff $GD(D, E) = R \cup \{d\}$, $GD(D, E') = R \cup \{d'\}$ where $d, d' \notin R$ and $d <' d'$. Let then χ select the minimal elements of \prec , i.e., the classical extensions E such that $E' \not\prec E$ for all other classical extensions E' . It can be shown that \prec is irreflexive, and moreover that \prec has some minimal element. Hence, $\chi(\Delta, <')$ selects some classical extension(s), if some exist. Moreover, by construction of \prec it is easily seen that χ satisfies Principle I.

5 Expressing Preferences in the Language

For several applications like legal reasoning it is important to reason not only with, but also about the preferences among default rules. Such preferences often depend on the particular context at hand, and it is not possible to assign preferences independently of a particular context.

To make reasoning about default priorities possible we must be able to refer to defaults explicitly, and we must introduce a special predicate symbol representing default preferences. We, therefore, extend our logical language in two respects.

1. We introduce a distinct set of rule names N . A naming function assigns a unique name to default rules. Formally, default names are simply ground terms in the underlying language.
2. We use the reserved two-place infix predicate symbol \prec to represent default priority. For instance, if d_1 and d_2 are default names, then $d_1 \prec d_2$ is a formula with the intended meaning: d_1 has priority over d_2 .

Definition 9. *A preferential default theory is a triple $\Delta = (D, W, name)$ where*

1. (D, W) is a default theory,
2. $name : D \rightarrow N$ is an injective function, and
3. W contains axioms guaranteeing that \prec is a strict partial order.

Note that we do not restrict the appearance of \prec to W . It is possible (and useful) to have defaults which derive priority relations among other defaults.

An extension of a preferential default theory $\Delta = (D, W, name)$ is just a classical extension of (D, W) . The question now is how to define preferred extensions for preferential default theories.

All the derived preference information now is contained in the extensions of Δ . What we need is a way to eliminate an extension if it contains priority information which is in conflict with the way the extension was generated.

Given the techniques developed for prioritized default theories, it is not difficult to see how this can be done. Basically, an extension E of a preferential default theory is preferred iff E is a preferred extension of a fully prioritized default theory $(D, W, <)$ such that $<$ is compatible with the preference information in E . Compatibility is tested by generating a syntactic description of $<$ in terms of \prec and checking whether this description is consistent with E .

Definition 10. *Let $\Delta = (D, W, name)$ be a preferential default theory, E a classical extension of Δ . We say $<$ is compatible with E if and only if*

$$E \cup \{d_i \prec d_k \mid r_i < r_k, name(r_i) = d_i, name(r_k) = d_k\}$$

is consistent.

Definition 11. *Let $\Delta = (D, W, name)$ be a preferential default theory. Then, a set of formulas E is a preferred extension of Δ if and only if E is a preferred extension of some fully prioritized default theory $(D, W, <)$ such that $<$ is compatible with E .*

Example 7. Let's consider the following scenario. Your mother expects you to visit her on sundays. Your wife likes the opera and expects you to join her whenever Mozart is played. Unfortunately, visiting your mother and simultaneously going to the opera is impossible. Normally, the rules representing your mother's wishes must have preference over those representing your wife's wishes. However, if it is your wife's birthday, then you definitely should give preference to her.

This scenario can be modeled as the following preferential default theory.¹ The set of defaults D contains the following three rules:

$$\begin{aligned} (d_1) \quad & \textit{sunday} : \textit{visit}(\textit{mother}) / \textit{visit}(\textit{mother}) \\ (d_2) \quad & \textit{play}(\textit{mozart}) : \textit{go}(\textit{opera}) / \textit{go}(\textit{opera}) \\ (d_3[d, d']) \quad & \textit{mother_rule}(d) \wedge \textit{wife_rule}(d') : d \prec d' / d \prec d' \end{aligned}$$

The following formulas are in the background theory W :

$$\begin{aligned} & \textit{birthday}(\textit{wife}) \wedge \textit{wife_rule}(d) \wedge \textit{mother_rule}(d') \rightarrow d \prec d', \\ & \textit{mother_rule}(d_1), \quad \textit{wife_rule}(d_2), \quad \neg(\textit{visit}(\textit{mother}) \wedge \textit{go}(\textit{opera})) \end{aligned}$$

Now assume the following facts hold in addition (i.e., are in W):

$$\textit{sunday}, \quad \textit{play}(\textit{mozart})$$

We obtain two classical extensions. Note that both extensions contain the preference information $d_1 \prec d_2$. It is easy to see that there is no total preference relation $<$ compatible with this information such that E_2 is a preferred extension of $(D, W, <)$. Only the first extension is preferred and you should visit your mother.

Now consider what happens if we add $\textit{birthday}(\textit{wife})$. Again we obtain two extensions, E'_1 containing $\textit{visit}(\textit{mother})$, and E'_2 containing $\textit{go}(\textit{opera})$. In this case the preference information in both extensions is $d_2 \prec d_1$. Note that the applicability of default $d_3[d_1, d_2]$ is blocked. Now E'_1 cannot be reconstructed as preferred extension of a fully prioritized theory $(D, W, <)$ such that $<$ is compatible with E'_1, E'_2 , on the other hand, can be reconstructed in such a way. You just have to use an ordering with $r_2 < r_1$ where $\textit{name}(r_1) = d_1$ and $\textit{name}(r_2) = d_2$. That is, you should join your wife and go to the opera. ■

6 Related Work and Conclusion

In this paper, we have presented an approach, based on the ideas of [7], to incorporating priority information into default logic. This approach overcomes problems of previous approaches with respect to general principles which, as we argue, any prioritized variant of default logic should satisfy. For space reasons, a detailed comparison of our approach to the many other variants of prioritized default logic is necessarily superficial.

Rintanen's approach and the approaches in [16, 5, 1, 6] have already been briefly mentioned. The latter handle priorities such that in the (re)construction of an extension, only some of the applicable defaults can be fired in each step.

In [9] priorities are handled by encoding them into the object-level rather than constraining the construction of extensions at the meta-level. It turns out that in this approach some reasonable default theories do not possess any preferred extensions at all. For instance, no preferred extension exists for Example 1.

Other approaches, somewhat less related to our work, are concerned with handling specificity by respecting logical entrenchment of rules. In [8], an approach to handling specificity is developed which rewrites the defaults, based on their logical entrenchment, such that more specific rules are preferred. This is in the spirit of early versions

¹ We define *name* implicitly by putting the name of a default in front of the default.

of Nute’s defeasible logic (cf. [20]). Nute distinguishes defeasible and certain rules and presents a semantics with strong proof theoretic flavor for answering queries to the system.

Further related work is present in the context of logic programming, where different proposals to enhance extended logic programs with priorities have been made, cf. [25, 28, 13]. Possible extensions of these approaches to full default logic remain to be explored; however, on the common fragment of extended logic programs, these approaches differ from ours. For a discussion of further approaches to priorities and specificity in default logic, see [8, 1, 6].

Several issues remain for future work. First of all, procedures for reasoning from prioritized default theories need to be investigated. In the finite propositional case, brave and cautious reasoning in prioritized and classical default logic are polynomial time equivalent, and thus, by the results in [14], Σ_2^p and Π_2^p -complete, respectively. In fact, a suitable full prioritization of $\Delta = (W, D, <)$ such that E is a preferred extension of $\Delta' = (W, D, <')$ can be guessed, and the condition $E = C(\Delta_E^{*E})$ can be checked in polynomial time with an NP oracle. As a consequence, theorem provers for Reiter’s default logic can be used after a polynomial transformation for solving reasoning tasks in prioritized default logic. The design of genuine algorithms for prioritized default logic remains to be explored.

Another issue are approximations of preferred extensions. As we have shown, consistent preference relaxations (CPRs) of preferred extensions are subject to certain limitations. It would be interesting to see to what extent relaxations satisfying Principle I and II are possible, as well as for which weakenings of the principles CPRs do exist.

References

1. F. Baader and B. Hollunder. Priorities on Defaults with Prerequisite and their Application in Treating Specificity in Terminological Default Logic. *JAR*, 15:41–68, 1995.
2. S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade. Inconsistency Management and Prioritized Syntax-Based Entailment. *Proc. IJCAI-93*, pp. 640–645, 1993.
3. G. Brewka. Preferred Subtheories: An Extended Logical Framework for Default Reasoning. *Proc. IJCAI '89*, pp. 1043–1048, 1989.
4. G. Brewka. Adding Priorities and Specificity to Default Logic. *Proc. JELIA '94*, LNAI 838, pp. 247–260, 1994.
5. G. Brewka. Reasoning About Priorities in Default Logic. *Proc. AAI '94*, 1994.
6. G. Brewka. Well-Founded Semantics for Extended Logic Programs with Dynamic Preferences. *JAIR*, 4:19–36, 1996.
7. G. Brewka and T. Eiter. Preferred Answer Sets for Extended Logic Programs. *Proc. KR '98*, pp. 86–97, 1998.
8. J. Delgrande and T. Schaub. A General Approach to Specificity in Default Reasoning. *Proc. KR '94*, pp. 146–157, 1994.
9. J. Delgrande and T. Schaub. Compiling Reasoning With and About Preferences into Default Logic. *Proc. IJCAI '97*, pp. 168–174, 1997.
10. T. Eiter and G. Gottlob. The Complexity of Logic-Based Abduction. *JACM*, 42(1):3–42, 1995.
11. M. Gelfond and V. Lifschitz. The Stable Model Semantics for Logic Programming. *Proc. ICSLP*, pp. 1070–1080, 1988.

12. M. Gelfond, H. Przymusinska, and T. Przymusinski. On the Relationship Between Circumscription and Negation as Failure. *AIJ*, 38:75–94, 1989.
13. M. Gelfond and R. Son. Reasoning with Prioritized Defaults. Manuscript, 1997.
14. G. Gottlob. Complexity Results for Nonmonotonic Logics. *J Logic and Computation*, 2(3):397–425, June 1992.
15. K. Konolige. Hierarchic Autoepistemic Theories for Nonmonotonic Reasoning. In *Proc. AAAI '88*, 1988.
16. W. Marek and M. Truszczyński. *Nonmonotonic Logics*. Springer, 1993.
17. J. McCarthy. Circumscription – A Form of Non-Monotonic Reasoning. *AIJ*, 13:27–39, 1980.
18. R. Moore. Semantical Considerations on Nonmonotonic Logics. *AIJ*, 25:75–94, 1985.
19. B. Nebel. How Hard is it to Revise a Belief Base ? In: *Handbook on Defeasible Reasoning and Uncertainty Management Systems. Volume 2: Belief Change*, to appear.
20. D. Nute. Defeasible Logic. *Handbook of Logic in Artificial Intelligence and Logic Programming*, III, pp. 353–395. 1994.
21. D. Poole. On the Comparison of Theories: Preferring the Most Specific Explanation. *Proc. IJCAI '85*, 1985.
22. H. Prakken. *Logical Tools for Modelling Legal Argument*. Dissertation, Vrije Universiteit Amsterdam, 1993.
23. R. Reiter. A Logic for Default Reasoning. *AIJ*, 13:81–132, 1980.
24. J. Rintanen. On Specificity in Default Logic. *Proc. IJCAI '95*, pp. 1474–1479, 1995.
25. C. Sakama and K. Inoue. Representing Priorities in Logic Programs. *Proc. IJCSLP-96*, 1996.
26. D. Touretzky. *The Mathematics of Inheritance*. Pitman Research Notes in AI, London, 1986.
27. D. Touretzky, R. Thomason, and J. Horty. A Skeptic's Menagerie: Conflictors, Preemptors, Reinstaters, and Zombies in Nonmonotonic Inheritance. *Proc. IJCAI '91*, 1991.
28. Y. Zhang and N. Foo. Answer Sets for Prioritized Logic Programs. *Proc. ILPS '97*, pp. 69–83, 1997.