

Priorities on Defaults with Prerequisites, and Their Application in Treating Specificity in Terminological Default Logic*

FRANZ BAADER

*Lehr- und Forschungsgebiet Theoretische Informatik, RWTH Aachen, Ahornstraße 55
52074 Aachen, Germany. e-mail: baader@informatik.rwth-aachen.de*

and

BERNHARD HOLLUNDER

*Deutsches Forschungszentrum für KI (DFKI), Stuhlsatzenhausweg 3, 66123 Saarbrücken
Germany. e-mail: hollunder@dfki.uni-sb.de*

(Received: 30 March 1994; in final form 1 September 1994)

Abstract. In a recent paper we have proposed terminological default logic as a formalism that combines means both for structured representation of classes and objects and for default inheritance of properties. The major drawback that terminological default logic inherits from general default logic is that it does not take precedence of more specific defaults over more general ones into account. This behavior has already been criticized in the general context of default logic, but it is all the more problematic in the terminological case where the emphasis lies on the hierarchical organization of concepts.

The present paper addresses the problem of modifying terminological default logic such that more specific defaults are preferred. We assume that the specificity ordering is induced by the hierarchical organization of concepts, which means that default information is not taken into account when computing priorities. It turns out that the existing approaches for expressing priorities between defaults do not seem to be appropriate for defaults with prerequisites. Therefore we shall consider an alternative approach for dealing with prioritization in the framework of Reiter's default logic. The formalism is presented in the general setting of default logic where priorities are given by an arbitrary partial ordering on the defaults. We shall exhibit some interesting properties of the new formalism, compare it with existing approaches, and describe an algorithm for computing extensions. In the terminological case, we thus obtain an automated default reasoning procedure that takes specificity into account.

Key words: terminological default logic, default theories with priorities, knowledge representation.

AMS Subject Classification: 68T30

1. Introduction

Early knowledge representation formalisms such as semantic networks and frames comprised means for structured representation of classes and objects and for

* This is an extended version of a paper presented at the *13th International Joint Conference on Artificial Intelligence*, August 1993, Chambéry, France.

default inheritance of properties. However, these formalisms did not have a well-defined formal semantics, and subsequent formalisms trying to overcome this problem usually concentrated on one of these two means of representation. Nonmonotonic inheritance networks are concerned with defeasible inheritance, sometimes in combination with strict inheritance, but the nodes in these networks are unstructured objects or classes.¹ Terminological representation formalisms, on the other hand, can be used to define the relevant concepts of a problem domain in a structured and well-formed way. This is done by building complex concept descriptions out of atomic concepts (unary predicates) and roles (binary predicates) with the help of operations provided by the concept language of the particular formalism. In addition, objects can be described with respect to their relation to concepts and their interrelation with each other. The concept descriptions are interpreted as universal statements, which means that they do not allow for exceptions. As a consequence, the terminological system can use descriptions to insert concepts automatically at the proper place in the concept hierarchy (classification), and it can use the facts stated about objects to deduce to which concepts they must belong, but objects cannot inherit properties by default.

The problem addressed in this paper is how to bring together both means of representation originally present in semantic networks and frames, without losing the advantages of terminological formalisms, such as being equipped with a formal and well-understood semantics and providing for automated reasoning services such as concept classification. An integration of defaults would often greatly enhance applicability of terminological systems, or would at least make their use more convenient in most applications (see, e.g., [18], which shows that embedding defaults into terminological systems is an important item on the wish list of users of such systems). For this reason, several existing terminological systems, such as BACK [16], CLASSIC [4], K-Rep [13], LOOM [14], or SB-ONE [11], have been or will be extended to provide the user with some kind of default reasoning facilities. As the designers of these systems themselves point out, however, these approaches usually have an ad hoc character and thus do not satisfy the requirement of having a formal semantics.

As a first attempt to give a formally well-founded solution to this problem, an integration of Reiter's default logic into a terminological formalism was proposed in [1]. One reason for selecting default logic, out of the wide range of nonmonotonic formalisms, was that Reiter's default rule approach fits well into the philosophy of terminological systems. Most of these systems already provide their users with a form of 'monotonic' forward rules, and it turned out that these rules can be viewed as specific default rules where the justifications are absent. A second pleasant feature of terminological default logic, as introduced in [1], is that it becomes decidable provided that applicability of default rules is restricted to objects explicitly present in the knowledge base. It should be noted that this constraint is also imposed on the monotonic rules in terminological systems.

Decidability of reasoning in the formalism is an indispensable requirement here; otherwise, the important inference services of terminological systems (such as classification) could not be implemented in a satisfactory way.

The major drawback, which terminological default logic inherits from general default logic, is that it does not take precedence of more specific defaults over more general ones into account. For example, assume that we have a default that says that penguins cannot fly,² and another one that says that birds can fly, and that classification shows that penguins are a subconcept of birds. Intuitively, for any penguin the more specific first default should be preferred, which means that there should be only one default extension in which the penguin cannot fly. However, in default logic the first default has no priority over the second one, which means that one also gets a second extension where the penguin can fly. This behavior has already been criticized in the general context of default logic, but it is all the more problematic in the terminological case where the emphasis lies on the hierarchical organization of concepts.

In the present paper we shall consider the problem of modifying terminological default logic such that more specific defaults are preferred. After a short recapitulation of default logic and its specialization, terminological default logic, in Section 2, we shall consider the existing approaches for expressing priorities between defaults and shall point out why they do not seem to be appropriate for our purpose (see Section 3). For this reason we present in Section 4 an alternative approach for dealing with prioritization in the framework of Reiter's default logic. The formalism is presented in the general setting of default logic where priorities are given by an arbitrary partial ordering on the defaults. For terminological default theories the priorities between defaults will be induced by the position of their prerequisites in the concept hierarchy. Thus, the specificity relation that we use is determined by the strict information. We do not consider specificity induced by the defaults.

We shall exhibit some interesting properties of the new formalism and shall compare it with existing approaches. It turns out that every extension according to our definition (P-extension) is an extension according to Reiter's definition (R-extension); however, R-extensions that are not compatible with the partial ordering on defaults are excluded by our formalism. Not all default theories with an R-extension have a P-extension, but every normal default theory has a P-extension. If the defaults are further restricted to prerequisite-free normal defaults, then our approach coincides with the 'ordered default theories' of Brewka and Junker [7, 10]. In Section 5 we describe a method for computing P-extensions. For terminological default theories with specificity, this method is effective, which shows that this type of default reasoning can be automated. During the preparation of this report we have learned that Brewka [6] has also proposed a generalization of his ordered default theories to the case of normal defaults with prerequisites. In Section 6 we shall briefly introduce Brewka's approach and point out the differences in our approach. All proofs are deferred to the appendix.

2. Default Logic

This section briefly reviews Reiter's default logic and its specialization, terminological default logic (see [19] and [1] for details).

2.1. REITER'S DEFAULT LOGIC

Reiter [19] deals with the problem of how to formalize nonmonotonic reasoning by introducing nonstandard, nonmonotonic inference rules, which he calls default rules. A *default rule* is any expression of the form

$$\frac{\alpha : \beta}{\gamma},$$

where α, β, γ are first-order formulae. Here α is called the *prerequisite* of the rule, β is its *justification*, and γ its *consequent*.³ For a set of default rules \mathcal{D} , we denote the sets of formulae occurring as prerequisites, justifications, and consequents in \mathcal{D} by $\text{Pre}(\mathcal{D})$, $\text{Jus}(\mathcal{D})$, and $\text{Con}(\mathcal{D})$, respectively.

A default rule is *closed* iff α, β, γ do not contain free variables. It is *semi-normal* iff its justification implies the consequent, and it is *normal* if its justification and consequent are identical. A *default theory* is a pair $(\mathcal{W}, \mathcal{D})$, where \mathcal{W} is a set of closed first-order formulae (the world description) and \mathcal{D} is a set of default rules. A default theory is *closed* iff all its default rules are closed.

Intuitively, a *closed* default rule can be applied, that is, its consequent is added to the current set of beliefs, if its prerequisite is already believed and its justification is consistent with the set of beliefs. Formally, the consequences of a *closed* default theory are defined with reference to the notion of an *extension* (called *R-extension* in this paper), which is a set of deductively closed first-order formulae defined by a fixed point construction (see [19], p. 89). In general, a closed default theory may have more than one R-extension, or even no extension. Depending on whether one wishes to employ skeptical or credulous reasoning, a closed formula δ is a *consequence of a closed default theory* iff it is in all R-extensions or if it is in at least one R-extension of the theory.

Reiter also gives an alternative characterization of an R-extension, which we shall use, in a slightly modified way, as the definition of R-extension. Here and in the following, $\text{Th}(\Gamma)$ stands for the deductive closure of a set of formulae Γ .

DEFINITION 2.1. Let \mathcal{E} be a set of closed formulae, and $(\mathcal{W}, \mathcal{D})$ be a closed default theory. We define $E_0 := \mathcal{W}$ and for all $i \geq 0$

$$E_{i+1} := E_i \cup \{\gamma \mid \alpha : \beta / \gamma \in \mathcal{D}, \alpha \in \text{Th}(E_i), \text{ and } \neg\beta \notin \mathcal{E}\}.$$

Then \mathcal{E} is an *R-extension* of $(\mathcal{W}, \mathcal{D})$ iff $\mathcal{E} = \bigcup_{i \geq 0} \text{Th}(E_i)$.

Note that the R-extension \mathcal{E} to be constructed by this iteration process occurs in the definition of each iteration step. Since we are only adding consequents

of defaults during the iteration, any R-extension \mathcal{E} of $(\mathcal{W}, \mathcal{D})$ is of the form $\text{Th}(\mathcal{W} \cup \text{Con}(\widehat{\mathcal{D}}))$ for a subset $\widehat{\mathcal{D}}$ of \mathcal{D} . An easy consequence of the definition is that $(\mathcal{W}, \mathcal{D})$ has an inconsistent R-extension iff \mathcal{W} is inconsistent.

To generalize the notion of an R-extension to arbitrary default theories, one just assumes that a default with free variables stands for all its ground instances. In Reiter's original semantics the world description and the consequents of all defaults have to be Skolemized before building ground instances (over the enlarged signature). As shown in [1], Skolemization leads to both semantic and algorithmic problems, which is the reason why we shall dispense with it in the case of terminological default theories.

2.2. TERMINOLOGICAL DEFAULT LOGIC

Instead of formally introducing a particular terminological language (see, e.g., [1] for details), we shall just mention the features of terminological languages that will be important for the following. The *terminological part* of such languages allows one to build complex concept descriptions out of atomic concepts (unary predicates) and roles (binary predicates). For our purposes it suffices to know that a concept description C can be regarded as a first-order formula $C(x)$ with one free variable x . The subsumption hierarchy between concepts corresponds to implication of formulae: C is subsumed by D iff $\forall x: C(x) \rightarrow D(x)$ is valid.

The *assertional part* of the language can be used to state that an object is an instance of a concept C or that two individuals are connected by a role R . Logically, this means that one has constant symbols a, b as names for objects and can build formulae $C(a)$ and $R(a, b)$ by respectively substituting a for the free variable in $C(x)$ and applying the binary predicate R to the constants a, b . A finite set of such formulae is called an *ABox*. Important inference problems for ABoxes are whether a given ABox is consistent (consistency problem) and whether an object a is an instance of a concept C (instantiation problem), that is, whether $C(a)$ is a logical consequence of the given ABox. It should be noted that the formulae $C(x)$ obtained as concept descriptions of a terminological language belong to a restricted subclass of all first-order formulae with one free variable. For this reason the subsumption, consistency, and instantiation problems are usually decidable for these languages.

A *terminological default theory* is a pair $(\mathcal{A}, \mathcal{D})$, where \mathcal{A} is an ABox and \mathcal{D} is a finite set of default rules whose prerequisites, justifications, and consequents are concept descriptions. Obviously, since ABoxes can be seen as sets of closed formulae, and since concept descriptions can be seen as formulae with one free variable,⁴ terminological default theories are subsumed by Reiter's notion of an open default theory. However, as motivated in Sections 3 and 4 of [1], we do not Skolemize before building ground instances. This means that an open default of a terminological default theory is interpreted as representing all closed defaults that can be obtained by instantiating the free variable by all object names occurring

in the ABox. With this interpretation, it is possible to compute all R-extensions of terminological default theories (see [1], Sections 5 and 6).

3. Previous Approaches to Prioritization

When conflicts occur in reasoning with defaults, it is quite obvious that the more specific information should prevail over the more general one. In the context of terminological default theories this means that, for an instance of the concepts C and D , a default with prerequisite C should be preferred if C is strictly subsumed by D . As mentioned in the introduction, this requirement is not taken into account by Reiter's approach. If we assume that *penguin*, *bird*, and *flies* are appropriately defined concept descriptions, where *penguin* is subsumed by *bird*, then the terminological default theory consisting of the world description $\{penguin(Danny)\}$ and the defaults

$$\frac{penguin(x) : \neg flies(x)}{\neg flies(x)} \quad \text{and} \quad \frac{bird(x) : flies(x)}{flies(x)}$$

has two R-extensions. One contains $flies(Danny)$ and the other $\neg flies(Danny)$, and the semantics gives no reason for preferring the second one, in which the more specific default was applied.

To overcome this kind of problem, several approaches for realizing priorities among defaults have been proposed in the literature. The priorities may be induced by specificity of prerequisites (as described above), but may also come from other sources (such as reliability of defaults). Note, however, that we assume this ordering to be given before applying defaults. Thus, specificity information induced by defaults will usually not be taken into account, unless there is some way of doing this independently and a priori.

3.1. PRIORITIZATION VIA SEMI-NORMAL DEFAULTS

Reiter and Criscuolo show how some kind of prioritization between defaults can be achieved without changing the formalism by encoding the priority information into the justifications of semi-normal defaults [20]. If the first (more specific) default of our example should be preferred over the second one, the negated prerequisite of the first default has to be conjoined with the justification of the second one. In other words, the second default has to be rewritten to

$$\frac{bird(x) : flies(x) \wedge \neg penguin(x)}{flies(x)}$$

Although our simple example can be handled with this approach, it is not clear how to treat more complex situations. For example, if there is no direct conflict between the consequents of two defaults, then the default of lower priority should not generally be blocked by the prerequisite of the one of higher priority.

Blocking of the default of lower priority should be activated only if one is in a context where both consequents together lead to a contradiction. Reiter and Criscuolo do not describe a general method for solving these problems; they just “focus on certain fairly simple patterns of default rules”. Another problem is that, even if one starts with normal defaults (as in our example), one ends up with semi-normal defaults when realizing priorities this way. But this means that one has to face the undesirable properties of nonnormal defaults, such as nonexistence of extensions. As an additional problem, Brewka [8] points out that “whenever additional knowledge requires blocking of a default, the default has to be rewritten”.

3.2. PRIORITIZED DEFAULT THEORIES

To avoid the introduction of semi-normal defaults, Brewka [8] takes the ideas underlying prioritized circumscription [12] and defines an iterated version of default logic, which he calls prioritized default logic (PDL). Instead of one set of defaults he takes a finite number of sets $\mathcal{D}_1, \dots, \mathcal{D}_n$ of closed defaults, with the intended meaning that defaults in \mathcal{D}_i have higher priority than those in \mathcal{D}_j if $i < j$. PDL-extensions are defined by iterated application of Reiter’s definition of an extension: A set of closed formulae \mathcal{E} is a PDL-extension of a prioritized default theory $(\mathcal{W}, \mathcal{D}_1, \dots, \mathcal{D}_n)$ iff for all $i, 1 \leq i \leq n, \mathcal{E}_i$ is an R-extension of $(\mathcal{E}_{i-1}, \mathcal{D}_i)$, where $\mathcal{E}_0 = \mathcal{W}$ and $\mathcal{E} = \mathcal{E}_n$.

As pointed out by Brewka himself, this approach makes sense only if it is restricted to prerequisite-free normal defaults. The problem caused by prerequisites is demonstrated by the following abstract example. Assume that we have two levels of priority, the first consisting of the default $d_1 = \beta : \gamma/\gamma$, and the second of $d_2 = : \beta/\beta$. If we start with the empty world description, then the default d_1 cannot be applied when constructing the R-extension on the first level. On the second level, d_2 fires, and we get β in the extension $\mathcal{E} = \mathcal{E}_2$. Now the default d_1 could fire, but it is no longer considered on the second level.

If restricted to prerequisite-free normal defaults, prioritized default logic yields a prioritized version of Poole’s approach to default reasoning [7], and it seems to exhibit a quite reasonable behavior. One reason why this is nevertheless not an appropriate formalism for treating specificity in terminological default theories is that the defaults have to be put into levels of priorities which are totally ordered. However, subsumption gives us only a partial ordering on defaults.

3.3. ORDERED DEFAULT THEORIES

In [7, 10] the approach just described is generalized to the situation where priorities are given by an arbitrary partial ordering on defaults.

To be more precise, an *ordered default theory* is a triple $(\mathcal{W}, \mathcal{D}, <)$, where \mathcal{W} is a set of closed first-order formulae, \mathcal{D} is a set of closed prerequisite-free normal

defaults, and $<$ is a strict partial ordering on \mathcal{D} such that $\{d' \in \mathcal{D} \mid d' < d\}$ is finite for every $d \in \mathcal{D}$.

The principal idea is to consider total extensions of the partial ordering when computing extensions of the ordered default theory (which we shall call *B-extensions* in the following). Any enumeration d_1, d_2, \dots of \mathcal{D} that is compatible with the partial ordering (i.e., $i < k$ if $d_i < d_k$)⁵ defines a *B-extension* as follows. One starts with \mathcal{W} , and in the i -th step of the iteration, the consequent β_i of the default $d_i = : \beta_i / \beta_i$ is added if β_i is consistent with the set of formulae obtained after step $i - 1$. Otherwise, the current set of formulae remains unchanged. The limit of this process is the extension.

Interestingly, even for this simple case of normal defaults without prerequisites, there are different reasonable ways of handling priorities. In fact, Brass [5] has proposed an approach that is different from the one described above.

Even though Brewka's ordered default theories allow for priorities given by a partial ordering, this approach (as well as the one described by Brass) cannot directly be used to realize specificity in terminological default theories. The reason is that the restriction to prerequisite-free defaults is too severe. In fact, for terminological default theories the priorities we wished to consider were induced by subsumption relationships between the concept descriptions in the prerequisites. But this means that for prerequisite-free terminological defaults we no longer have a need for prioritization.

The situation is, however, not as bad as it seems. As shown in [3, 9], the closed normal default $\alpha : \beta / \beta$ can be approximated by the closed prerequisite-free normal default $\alpha \rightarrow \beta / \alpha \rightarrow \beta$. Thus one could start with a normal terminological default theory, determine the priorities between defaults from their prerequisites, and then transform the defaults into the corresponding ones without prerequisites. This way one ends up with an ordered default theory, which approximates the terminological default theory and which handles priorities induced by specificity of prerequisites in the terminological default theory.

However, we claim that this approach is still not satisfactory because it gives us a lot more than we bargained for. As pointed out in [9], the approximation not only gets rid of prerequisites but also equips the defaults with properties of classical implication, such as reasoning by cases and reasoning using contrapositives of the original defaults. For example, assume that, in addition to the concept descriptions *penguin*, *bird*, and *flies*, we have a description *winged* for objects having wings and that the only subsumption relation is the one between *penguin* and *bird*. If we consider the terminological default theory consisting of the world description $\{penguin(Danny)\}$ and the defaults

$$\frac{penguin(x) : \neg flies(x)}{\neg flies(x)}, \quad \frac{bird(x) : winged(x)}{winged(x)}, \quad \frac{winged(x) : flies(x)}{flies(x)},$$

then the preferred extension should be the one in which Danny has wings but does not fly.⁶

The approach we have described yields this extension; but it also yields another one in which Danny does not have wings, because as soon as the (approximation of the) first default has fired, the contrapositive of the third one can be fired, which gives us $\neg winged(Danny)$.

This shows that in this approach the defaults no longer behave like simple forward rules. But the similarity of default rules with the monotonic forward rules of terminological systems was one of our reasons for choosing default logic in the first place.

4. Default Theories with Priorities

To overcome the problems pointed out in the preceding section, we shall now propose a new approach for handling priorities among defaults with prerequisites. The semantics will be very close to Reiter's semantics, and the properties of our theory will also resemble those of Reiter's theory.

DEFINITION 4.1. A *default theory with priorities* is a triple $(\mathcal{W}, \mathcal{D}, <)$ consisting of a closed default theory $(\mathcal{W}, \mathcal{D})$ and a strict partial ordering $<$ on \mathcal{D} such that $\{d' \in \mathcal{D} \mid d' < d\}$ is finite for every $d \in \mathcal{D}$.

In the terminological case, \mathcal{W} is an ABox, and \mathcal{D} is obtained by instantiating the terminological default rules by all constants occurring in the ABox. For two instantiated terminological default rules d_1, d_2 with prerequisites $C_1(a_1), C_2(a_2)$ we have $d_1 < d_2$ iff they are concerned with the same object (i.e., $a_1 = a_2$) and C_1 is more specific than C_2 (i.e., C_1 is subsumed by C_2 but not vice versa). The restriction on the ordering is satisfied because \mathcal{D} is finite by definition of terminological default theories. In this case, $(\mathcal{W}, \mathcal{D}, <)$ is called *terminological default theory with specificity*.

Our definition of an extension for a default theory with priorities is modeled on Reiter's iterative characterization of R-extensions (see Definition 2.1). The main idea for treating priorities is that the consequent of a default can only be added during an iteration step if the default is not *delayed* by a preferred default, that is, there does not exist a smaller default that is currently active.

DEFINITION 4.2. For a set E of closed formulae, and a closed default $d = \alpha : \beta/\gamma$ we say that d is *active* in E iff its prerequisite is a consequence of E (i.e., $\alpha \in \text{Th}(E)$), its justification is consistent with E (i.e., $\neg\beta \notin \text{Th}(E)$), and its consequent is not a consequence of E (i.e., $\gamma \notin \text{Th}(E)$).

Using this notation, we can now give the main definition of this article.

DEFINITION 4.3. Let $(\mathcal{W}, \mathcal{D}, <)$ be a default theory with priorities, and let \mathcal{E} be a set of closed formulae. We define $E_0 := \mathcal{W}$, and for all $i \geq 0$

$$E_{i+1} := E_i \cup \{\gamma \mid \exists d \in \mathcal{D}: d = \alpha : \beta/\gamma, \alpha \in \text{Th}(E_i), \neg\beta \notin \mathcal{E}, \\ \text{and all } d' < d \text{ are not active in } E_i\}.$$

Then \mathcal{E} is a P-extension of $(\mathcal{W}, \mathcal{D}, <)$ iff $\mathcal{E} = \bigcup_{i \geq 0} \text{Th}(E_i)$.

The only difference in Reiter's characterization is the additional requirement that smaller defaults must not be active in the current state of the iteration. With this definition of an extension we get the intuitively correct result in our example with the three defaults concerning penguins, birds, and objects with wings. In fact, for any penguin the second default (asserting that birds normally have wings) can fire only after the more specific default (asserting that penguins normally cannot fly) has been applied. But this means that the third default (asserting that winged objects normally can fly) will never become applicable for a penguin (before its prerequisite becomes derivable, the negation of its justification must have been added). This means that our definition of a P-extension chooses from the two existing R-extensions the one that respects priorities.

Our first theorem states that this will always be the case, that is, that the set of all P-extensions is always a subset of the set of all R-extensions.

THEOREM 4.4. *Let \mathcal{E} be a P-extension of the default theory with priorities $(\mathcal{W}, \mathcal{D}, <)$. Then \mathcal{E} is an R-extension of $(\mathcal{W}, \mathcal{D})$.*

The proof is given in the appendix. The main idea is to take a P-extension \mathcal{E} that has been obtained from the sequence E_0, E_1, \dots and to use it to construct a sequence F_0, F_1, \dots as in the characterization of R-extensions. It is easy to see that $E_i \subseteq F_i$ for all $i \geq 0$, but the converse is not true. In fact, the consequent γ of a default d may be added to F_i but not to E_i because d is delayed by a smaller default that is active. A straightforward way to prove that $F_i \subseteq \mathcal{E} = \bigcup_{i \geq 0} \text{Th}(E_i)$ would thus be to show that the set of active defaults delaying d decreases along our F -iteration. Unfortunately, the set of defaults delaying d may also increase because prerequisites of smaller defaults that have not been derivable at step i may become derivable in a later step of the iteration. In the proof we shall circumvent this problem by considering the set of defaults that may potentially delay d , namely, defaults smaller than d that are currently active or may become active as soon as their prerequisite is derivable (see Appendix A.1 for details).

Since not all default theories have R-extensions, it follows that a default theory with priorities need not have a P-extension. But even if we have R-extensions, there need not exist P-extensions of a default theory with priorities. This is demonstrated by the following example.

EXAMPLE 4.5. Assume that \mathcal{W} is empty, and consider the three defaults

$$\frac{:\beta}{\beta}, \quad \frac{:\neg\beta}{\neg\beta} \quad \text{and} \quad \frac{\beta:\alpha}{\neg\alpha}.$$

We assume that the first default is smaller than the second one and that there are no other comparabilities with respect to $<$. This default theory has the R-extension $\text{Th}(\{\neg\beta\})$, but it does not have a P-extension. In fact, a P-extension

would prefer the first default, which yields β ; but then the third default (which is a modified version of the well-known one-rule example of a default theory having no R-extension) would become relevant.

As in the case without priorities, normal default theories with priorities have much nicer properties than arbitrary default theories with priorities.

THEOREM 4.6. *Every closed normal default theory with priorities has a P-extension.*

The proof, as given in the appendix, is a relatively straightforward adaption of Reiter's proof for R-extensions. To construct an R-extension $\mathcal{E} = \bigcup_{i \geq 0} \text{Th}(E_i)$ of a normal default theory $(\mathcal{W}, \mathcal{D})$, Reiter starts with the world description (i.e., $E_0 := \mathcal{W}$) and in successive steps adds as many consequents of active defaults as is possible without destroying consistency (i.e., $E_{i+1} := E_i \cup \text{Con}(\widehat{\mathcal{D}})$, where $\widehat{\mathcal{D}}$ is a maximal subset of the set of defaults active in E_i such that $E_i \cup \text{Con}(\widehat{\mathcal{D}})$ is consistent). To take priority information given by a strict partial order $<$ on \mathcal{D} into account, this construction is simply modified by considering only those active defaults that are minimal with respect to $<$ (see Appendix A.2 for details).

In Appendix A.2 we shall show that this construction always yields a P-extension of a normal default theory with priorities. But in general not all P-extensions can be obtained this way. The following example shows that this is true even for normal default theories without priorities.

EXAMPLE 4.7. Assume that \mathcal{W} is empty and \mathcal{D} contains the defaults

$$d_1 = \frac{:\alpha}{\alpha}, \quad d_2 = \frac{:\beta}{\beta} \quad \text{and} \quad d_3 = \frac{\alpha : \neg\beta}{\neg\beta}.$$

We assume that the ordering on defaults is empty, which means that here the notions R-extension and P-extension coincide. The default theory has two R-extensions, namely, $\mathcal{E}_1 = \text{Th}(\{\alpha, \beta\})$ and $\mathcal{E}_2 = \text{Th}(\{\alpha, \neg\beta\})$. However, \mathcal{E}_2 cannot be obtained by successively adding maximally consistent sets of consequents of active defaults. In fact, d_1 and d_2 are the only defaults that are active in $E_0 = \emptyset$. Since $E_0 \cup \{\alpha, \beta\}$ is consistent, there is exactly one maximal set $\widehat{\mathcal{D}} \subseteq \{d_1, d_2\}$ such that $E_0 \cup \text{Con}(\widehat{\mathcal{D}})$ is consistent, namely, $\{d_1, d_2\}$ itself. Thus the only set E_1 that one can get this way is $E_1 = \{\alpha, \beta\}$. This shows that extension \mathcal{E}_2 (which does not contain β) cannot be obtained.

Note that an alternative strategy, which adds only one consequent of a minimal active default in each step of the iteration, may result in an R-extension that is not a P-extension (see Example 6.2).

If we further restrict the attention to normal defaults *without prerequisites*, then the notion of a P-extension coincides with that of a B-extension, which shows that our approach is a generalization of ordered default theories.

THEOREM 4.8. *Let \mathcal{D} be a set of closed prerequisite-free normal defaults. Then \mathcal{E} is a P-extension of the default theory with priorities $(\mathcal{W}, \mathcal{D}, <)$ iff \mathcal{E} is a B-extension of the ordered default theory $(\mathcal{W}, \mathcal{D}, <)$.*

If \mathcal{D} was assumed to be finite, the proof would be relatively easy; but in the general case of possibly infinite sets of defaults, it becomes more involved (see Appendix A.3).

5. Computing P-Extensions

Since all P-extensions are R-extensions, one could first generate all R-extensions of a default theory, and then for each R-extension \mathcal{E} directly use the definition of P-extensions to check whether \mathcal{E} is a P-extension. For terminological default theories this provides us with an effective procedure for computing all P-extensions. In fact, in [1] it is shown how to compute all R-extensions of a terminological default theory. Since one has only finitely many closed defaults, and since the instantiation problem for the terminological languages we use in [1] is decidable, the iteration in the definition of a P-extension is effective as well.

However, there may exist a lot more R-extensions than P-extensions, and computing R-extensions is rather expensive. For this reason, it would be preferable to have an algorithm for directly computing P-extensions. The idea behind the algorithm presented below is to make an iteration similar to the one in the definition of a P-extension, but without already having the final set \mathcal{E} for controlling which consequents of defaults are added. After the iteration becomes stable (which will always be the case for finite sets of closed defaults), one has to check an additional condition to make sure that the result really is a P-extension.

The main problem is to determine which sets of consequents are candidates for being added in each step of the iteration. Of course there can be more than one correct choice because there may exist more than one P-extension. If we look at the definition of E_{i+1} in Definition 4.3, we see that the defaults whose consequents are added are defaults active in E_i that are minimal w.r.t. the priority order $<$. Which subset of their consequents is taken depends on the set \mathcal{E} used for the iteration. Since our algorithm does not know the final \mathcal{E} , it has to consider arbitrary subsets; but we shall see that there are some constraints that reduce the number of possible choices.

It should be noted that neither a greedy procedure (which takes maximal subsets that are consistent with what has already been computed) nor an overly modest procedure (which adds only one consequent in each step) would be complete. Example 4.7 demonstrates this for the greedy procedure, even in the absence of priority information. Examples that illustrate why the overly modest procedure is not appropriate for computing P-extensions can be found in the next section.

In the following (nondeterministic) algorithm, E_i will always be a subset of $\mathcal{W} \cup \text{Con}(\mathcal{D})$, and J_i will be a subset of $\neg\text{Jus}(\mathcal{D})$ (where, for a set \mathcal{F} of formulae, $\neg\mathcal{F} := \{\neg\beta \mid \beta \in \mathcal{F}\}$).

ALGORITHM 5.1. *Let $(\mathcal{W}, \mathcal{D}, <)$ be a closed default theory with priorities. If \mathcal{W} is inconsistent, then $\text{Th}(\mathcal{W})$ is the only P-extension. Otherwise we define $E_0 := \mathcal{W}$ and $J_0 := \emptyset$.*

Now assume that E_i ($i \geq 0$) is already defined. Consider

$$\mathcal{D}_{i+1} := \{d \in \mathcal{D} \mid d \text{ is active in } E_i \text{ and no } d' < d \text{ is active in } E_i\},$$

and choose a nonempty subset $\widehat{\mathcal{D}}_{i+1}$ of \mathcal{D}_{i+1} that satisfies

$$\neg\beta \notin \text{Th}(E_i \cup \text{Con}(\widehat{\mathcal{D}}_{i+1}) \cup J_i \cup \neg\text{Jus}(\mathcal{D}_{i+1} \setminus \widehat{\mathcal{D}}_{i+1})) \quad (*)$$

for all $\beta \in \text{Jus}(\widehat{\mathcal{D}}_{i+1})$.

If there is no such set, then $E_{i+1} := E_i$, $J_{i+1} := J_i$, and the iteration process becomes stable.⁷ Otherwise each choice yields new sets $E_{i+1} := E_i \cup \text{Con}(\widehat{\mathcal{D}}_{i+1})$ and $J_{i+1} := J_i \cup \neg\text{Jus}(\mathcal{D}_{i+1} \setminus \widehat{\mathcal{D}}_{i+1})$.

The set $\mathcal{E} := \bigcup_{i \geq 0} \text{Th}(E_i)$ is a P-extension iff

- (1) *for all $d = \alpha : \beta/\gamma \in \bigcup_{i \geq 1} \widehat{\mathcal{D}}_i$ we have $\neg\beta \notin \mathcal{E}$, and*
- (2) *for all $\neg\beta \in \bigcup_{i \geq 1} J_i$ we have $\neg\beta \in \mathcal{E}$.*

A proof of soundness and completeness of this algorithm is given in Appendix A.4. The idea behind the sets J_i is as follows. If the consequent of a minimal active default is not included in E_{i+1} , then the reason must be that its justification is not consistent with the final extension. Thus, if we exclude such a default from $\widehat{\mathcal{D}}_{i+1}$, we know that the negation of its justification must belong to the extension. The condition (*) on $\widehat{\mathcal{D}}_{i+1}$ corresponds to the fact that defaults whose consequents are added to a P-extension must have justifications that are consistent with the extension. This condition can ensure only local correctness of our choices. For this reason we must check the two conditions on \mathcal{E} to ensure global correctness. Because of this global test, the algorithm would also be correct without testing the local condition (*). Thus, testing this local condition is already an optimization of the brute force method, since it reduces the number of subsets of \mathcal{D}_{i+1} that must be considered.

For terminological default theories with specificity, all the steps of the algorithm are effective, provided that the consistency and instantiation problem for the underlying terminological language is decidable (an assumption which is usually satisfied). In addition, since one has only finitely many closed defaults, the iteration will become stable after finitely many steps.

A possible optimization of the algorithm for computing extensions is to partition the set of defaults into subsets that do not ‘influence’ each other. Then extensions can be computed separately for each subset, and only afterwards be put together to extensions of the whole set of defaults. We make this idea more

precise for the case of terminological default theories with specificity $(\mathcal{W}, \mathcal{D}, <)$. Recall that \mathcal{D} is obtained by instantiating open terminological default rules by all individual names (i.e., constants) occurring in the ABox \mathcal{W} . For two instantiated terminological default rules d_1, d_2 with prerequisites $C_1(a_1), C_2(a_2)$ we have $d_1 < d_2$ iff they are concerned with the same object (i.e., $a_1 = a_2$) and C_1 is more specific than C_2 (i.e., C_1 is subsumed by C_2 but not vice versa).

The set N of individual names in \mathcal{W} can be partitioned into connectivity classes as follows: we say that $a \in N$ is connected with $b \in N$ iff $a = b$ or there exist $n \geq 1$, role names R_1, \dots, R_n , and individual names $a_0, \dots, a_n \in N$ such that $a = a_0$, $b = a_n$, and for all $1 \leq i \leq n$, $R_i(a_{i-1}, a_i) \in \mathcal{W}$. Let N_1, \dots, N_k be the equivalence classes of the relation ‘is-connected’. This partition of N induces a partition $\mathcal{D}_1, \dots, \mathcal{D}_k$ of \mathcal{D} : the set \mathcal{D}_i consists of all instances of the original open terminological default rules with individuals from N_i . Obviously, if d_1 and d_2 belong to different such subsets, then there cannot be a priority relation between d_1 and d_2 . In addition, adding the consequent of d_1 to a partially constructed extension has no influence on the applicability of the default d_2 . In fact, assume that a and b are not connected in an ABox \mathcal{A} , and that $\mathcal{A} \cup \{C(a)\}$ is consistent. Then $D(b)$ is a consequence of $\mathcal{A} \cup \{C(a)\}$ iff it is already a consequence of \mathcal{A} .⁸

For $i = 1, \dots, k$, let G_i be the set of all extensions of $(\mathcal{W}, \mathcal{D}_i, <_i)$, where $<_i$ is the restriction of $<$ to \mathcal{D}_i . The set G of all extensions of $(\mathcal{W}, \mathcal{D}, <)$ is obtained by considering all possible unions of extensions from G_1, \dots, G_k , i.e.,

$$G = \{\text{Th}(\mathcal{E}_1 \cup \dots \cup \mathcal{E}_k) \mid \mathcal{E}_1 \in G_1, \dots, \mathcal{E}_k \in G_k\}.$$

If n_1, \dots, n_k are the cardinalities of the sets G_1, \dots, G_k , then G has cardinality $n_1 \times \dots \times n_k$. However, our optimized algorithm must only compute $n_1 + \dots + n_k$ extensions of smaller default theories.

We have the feeling that, compared with the case of default theories without priorities, it is rather hard to come up with additional conceptual optimizations. Since the definition of an extension strongly depends on the iteration process, it is hard to conceive of an algorithm for computing extensions that is not based on this process.⁹

6. Related Work

During the preparation of this report we learned that Brewka [6] has proposed a generalization of his ordered default theories (as described in Section 3.3) to the case of normal defaults with prerequisites. In this section we shall briefly introduce Brewka’s approach and point out the differences in our approach.

Brewka considers default theories with priorities $(\mathcal{D}, \mathcal{W}, <)$ where \mathcal{D} is a *finite* set of closed *normal* defaults. Any total extension \ll of $<$ defines an extension

$\bigcup_{i \geq 0} \text{Th}(F_i)$ (called B^* -extension in the following) as follows. Let $F_0 := \mathcal{W}$, and for $i \geq 0$

$$F_{i+1} := \begin{cases} F_i & \text{if there exists no default that is active in } F_i, \\ F_i \cup \{\beta\} & \text{otherwise, where } \beta \text{ is the consequent of the} \\ & \ll\text{-minimal default that is active in } F_i. \end{cases}$$

From the results in [6] it follows that every B^* -extension is an R-extension. This means that, as in our approach, Brewka takes a subset of the set of all R-extensions as admissible extensions of a default theory with priorities. The following two examples show, however, that in general he takes another subset than we do. The first example demonstrates that one may obtain more P-extensions than B^* -extensions.

EXAMPLE 6.1. Assume that \mathcal{W} is empty and that \mathcal{D} consists of the normal defaults

$$d_1 = \frac{: \alpha}{\alpha}, \quad d_2 = \frac{\beta : \neg \alpha}{\neg \alpha}, \quad d_3 = \frac{: \beta}{\beta}, \quad \text{and} \quad d_4 = \frac{\alpha : \neg \beta}{\neg \beta},$$

where $d_2 < d_1, d_4 < d_3$, and these are the only $<$ -relationships between defaults in \mathcal{D} . First, we show that $\mathcal{E} := \text{Th}(\{\alpha, \beta\})$ is a P-extension. In fact, using \mathcal{E} in the definition of P-extensions yields $E_0 = \emptyset, E_1 = \{\alpha, \beta\}$, and $E_2 = E_1$. Since d_2 and d_4 are not active in E_0 (their prerequisites are not deducible), d_1 and d_3 are the minimal defaults that are active in E_0 . In addition, we have $\neg \alpha, \neg \beta \notin \mathcal{E}$, which shows that $E_1 = \{\alpha, \beta\}$. But then d_2 and d_4 are not active in E_1 (their negated justifications are deducible), which shows that $E_1 = E_2$.

With Brewka's definition of an extension, either α or β is added in the first step of the iteration, depending on whether $d_1 \ll d_3$ or $d_3 \ll d_1$. This is so because again d_2 and d_4 are not active in $F_0 = \emptyset$. We restrict our attention to the case $F_1 = \{\alpha\}$. (The case $F_1 = \{\beta\}$ is symmetric.) Now d_3 and d_4 are active in F_1 ; but since \ll is an extension of $<$, we know that d_4 is the minimal default that is active. For this reason, we get $F_2 = \{\alpha, \neg \beta\}$, which shows that we cannot get $\mathcal{E} = \text{Th}(\{\alpha, \beta\})$ this way.

It is easy to see that in this example the two B^* -extensions $\text{Th}(\{\alpha, \neg \beta\})$ and $\text{Th}(\{\neg \alpha, \beta\})$ are also P-extensions. But in general, B^* -extensions need not be P-extensions. This is demonstrated by the next example.

EXAMPLE 6.2. Assume that \mathcal{W} is empty and that \mathcal{D} consists of the normal defaults

$$d_1 = \frac{: \alpha}{\alpha}, \quad d_2 = \frac{: \beta}{\beta}, \quad d_3 = \frac{\beta : \gamma}{\gamma}, \quad \text{and} \quad d_4 = \frac{\alpha : \neg \gamma}{\neg \gamma}.$$

The only $<$ -relationship that exists is $d_3 < d_4$. Now $\mathcal{E} := \text{Th}(\{\alpha, \beta, \neg \gamma\})$ is a B^* -extension, but not a P-extension.

To show that \mathcal{E} is a B^* -extension, we consider the total extension $d_1 \ll d_3 \ll d_4 \ll d_2$ of \langle . With this ordering, we obviously get $F_0 = \emptyset$, and $F_1 = \{\alpha\}$. Now d_4 is active in F_1 , and the only smaller default, d_3 , is not active. Thus we get $F_2 = \{\alpha, \neg\gamma\}$. In the next step, the only active default is d_2 , which means that $F_3 = \{\alpha, \neg\gamma, \beta\}$. In F_3 , no defaults are active, and thus $F_4 = F_3$. This shows that $\mathcal{E} = \text{Th}(F_3)$ is a B^* -extension.

Let us now show that \mathcal{E} is not a P-extension. Using \mathcal{E} in the iterative definition of P-extensions, we obviously get $E_0 = \mathcal{W}$, and $E_1 = \{\alpha, \beta\}$. Now observe that $d_3 < d_4$, and d_3 is active in E_1 . But this means that neither the consequent of d_3 can be added (since the negation of d_3 's justification is in \mathcal{E}) nor the consequent of d_4 (since d_4 is delayed by d_3). Thus $E_1 = E_2$, which shows that $\bigcup_{i \geq 0} \text{Th}(E_i) = \text{Th}(\{\alpha, \beta\}) \neq \mathcal{E}$.

Brewka's and our approach use different (but equivalent) characterizations of R-extensions as starting point. In fact, for characterizing all R-extensions one can use either the iteration process described in Section 2 above (which adds as many consequents of defaults as possible in each step), or a similar iteration process that adds only one consequent in each step. Our approach for handling priorities generalizes the first characterization, whereas Brewka's approach generalizes the second characterization. From the abstract examples, it is not clear which generalization is more appropriate, and we do not believe that the question of which approach is better can be resolved in general.

In [6], Brewka also considers an alternative approach for handling priorities. In Example 6.2, this approach yields the same extensions as ours, namely, $\text{Th}(\{\alpha, \beta, \gamma\})$. However, it is based on an intuition on when to delay a default, which is very different from our approach and Brewka's B^* -extension approach: A smaller default d_1 can delay an active default d_2 in E_i even if the prerequisite of d_1 is not in $\text{Th}(E_i)$. It must only be the case that the prerequisite of d_1 will be contained in the final extension. Thus, in Example 6.2, d_4 is delayed by d_3 in F_1 , even though the prerequisite β of d_3 is not yet present in F_1 . This is the reason why the B^* -extension $\mathcal{E} = \text{Th}(F_3)$ is not admissible in Brewka's alternative approach.

It is easy to see, however, that this alternative approach is also orthogonal to ours. For example, the P-extension $\mathcal{E} = \text{Th}(\{\alpha, \beta\})$ in Example 6.1 is not an extension with respect this approach. In fact, since α and β are in \mathcal{E} , the defaults d_1 and d_3 are delayed by d_2 and d_4 .

7. Conclusion

We have addressed the question of how to prefer more specific defaults over more general ones. This problem is of general interest for default reasoning but is even more important in the terminological case where the emphasis lies on the hierarchical organization of concepts. Of the previously existing approaches for handling priorities among defaults, Brewka's ordered default theories turned out

to come nearest to what is needed for solving the specificity problem in terminological default theories. But its restriction to prerequisite-free normal defaults seems to be too severe to make it an adequate solution in the terminological case.

Therefore we have proposed a new approach, called default theories with priorities, for handling priorities among defaults *with prerequisites*. The properties we could prove for this formalism demonstrate that it is a quite reasonable generalization of Reiter's default logic and of Brewka's ordered default theories. In addition, it correctly handles examples for which the other approaches give unintuitive results.

Brewka's independently developed generalization of ordered default theories to the case of normal defaults with prerequisites turned out to be orthogonal to our approach, in the sense that there are extensions obtained with his approach that are not obtained with ours and vice versa.

We have also described a method for generating the extensions of a default theory with priorities. This method is effective provided that the base logic is decidable and one has only finitely many closed defaults. These restrictions are satisfied in the terminological case, which means that terminological default logic with specificity is decidable. Consequently, it is possible to equip a terminological representation system with appropriate automated default reasoning services.

The priority ordering we have proposed for terminological default theories takes into account only the strict subsumption links between prerequisites of terminological defaults. If one wishes to consider specificity induced by defaults as well, one can for example adapt the method proposed by Brewka ([6], Definition 3) to our approach. Another interesting point for further research is to consider priorities on terminological defaults that not only take into account subsumption between prerequisites of defaults, but also consider the role relationships in ABoxes.

Appendix. Proofs of Theorems

A.1. PROOF OF THEOREM 4.2

THEOREM 4.4. *Let \mathcal{E} be a P-extension of the default theory with priorities $(\mathcal{W}, \mathcal{D}, <)$. Then \mathcal{E} is an R-extension of $(\mathcal{W}, \mathcal{D})$.*

Assume that $\mathcal{E} = \bigcup_{i \geq 0} \text{Th}(E_i)$ is a P-extension of $(\mathcal{W}, \mathcal{D}, <)$, obtained by iteratively generating sets E_0, E_1, \dots as described in Definition 4.3. Before we can prove that \mathcal{E} is also an R-extension, we need a technical definition and two lemmas.

For all $d \in \mathcal{D}$ and $i \geq 0$ we define

$$D_i^d = \left\{ d' = \frac{\alpha' : \beta'}{\gamma'} \mid d' < d \text{ and } (\alpha' \notin \text{Th}(E_i) \text{ or } d' \text{ is active in } E_i) \right\}.$$

The set D_i^d contains all defaults that either delay d in step i of the iteration or may delay d in a later step when their prerequisite becomes deducible. By our assumption on the strict partial ordering $<$, there exist only finitely many $d' < d$, which means that the sets D_i^d are always finite. The first lemma shows that D_i^d stays the same or gets smaller when the index is increased.

LEMMA A.1. *For all $d \in \mathcal{D}$ and all $i \geq 0$ we have $D_i^d \supseteq D_{i+1}^d$.*

Proof. Let $d' = \alpha' : \beta' / \gamma'$ be a default contained in D_{i+1}^d . First assume that $\alpha' \notin \text{Th}(E_{i+1})$. Now $E_i \subseteq E_{i+1}$ yields $\alpha' \notin \text{Th}(E_i)$, and therefore d' is contained in D_i^d .

On the other hand, assume that d' is active in E_{i+1} , i.e., $\alpha' \in \text{Th}(E_{i+1})$, $\neg\beta' \notin \text{Th}(E_{i+1})$, and $\gamma' \notin \text{Th}(E_{i+1})$. If $\alpha' \notin \text{Th}(E_i)$, then d' is in D_i^d , and we are done. For $\alpha' \in \text{Th}(E_i)$, the default d' is active in E_i , since $\neg\beta' \notin \text{Th}(E_{i+1})$ and $\gamma' \notin \text{Th}(E_{i+1})$ together with $E_i \subseteq E_{i+1}$ imply $\neg\beta' \notin \text{Th}(E_i)$ and $\gamma' \notin \text{Th}(E_i)$. \square

Our next lemma states that a default whose prerequisite is believed in some state of the iteration and whose justification is consistent with \mathcal{E} will eventually fire during the iteration.

LEMMA A.2. *Let $d = \alpha : \beta / \gamma \in \mathcal{D}$ and $i \geq 0$ be such that $\alpha \in \text{Th}(E_i)$ and $\neg\beta \notin \mathcal{E}$. Then there exists an index $j > i$ such that $\gamma \in \text{Th}(E_j)$.*

Proof. The lemma is proved by induction on the cardinality of D_i^d . Let A_i^d be the set of all defaults smaller than d that are active in E_i , namely,

$$A_i^d := \{d' < d \mid d' \text{ is active in } E_i\}.$$

Obviously, A_i^d is contained in D_i^d for all $i \geq 0$. First assume that A_i^d is empty, in other words, all defaults smaller than d are not active in E_i . This, together with our assumption that $\alpha \in \text{Th}(E_i)$ and $\neg\beta \notin \mathcal{E}$, implies $\gamma \in E_{i+1}$. Thus we can take $j = i + 1$.

Now assume that A_i^d is not empty. Let $d' = \alpha' : \beta' / \gamma'$ be a minimal default with respect to $<$ in A_i^d . We distinguish two cases.

Case 1: $\neg\beta' \notin \mathcal{E}$. Since d' is active, we also know that $\alpha' \in \text{Th}(E_i)$, and minimality of d' implies that it cannot be delayed by a smaller default that is active in E_i . This shows that $\gamma' \in E_{i+1}$, and thus $d' \notin D_{i+1}^d$. Together with Lemma A.1 this yields $D_i^d \supset D_{i+1}^d$.

Case 2: $\neg\beta' \in \mathcal{E}$. Since $\mathcal{E} = \bigcup_{j \geq 0} \text{Th}(E_j)$, there exists $j \geq 0$ such that $\neg\beta' \in \text{Th}(E_j)$. In addition, the fact that d' is active in E_i implies that $\neg\beta' \notin \text{Th}(E_i)$, which yields $j > i$. From $\neg\beta' \in \text{Th}(E_j)$ and $\alpha' \in \text{Th}(E_i) \subseteq \text{Th}(E_j)$ we can deduce $d' \notin D_j^d$. This yields $D_i^d \supset D_j^d$.

We have seen that in both cases there exists an index $j > i$ such that $D_i^d \supset D_j^d$. Obviously, $j > i$ implies that $\alpha \in \text{Th}(E_i) \subseteq \text{Th}(E_j)$, which shows that j also

satisfies the assumption of the lemma. By induction we get an index $j' > j > i$ with $\gamma \in \text{Th}(E_{j'})$. \square

To prove Theorem 4.4, we take the P-extension \mathcal{E} , which has been obtained from the sequence E_0, E_1, \dots , and use it to construct a sequence F_0, F_1, \dots as described in the characterization of R-extensions, in other words, $F_0 := \mathcal{W}$, and for all $i \geq 0$

$$F_{i+1} := F_i \cup \{\gamma \mid \alpha : \beta/\gamma \in \mathcal{D}, \alpha \in \text{Th}(F_i), \text{ and } \neg\beta \notin \mathcal{E}\}.$$

To show that \mathcal{E} is an R-extension, we have to prove that $\mathcal{E} = \bigcup_{i \geq 0} \text{Th}(F_i)$.

LEMMA A.3. *For all $i \geq 0$ we have $E_i \subseteq F_i$.*

Proof. This can easily be proved by induction on i . \square

In general, the other direction does not hold: F_i is not necessarily a subset of E_i . But for all $i \geq 0$ we get $F_i \subseteq \mathcal{E} = \bigcup_{j \geq 0} \text{Th}(E_j)$ as an immediate consequence of the next lemma.

LEMMA A.4. *For every $i \geq 0$ and every $\gamma \in F_i$ there exists an index j such $\gamma \in \text{Th}(E_j)$.*

Proof. The lemma is proved by induction on i . For $i = 0$ there is nothing to show because $F_0 = \mathcal{W} = E_0$. Now assume that $i > 0$.

Let γ be an element of F_i . If $\gamma \in F_{i-1}$, we know by induction that $\gamma \in \text{Th}(E_j)$ for some j . Thus assume that $\gamma \notin F_{i-1}$. Consequently, γ is the consequent of a default $\alpha : \beta/\gamma$ whose prerequisite α is in $\text{Th}(F_{i-1})$ and whose justification is consistent with \mathcal{E} .

Let $\gamma_1, \dots, \gamma_n$ be formulae in F_{i-1} such that $\alpha \in \text{Th}(\{\gamma_1, \dots, \gamma_n\})$. By induction we know that for each of these formulae $\gamma_k \in \text{Th}(F_{i-1})$ there exists an index j_k such that $\gamma_k \in \text{Th}(E_{j_k})$. For $j = \max\{j_1, \dots, j_n\}$ we have $\{\gamma_1, \dots, \gamma_n\} \subseteq \text{Th}(E_j)$, which implies $\alpha \in \text{Th}(E_j)$.

Since we also know that $\neg\beta \notin \mathcal{E}$, the assumptions of Lemma A.2 are satisfied for d and j . Thus we can conclude that there exists an index $j' > j$ such that $\gamma \in \text{Th}(E_{j'})$. \square

A.2. PROOF OF THEOREM 4.4

THEOREM 4.6. *Every closed normal default theory with priorities has a P-extension.*

Let $(\mathcal{W}, \mathcal{D}, <)$ be a closed normal default theory with priorities. If \mathcal{W} is inconsistent, then $\text{Th}(\mathcal{W})$ is a P-extension. Thus assume that \mathcal{W} is consistent.

A P-extension of $(\mathcal{W}, \mathcal{D}, <)$ is inductively constructed as follows. We define $E_0 := \mathcal{W}$, and for all $i \geq 0$

$$\mathcal{D}_{i+1} := \{d \in \mathcal{D} \mid d \text{ is active in } E_i, \text{ and no } d' < d \text{ is active in } E_i\}.$$

Let $E_{i+1} := E_i \cup \text{Con}(\widehat{\mathcal{D}}_{i+1})$, where $\widehat{\mathcal{D}}_{i+1}$ is a maximal subset of \mathcal{D}_{i+1} such that $E_i \cup \text{Con}(\widehat{\mathcal{D}}_{i+1})$ is consistent.

By compactness, we know that $\mathcal{E} := \bigcup_{i \geq 0} \text{Th}(E_i)$ is also consistent. To show that \mathcal{E} is a P-extension, we have to prove $\mathcal{E} = \bigcup_{i \geq 0} \text{Th}(F_i)$, where $F_0 := \mathcal{W}$, and

$$F_{i+1} := F_i \cup \{\beta \mid \exists d \in \mathcal{D} : d = \alpha : \beta / \beta, \alpha \in \text{Th}(F_i), \neg\beta \notin \mathcal{E}, \\ \text{and all } d' < d \text{ are not active in } F_i\}.$$

This is an immediate consequence of the following lemma. Note that we need not have $E_i = F_i$. In fact, $F_i \setminus F_{i-1}$ may contain elements γ of $\text{Th}(F_{i-1}) \setminus F_{i-1}$, whereas elements of $\text{Th}(E_{i-1}) \setminus E_{i-1}$ are not in E_i , by definition of active.

LEMMA A.5. *For all $i \geq 0$ we have $\text{Th}(E_i) = \text{Th}(F_i)$.*

Proof. The lemma is proved by induction on i . For $i = 0$ there is nothing to show, since $E_0 = \mathcal{W} = F_0$. Thus assume that $i \geq 0$ and that we already know that $\text{Th}(E_i) = \text{Th}(F_i)$.

First, we show that $E_{i+1} \subseteq \text{Th}(F_{i+1})$. Let β be an element of E_{i+1} . If $\beta \in \text{Th}(E_i)$, we know by induction that $\beta \in \text{Th}(F_i) \subseteq \text{Th}(F_{i+1})$.

Now assume that $\beta \in E_{i+1} \setminus \text{Th}(E_i)$. Thus $\beta \in \text{Con}(\widehat{\mathcal{D}}_{i+1})$, and hence there exists a default $d = \alpha : \beta / \beta$ in \mathcal{D} such that d is active in E_i and no default $d' < d$ is active in E_i .

To prove that $\beta \in F_{i+1}$, and thus $\beta \in \text{Th}(F_{i+1})$, we have to show that no default $d' < d$ is active in F_i , $\alpha \in \text{Th}(F_i)$, and $\neg\beta \notin \mathcal{E}$. The first two properties follow immediately from what we know about d , since $\text{Th}(F_i) = \text{Th}(E_i)$ (by induction). The third property follows from $\beta \in E_{i+1} \subseteq \mathcal{E}$ and the fact that \mathcal{E} is consistent. This concludes the proof of $E_{i+1} \subseteq \text{Th}(F_{i+1})$.

Now let us show that $F_{i+1} \subseteq \text{Th}(E_{i+1})$. Let β be an element of F_{i+1} . Again, the case $\beta \in \text{Th}(F_i)$ is trivial. Thus assume that $\beta \in F_{i+1} \setminus \text{Th}(F_i)$. This means that there exists a default $d = \alpha : \beta / \beta$ in \mathcal{D} such that no default $d' < d$ is active in F_i , $\alpha \in \text{Th}(F_i)$, and $\neg\beta \notin \mathcal{E}$. Now $\text{Th}(F_i) = \text{Th}(E_i)$ yields $d \in \mathcal{D}_{i+1}$. Note that we really need to know that $\beta \notin \text{Th}(F_i) = \text{Th}(E_i)$ to conclude that d is active in E_i .

It remains to be shown that d is in fact an element of $\widehat{\mathcal{D}}_{i+1}$. Assume to the contrary that $d \notin \widehat{\mathcal{D}}_{i+1}$. By maximality of $\widehat{\mathcal{D}}_{i+1}$ this means that $E_{i+1} \cup \{\beta\}$ is inconsistent, which in turn means that $\neg\beta \in \text{Th}(E_{i+1})$. Since $\beta \in \mathcal{E}$ and $\text{Th}(E_{i+1}) \subseteq \mathcal{E}$, this contradicts the fact that \mathcal{E} is consistent. \square

A.3. PROOF OF THEOREM 4.6

THEOREM 4.8. *Let \mathcal{D} be a set of closed prerequisite-free normal defaults. Then \mathcal{E} is a P-extension of the default theory with priorities $(\mathcal{W}, \mathcal{D}, <)$ iff \mathcal{E} is a B-extension of the ordered default theory $(\mathcal{W}, \mathcal{D}, <)$.*

If \mathcal{W} is inconsistent, then $\text{Th}(\mathcal{W})$ is the only P-extension and the only B-extension. In the following we assume that \mathcal{W} is consistent. Note that this means that all P-extensions and B-extensions are consistent.

To prove the theorem for consistent \mathcal{W} , we first show that any B-extension of the normal, prerequisite-free theory $(\mathcal{W}, \mathcal{D}, <)$ is also a P-extension. Let

$$d_1 = \frac{: \beta_1}{\beta_1}, \quad d_2 = \frac{: \beta_2}{\beta_2}, \quad d_3 = \frac{: \beta_3}{\beta_3}, \dots$$

be an enumeration of \mathcal{D} that is compatible with $<$, and let \mathcal{E} be the B-extension defined by this enumeration. This means that $\mathcal{E} = \bigcup_{i \geq 0} \text{Th}(F_i)$, where $F_0 := \mathcal{W}$, and for $i \geq 0$

$$F_{i+1} := \begin{cases} F_i & \text{if } \neg\beta_{i+1} \in \text{Th}(F_i), \\ F_i \cup \{\beta_{i+1}\} & \text{otherwise.} \end{cases}$$

We use \mathcal{E} to define the sets E_i as described in the definition of a P-extension (see Definition 4.3). What remains to be shown is that this iteration really yields \mathcal{E} , i.e., that $\mathcal{E} = \bigcup_{i \geq 0} \text{Th}(E_i)$. This is an immediate consequence of the following two lemmas.

LEMMA A.6. *For all $i \geq 0$ we have $E_i \subseteq \mathcal{E}$.*

Proof. The proof is by induction on i . For $i = 0$ we have $E_0 = \mathcal{W} \subseteq \mathcal{E}$. Now assume that $i \geq 0$ and consider $\beta \in E_{i+1}$. For $\beta \in E_i$ we know $\beta \in \mathcal{E}$ by induction. For $\beta \in E_{i+1} \setminus E_i$ we know (beside other things) that $:\beta/\beta \in \mathcal{D}$ and $\neg\beta \notin \mathcal{E}$. Since the default $:\beta/\beta$ occurs in the enumeration d_1, d_2, \dots , there exists $j \geq 1$ such that $\beta = \beta_j$. In addition, $\neg\beta \notin \mathcal{E}$ implies $\neg\beta \notin \text{Th}(F_{j-1})$, which yields $\beta = \beta_j \in F_j \subseteq \mathcal{E}$. \square

LEMMA A.7. *For all $i \geq 0$ we have $F_i \subseteq E_i$.*

Proof. Again, the proof is by induction on i . For $i = 0$ we have $F_0 = \mathcal{W} = E_0$. Now assume that $i \geq 0$, and consider $\beta \in F_{i+1}$. For $\beta \in F_i$, induction yields $\beta \in F_i \subseteq E_i \subseteq E_{i+1}$.

On the other hand, $\beta \in F_{i+1} \setminus F_i$ means that $\beta = \beta_{i+1}$. Since $\beta = \beta_{i+1}$, we know that $\beta \in \mathcal{E}$, and thus $\neg\beta \notin \mathcal{E}$. In order to get $\beta \in E_{i+1}$, it remains to be shown that all $d' < d_{i+1}$ are not active in E_i . Recall that a default $d' = : \beta'/\beta'$ is active in E_i iff $\beta' \notin \text{Th}(E_i)$ and $\neg\beta' \notin \text{Th}(E_i)$.

From $d' < d_{i+1}$ we can deduce that $d' = d_j$ for some $j < i + 1$. If $\beta_j \in F_j$ then $F_j \subseteq E_j$ (by induction) and $E_j \subseteq E_i$ (because $j \leq i$) yield $\beta_j \in E_i$, which implies that $d' = : \beta_j/\beta_j$ is not active in E_i . Finally, assume that $\beta_j \notin F_j$. This

means that $\neg\beta_j \in \text{Th}(E_{j-1})$, and thus $F_{j-1} \subseteq E_{j-1} \subseteq E_i$ (by induction and $j-1 < i$) yields $\neg\beta_j \in \text{Th}(E_i)$. Again, this implies that $d' = : \beta_j/\beta_j$ is not active in E_i . \square

This concludes the proof that any B-extension is a P-extension. For the other direction, assume that $\mathcal{E} = \bigcup_{i \geq 0} \text{Th}(E_i)$ is a P-extension of the normal, prerequisite-free theory $(\mathcal{W}, \mathcal{D}, <)$, obtained by iteratively generating the sets E_i as described in Definition 4.3. The first lemma says that \mathcal{E} satisfies a property that is obviously true for B-extensions but is not so trivial for P-extensions.

LEMMA A.8. *For $d = : \beta/\beta$ we have either $\beta \in \mathcal{E}$ or $\neg\beta \in \mathcal{E}$.*

Proof. The proof is by induction on $<$. Assume that $\neg\beta$ is not in \mathcal{E} . Thus the only reason for d not to fire in the i th step of the iteration can be that there exists a default $d' = : \beta'/\beta' < d$ that is active in E_{i-1} . By induction, we know that $\beta' \in \mathcal{E}$ or $\neg\beta' \in \mathcal{E}$, and hence there exists an index j such that $\beta' \in \text{Th}(E_j)$ or $\neg\beta' \in \text{Th}(E_j)$. Thus we have seen that for all $d' < d$ there exists an index j such that d' is no longer active after step j of the iteration. Since there exist only finitely many defaults smaller than d , this means that d will eventually fire, which shows that $\beta \in \mathcal{E}$. \square

To show that \mathcal{E} is a B-extension, we shall define a strict partial ordering \ll that extends $<$, and show that any enumeration of \mathcal{D} that is compatible with \ll yields \mathcal{E} as B-extension. The main idea behind the definition of \ll is that we must prevent the consequents β of defaults $d = : \beta/\beta$ with $\beta \notin \mathcal{E}$ from being added to the B-extension. For this reason, we will make sure that there are defaults $d_1 = : \beta_1/\beta_1 \ll d, \dots, d_n = : \beta_n/\beta_n \ll d$ such that $\beta_1, \dots, \beta_n \in \mathcal{E}$ and $\neg\beta \in \text{Th}(\mathcal{W} \cup \{\beta_1, \dots, \beta_n\})$.

The main technical problem will be to show that \ll has the finiteness property, in other words, that for all $d \in \mathcal{D}$ the set $\{d' \in \mathcal{D} \mid d' \ll d\}$ is finite. This property is necessary because otherwise there would not exist an enumeration of \mathcal{D} that is compatible with \ll .

For $i \geq 1$, we define \mathcal{D}_i as the set of all defaults that fire at step i of the iteration:

$$\mathcal{D}_i := \{d = : \beta/\beta \notin \text{Th}(E_{i-1}), \neg\beta \notin \mathcal{E}, \\ \text{and all } d' < d \text{ are not active in } E_{i-1}\}.$$

Note that $d = : \beta/\beta \in \mathcal{D}_i$ implies $\beta \in \mathcal{E}_i$, and thus $\beta \in \mathcal{E}$. Let $\mathcal{D}_{\text{fired}} := \bigcup_{i \geq 1} \mathcal{D}_i$ and $\mathcal{D}_{\text{out}} := \{d = : \beta/\beta \mid \beta \notin \mathcal{E}\}$. Obviously, $\mathcal{D}_{\text{fired}}$ and \mathcal{D}_{out} are disjoint, but note that their union can be a strict subset of \mathcal{D} .

LEMMA A.9. *For all $d = : \beta/\beta \in \mathcal{D}_{\text{out}}$ there exist a nonnegative integer $i(d)$ and a finite set of defaults $M(d) \subseteq \mathcal{D}_{\text{fired}}$ such that one of the following two properties holds:*

- (1) $i(d) = 0$, $M(d) = \emptyset$, and $\neg\beta \in \text{Th}(\mathcal{W})$.
- (2) $i(d) \geq 1$, $M(d) \subseteq \bigcup_{i \leq i(d)} \mathcal{D}_i$, $\neg\beta \notin \text{Th}(E_{i(d)-1})$, and $\neg\beta \in \text{Th}(\mathcal{W} \cup \text{Con}(M(d)))$.

Proof. Since $d \in \mathcal{D}_{\text{out}}$, we know (by Lemma A.8) that $\neg\beta \in \mathcal{E} = \bigcup_{i \geq 0} \text{Th}(E_i)$. We define $i(d)$ to be the smallest i such that $\neg\beta \in \text{Th}(E_i)$. If $i(d) = 0$, we define $M(d) := \emptyset$. Now $\neg\beta \in \text{Th}(\mathcal{W})$ is satisfied because $E_0 = \mathcal{W}$.

Assume that $i(d) > 0$. Minimality of $i(d)$ yields $\neg\beta \notin \text{Th}(E_{i(d)-1})$. From $\neg\beta \in \text{Th}(E_{i(d)})$ we can deduce that there exist defaults $d_1, \dots, d_n \in \bigcup_{i \leq i(d)} \mathcal{D}_i$ such that $\neg\beta \in \text{Th}(\mathcal{W} \cup \text{Con}(\{d_1, \dots, d_n\}))$. Thus we define $M(d) := \{d_1, \dots, d_n\}$. \square

Now we are ready to define the extension of $<$ we are looking for. Let \ll be the transitive closure of the relation $< \cup \prec$, where \prec is defined by

$$d' \prec d \text{ iff } d' \in \mathcal{D}_{\text{fired}}, d \in \mathcal{D}_{\text{out}}, \text{ and } d' \in M(d).$$

Obviously, \ll is an extension of $<$, i.e., $d < d'$ implies $d \ll d'$. It remains to be shown that \ll is appropriate for our purposes.

LEMMA A.10. *The relation \ll satisfies the following properties:*

- (1) \ll Noetherian; that is, there does not exist an infinitely descending chain $d_1 \gg d_2 \gg \dots$.
- (2) \ll is a partial ordering.
- (3) \ll satisfies the finiteness property; that is, for all $d \in \mathcal{D}$ the set $\{d' \in \mathcal{D} \mid d' \ll d\}$ is finite.

Proof. (1) First note that $<$ and \prec are strict partial orderings satisfying the finiteness property. For $<$, this is just the condition that a partial ordering has to satisfy to be admissible in a theory $(\mathcal{W}, \mathcal{D}, <)$. For \prec , transitivity and irreflexivity follow from disjointness of $\mathcal{D}_{\text{fired}}$ and \mathcal{D}_{out} . In fact, $\mathcal{D}_{\text{fired}} \cap \mathcal{D}_{\text{out}} = \emptyset$ implies that neither the situation $d \prec d$ nor the situation $d \prec d' \prec d''$ can occur. The finiteness property for \prec is now an immediate consequence of the fact that $M(d)$ is finite for all $d \in \mathcal{D}_{\text{out}}$.

Since $<$ and \prec are irreflexive and satisfy the finiteness property, they are Noetherian as well. For this reason, an infinitely descending chain for \ll must be (without loss of generality) of the form

$$d_1 \succ d'_1 > d_2 \succ d'_2 > \dots d_j \succ d'_j > d_{j+1} \succ \dots$$

For $j \geq 1$ we have $d_j \in \mathcal{D}_{\text{out}}$ and $d'_j \in \mathcal{D}_{\text{fired}}$. To prove that such a chain cannot be infinite, we show that $i(d_j)$ is larger than $i(d_{j+1})$.

Let $i_0 = i(d_j)$. Since $d_j \succ d'_j$, we know that $d'_j \in M(d_j)$, and thus $d'_j \in \bigcup_{i \leq i_0} \mathcal{D}_i$. Let $i_1 \leq i_0$ be such that $d'_j \in \mathcal{D}_{i_1}$. Because $d_{j+1} < d'_j$, we thus know that d_{j+1} cannot be active in E_{i_1-1} . But this means that the justification of d_{j+1} , say β_{j+1} , is already inconsistent with E_{i_1-1} , in other word, $\neg\beta_{j+1} \in \text{Th}(E_{i_1-1})$. This shows that $i(d_{j+1})$ is smaller than i_1 , and thus smaller than $i_0 = i(d_j)$.

(2) The relation \ll is transitive by definition. Irreflexivity follows from the fact that \ll is Noetherian.

(3) The finiteness property of \ll is shown by Noetherian induction on \ll . For $d \in \mathcal{D}$ we consider the sets $\{d' \in \mathcal{D} \mid d' < d\}$ and $\{d'' \in \mathcal{D} \mid d'' \prec d\}$. Both sets are finite, since $<$ and \prec satisfy the finiteness property. The \ll -successors of d are the elements of these sets, together with their \ll -successors. By induction, we know that the elements of these sets have only finitely many \ll -successors, which completes the proof of the lemma. \square

Now let $d_1 = : \beta_1/\beta_1$, $d_2 = : \beta_2/\beta_2, \dots$ be an enumeration of \mathcal{D} that is compatible with \ll . The finiteness property for \ll guarantees that such an enumeration exists. Let $\mathcal{F} = \bigcup_{i \geq 0} \text{Th}(F_i)$ be the B-extension of $(\mathcal{W}, \mathcal{D}, \ll)$ defined by this enumeration. Since \ll extends $<$, \mathcal{F} is an extension of $(\mathcal{W}, \mathcal{D}, <)$ as well. It remains to be shown that $\mathcal{F} = \mathcal{E}$.

LEMMA A.11. *For all $i \geq 0$ we have $F_i \subseteq \mathcal{E}$.*

Proof. The proof is by induction on i . For $i = 0$, $F_0 = \mathcal{W} \subseteq \mathcal{E}$. Now let $i > 0$. If $F_i = F_{i-1}$, we have $F_i = F_{i-1} \subseteq \mathcal{E}$ by induction.

Now assume that $F_i = F_{i-1} \cup \{\beta_i\}$, but $\beta_i \notin \mathcal{E}$. Thus $d_i \in \mathcal{D}_{\text{out}}$, and there exists a finite set of defaults $M(d_i)$ such that $\neg\beta_i \in \text{Th}(\mathcal{W} \cup \text{Con}(M(d_i)))$. For $d = : \beta/\beta \in M(d_i)$ we have $d \ll d_i$ and $\beta \in \mathcal{E}$. Since the enumeration d_1, d_2, \dots is compatible with \ll , this means that $\beta \in F_{i-1}$ or $\neg\beta \in F_{i-1}$. The second case cannot occur because this would mean that $\neg\beta \in \mathcal{E}$, by induction.

To sum up, we have seen that the consequents of all defaults in $M(d)$ are in F_{i-1} . But this shows that $\neg\beta_i \in \text{Th}(F_{i-1})$, which contradicts our assumption that $F_i = F_{i-1} \cup \{\beta_i\}$. \square

The next lemma completes the proof of Theorem 4.8.

LEMMA A.12. $\mathcal{E} = \mathcal{F}$.

Proof. Because of the previous lemma, we know that $\mathcal{F} \subseteq \mathcal{E}$. If $\mathcal{F} \neq \mathcal{E}$, then there exists a default $d = : \beta/\beta \in \mathcal{D}$ such that $\beta \in \mathcal{E} \setminus \mathcal{F}$. But $\beta \notin \mathcal{F}$ implies $\neg\beta \in \mathcal{F}$, which together with $\mathcal{F} \subseteq \mathcal{E}$ contradicts the fact that \mathcal{E} is consistent. \square

A.4. SOUNDNESS AND COMPLETENESS OF ALGORITHM 5.1

ALGORITHM 5.1. *Let $(\mathcal{W}, \mathcal{D}, <)$ be a closed default theory with priorities. If \mathcal{W} is inconsistent, then $\text{Th}(\mathcal{W})$ is the only P-extension. Otherwise we define $E_0 := \mathcal{W}$ and $J_0 := \emptyset$.*

Now assume that E_i ($i \geq 0$) is already defined. Consider

$$\mathcal{D}_{i+1} := \{d \in \mathcal{D} \mid d \text{ is active in } E_i \text{ and no } d' < d \text{ is active in } E_i\},$$

and choose a nonempty subset $\widehat{\mathcal{D}}_{i+1}$ of \mathcal{D}_{i+1} that satisfies

$$\neg\beta \notin \text{Th}(E_i \cup \text{Con}(\widehat{\mathcal{D}}_{i+1}) \cup J_i \cup \neg\text{Jus}(\mathcal{D}_{i+1} \setminus \widehat{\mathcal{D}}_{i+1}))$$

for all $\beta \in \text{Jus}(\widehat{\mathcal{D}}_{i+1})$.

If there is no such set, then $E_{i+1} := E_i$, $J_{i+1} := J_i$. Otherwise each choice yields new sets $E_{i+1} := E_i \cup \text{Con}(\widehat{\mathcal{D}}_{i+1})$ and $J_{i+1} := J_i \cup \neg\text{Jus}(\mathcal{D}_{i+1} \setminus \widehat{\mathcal{D}}_{i+1})$.

The set $\mathcal{E} := \bigcup_{i \geq 0} \text{Th}(E_i)$ is a P-extension iff

- (1) for all $d = \alpha : \beta/\gamma \in \bigcup_{i \geq 1} \widehat{\mathcal{D}}_i$ we have $\neg\beta \notin \mathcal{E}$, and
- (2) for all $\neg\beta \in \bigcup_{i \geq 1} J_i$ we have $\neg\beta \in \mathcal{E}$.

If \mathcal{W} is inconsistent, $\text{Th}(\mathcal{W})$ is the only P-extension, and this is what the algorithm yields in this case. Thus we may assume without loss of generality that \mathcal{W} is consistent.

To prove *soundness*, assume that $\mathcal{E} = \bigcup_{i \geq 0} \text{Th}(E_i)$ where the E_i are obtained as described in the algorithm, and that \mathcal{E} is accepted as admissible output because it satisfies the two conditions that are checked at the end of the algorithm.

To show that \mathcal{E} is a P-extension, we use it to generate sets F_i as described in the definition of P-extensions. This means that we define $F_0 := \mathcal{W}$, and for all $i \geq 0$

$$F_{i+1} := F_i \cup \{\gamma \mid \exists d \in \mathcal{D}: d = \alpha : \beta/\gamma, \alpha \in \text{Th}(F_i), \neg\beta \notin \mathcal{E}, \text{ and all } d' < d \text{ are not active in } F_i\}.$$

Now \mathcal{E} is a P-extension if $\mathcal{E} = \bigcup_{i \geq 0} \text{Th}(F_i)$. This is shown by proving, for all $i \geq 0$, that $\text{Th}(F_i) = \text{Th}(E_i)$. We proceed by induction on i . For $i = 0$, we have $F_0 = \mathcal{W} = E_0$. Now assume that $\text{Th}(F_i) = \text{Th}(E_i)$ is already known.

LEMMA A.13. $E_{i+1} \subseteq \text{Th}(F_{i+1})$.

Proof. Let γ be an element of E_{i+1} . If $\gamma \in \text{Th}(E_i)$, we have $\gamma \in \text{Th}(F_i) \subseteq \text{Th}(F_{i+1})$ by induction. Thus assume that $\gamma \in E_{i+1} \setminus \text{Th}(E_i)$. This means that there exists a default $d = \alpha : \beta/\gamma \in \widehat{\mathcal{D}}_{i+1}$ that is the reason for γ being in E_{i+1} .

By definition of \mathcal{D}_{i+1} we know that $\alpha \in \text{Th}(E_i) = \text{Th}(F_i)$, and that no default $d' < d$ is active in E_i . Since $\text{Th}(F_i) = \text{Th}(E_i)$, this means that $d' < d$ is not active in F_i as well. To get $\gamma \in F_{i+1}$, and thus $\gamma \in \text{Th}(F_{i+1})$, it remains to be shown that $\neg\beta \notin \mathcal{E}$. But this is the case because $d \in \bigcup_{j \geq 1} \widehat{\mathcal{D}}_j$, and \mathcal{E} satisfies the first condition checked at the end of the algorithm. \square

LEMMA A.14. $F_{i+1} \subseteq \text{Th}(E_{i+1})$.

Proof. Assume that $\gamma \in F_{i+1}$. By induction, the case where $\gamma \in \text{Th}(F_i)$ is again trivial. For $\gamma \in F_{i+1} \setminus \text{Th}(F_i)$ we know that there exists a default $d = \alpha : \beta/\gamma$ such that $\alpha \in \text{Th}(F_i) = \text{Th}(E_i)$, $\neg\beta \notin \mathcal{E}$, and all $d' < d$ are not active in F_i . Obviously, $\neg\beta \notin \mathcal{E}$ implies $\neg\beta \notin \text{Th}(E_i)$, and thus $\text{Th}(F_i) = \text{Th}(E_i)$ yields $d \in \mathcal{D}_{i+1}$.

In order to get $\gamma \in E_{i+1}$ it remains to be shown that $d \in \widehat{\mathcal{D}}_{i+1}$. Assume that $d \in \mathcal{D}_{i+1} \setminus \widehat{\mathcal{D}}_{i+1}$. Then $\neg\beta \in J_{i+1}$, and thus the second condition checked at the end of the algorithm yields $\neg\beta \in \mathcal{E}$. But we already know that $\neg\beta \notin \mathcal{E}$. \square

This completes the proof of soundness. To show *completeness*, we assume that $\mathcal{E} = \bigcup_{i \geq 0} \text{Th}(F_i)$ is a P-extension obtained from the sequence F_0, F_1, \dots as described in the definition of P-extensions. We use \mathcal{E} and the F_i to choose the right sets $\widehat{\mathcal{D}}_i$ in the algorithm.

Assume that $\widehat{\mathcal{D}}_i$ and the corresponding sets E_i, J_i are already defined and that these sets satisfy $\text{Th}(E_i) = \text{Th}(F_i)$ and $J_i \subseteq \mathcal{E}$. Note that for $i = 0$ this is trivially satisfied, since $E_0 = \mathcal{W}$ and $J_0 = \emptyset$. We define

$$\widehat{\mathcal{D}}_{i+1} := \{d = \alpha : \beta/\gamma \mid d \in \mathcal{D}_{i+1} \text{ and } \neg\beta \notin \mathcal{E}\},$$

where $\mathcal{D}_{i+1} = \{d \mid d \text{ is active in } E_i \text{ and no } d' < d \text{ is active in } E_i\}$. Obviously, $\widehat{\mathcal{D}}_{i+1}$ is a subset of \mathcal{D}_{i+1} .

If $\widehat{\mathcal{D}}_{i+1}$ is empty, then it is easy to see (using the induction hypothesis $\text{Th}(E_i) = \text{Th}(F_i)$) that $\text{Th}(F_i) = \text{Th}(F_{i+1})$, and thus $\mathcal{E} = \text{Th}(F_i)$. To get $E_{i+1} = E_i$, and thus $\text{Th}(E_{i+1}) = \text{Th}(E_i) = \text{Th}(F_i) = \text{Th}(F_{i+1})$, we have to show that there cannot be a nonempty subset \mathcal{D}'_{i+1} of \mathcal{D}_{i+1} satisfying

$$\neg\beta \notin \text{Th}(E_i \cup \text{Con}(\mathcal{D}'_{i+1}) \cup J_i \cup \neg\text{Jus}(\mathcal{D}_{i+1} \setminus \mathcal{D}'_{i+1}))$$

for all $\beta \in \text{Jus}(\mathcal{D}'_{i+1})$. But if the justification β of a default $d \in \mathcal{D}_{i+1}$ satisfies this condition, $\text{Th}(E_i) = \text{Th}(F_i) = \mathcal{E}$ shows that $\neg\beta \notin \mathcal{E}$. This contradicts our assumption that $\widehat{\mathcal{D}}_{i+1}$ is empty.

Now assume that $\widehat{\mathcal{D}}_{i+1}$ is not empty. We have to show that all $\beta \in \text{Jus}(\widehat{\mathcal{D}}_{i+1})$ satisfy the condition

$$\neg\beta \notin \text{Th}(E_i \cup \text{Con}(\widehat{\mathcal{D}}_{i+1}) \cup J_i \cup \neg\text{Jus}(\mathcal{D}_{i+1} \setminus \widehat{\mathcal{D}}_{i+1})).$$

This is an immediate consequence of the next lemma, since $\beta \in \text{Jus}(\widehat{\mathcal{D}}_{i+1})$ satisfies $\neg\beta \notin \mathcal{E}$ (by definition of $\widehat{\mathcal{D}}_{i+1}$).

LEMMA A.15. $\text{Th}(E_i \cup \text{Con}(\widehat{\mathcal{D}}_{i+1}) \cup J_i \cup \neg\text{Jus}(\mathcal{D}_{i+1} \setminus \widehat{\mathcal{D}}_{i+1})) \subseteq \mathcal{E}$.

Proof. Since the P-extension \mathcal{E} is deductively closed, it is sufficient to prove $E_i \cup \text{Con}(\widehat{\mathcal{D}}_{i+1}) \cup J_i \cup \neg\text{Jus}(\mathcal{D}_{i+1} \setminus \widehat{\mathcal{D}}_{i+1}) \subseteq \mathcal{E}$.

By induction, we know $J_i \subseteq \mathcal{E}$ and $\text{Th}(E_i) = \text{Th}(F_i) \subseteq \mathcal{E}$. For $\gamma \in \text{Con}(\widehat{\mathcal{D}}_{i+1})$ we have $\gamma \in F_{i+1} \subseteq \mathcal{E}$ by definition of $\widehat{\mathcal{D}}_{i+1}$. Finally, assume that $d = \alpha : \beta/\gamma \in \mathcal{D}_{i+1} \setminus \widehat{\mathcal{D}}_{i+1}$, that is, $\neg\beta \in \neg\text{Jus}(\mathcal{D}_{i+1} \setminus \widehat{\mathcal{D}}_{i+1})$. We have $\neg\beta \in \mathcal{E}$ because otherwise d would be in $\widehat{\mathcal{D}}_{i+1}$. \square

Thus we have seen that $\widehat{\mathcal{D}}_{i+1}$ is an admissible subset of \mathcal{D}_{i+1} . The lemma shows that $J_{i+1} := J_i \cup \neg\text{Jus}(\mathcal{D}_{i+1} \setminus \widehat{\mathcal{D}}_{i+1})$ (as defined in the algorithm) is a subset of \mathcal{E} . In addition, by definition of $\widehat{\mathcal{D}}_{i+1}$, $E_{i+1} := E_i \cup \text{Con}(\widehat{\mathcal{D}}_{i+1})$ satisfies $\text{Th}(E_{i+1}) =$

$\text{Th}(F_{i+1})$. In fact, the only difference possible between $F_{i+1} \setminus F_i$ and $E_{i+1} \setminus E_i$ is that $F_{i+1} \setminus F_i$ may contain some additional elements, which are, however, elements of $\text{Th}(F_i) = \text{Th}(E_i)$.

To sum up, we have shown by induction that in each step of the algorithm one can choose an admissible set $\widehat{\mathcal{D}}_i \subseteq \mathcal{D}_i$ such that the set $\text{Th}(E_i)$ obtained by this choice coincides with $\text{Th}(F_i)$. Thus we have $\mathcal{E} = \bigcup_{i \geq 0} \text{Th}(E_i)$. It remains to be shown that the two conditions at the end of the algorithm are satisfied.

For $d = \alpha : \beta/\gamma \in \bigcup_{i \geq 1} \widehat{\mathcal{D}}_i$ we have $\neg\beta \notin \mathcal{E}$ by definition of the sets $\widehat{\mathcal{D}}_i$. Finally, Lemma A.15 implies $\bigcup_{i \geq 1} J_i \subseteq \mathcal{E}$.

Acknowledgements

We thank Peter Patel-Schneider for interesting discussions on specificity of defaults and Bernhard Nebel and the referees for helpful comments on previous versions of this paper. This work has been supported by the German Ministry for Research and Technology (BMFT) under research contract ITW 92 01.

Notes

¹ There are some attempts to generalize this approach to structured classes, but they work in a very restricted setting, and it is not clear how to obtain more general results in this direction (see, e.g., [17]).

² The reader who is surprised that this is only taken as a default property of penguins should have a look at the cover of [15].

³ For the sake of simplicity we consider only defaults with one justification. However, our results can easily be extended to the general case of defaults with finitely many justifications.

⁴ The formulae occurring in one rule are assumed to have identical free variables.

⁵ Note that the finiteness condition “ $\{d' \in \mathcal{D} \mid d' < d\}$ is finite for every $d \in \mathcal{D}$ ” is necessary and sufficient for the existence of an enumeration compatible with $<$.

⁶ Note that the ‘default subsumption’ between *bird* and *winged* is assumed to have no influence on the priority ordering.

⁷ If $E_{i+1} = E_i$ and $J_{i+1} = J_i$, then for all $j \geq i$, $E_j = E_i$ and $J_j = J_i$.

⁸ Note that this is a consequence of the restriction to a particular type of first-order formulae $C(x)$, $D(x)$ in terminological languages.

⁹ For default theories without priorities, extensions can be characterized without reference to the iteration process, which allows for alternative ways of computing extensions (see, e.g., [1]).

References

1. Baader, F. and Hollunder, B.: Embedding defaults into terminological knowledge representation formalisms, in *Proc. 3rd Int. Conf. on Knowledge Representation and Reasoning*, Cambridge, MA, 1992.
2. Baader, F. and Hollunder, B.: How to prefer more specific defaults in terminological default logic, in *Proc. 13th Int. Joint Conf. on Artificial Intelligence*, Chambéry, France, 1993.
3. Besnard, P.: *An Introduction to Default Logic*, Symbolic Computation Series, Springer, 1989.
4. Brachman, R. J., McGuinness, D. L., Patel-Schneider, P. F., Resnick, L. A., and Borgida, A.: Living with CLASSIC: When and how to use a KL-ONE-like languages, in J. Sowa (ed.), *Principles of Semantic Networks*, Morgan Kaufmann, San Mateo, CA, 1991, pp. 401–456.

5. Brass, S.: Deduction with Supernormal Defaults, in G. Brewka, K. P. Jantke and P. H. Schmitt (eds), *Nonmonotonic and Induction Logics, 2nd Int. Workshop*, Springer LNCS 659, 1992.
6. Brewka, G.: Adding priorities and specificity to default logic, in C. MacNish, D. Pearce and L. M. Pereira (eds), *Logics in Artificial Intelligence, European Workshop, JELIA'94*, York, UK, Springer LNAI 838, 1994, pp. 247–260.
7. Brewka, G.: Preferred subtheories: An extended logical framework for default reasoning, in *Proc. 11th Int. Joint Conf. on Artificial Intelligence*, Detroit, MI, 1989.
8. Brewka, G.: *Nonmonotonic Reasoning: Logical Foundations of Commonsense*, Cambridge University Press, Cambridge, 1991.
9. Delgrande, J. P. and Jackson, W. K.: Default logic revisited, in *Proc. 2nd Int. Conf. on Knowledge Representation and Reasoning*, Cambridge, MA, 1991.
10. Junker, U. and Brewka, G.: Handling partially ordered defaults in TMS, in *Proc. 1st European Conf. on Symbolic and Quantitative Approaches for Uncertainty*, Marseille, France, 1991.
11. Kobsa, A.: The SB-ONE knowledge representation workbench, in *Preprint of the Workshop on Formal Aspects of Semantic Networks*, Two Harbours, CA, 1989.
12. Lifschitz, V.: Computing circumscription, in *Proc. 9th Int. Joint Conf. on Artificial Intelligence*, Los Angeles, CA, 1985.
13. Mays, E. and Dionne, B.: Making KR systems useful, in *Terminological Logic Users Workshop, Proc.*, KIT-Report 95, TU, Berlin, 1991, pp. 11–12.
14. MacGregor, R.: Statement of interest, in K. von Luck, B. Nebel and C. Peltason (eds), *Statement of Interest 2nd Int. Workshop on Terminological Logics*, Document D-91-13, DFKI Kaiserslautern, 1991.
15. Ginsberg, M. L. (ed.), *Readings in Nonmonotonic Reasoning*, Morgan Kaufmann, Los Altos, CA, 1987.
16. μ BACK, System presentation, in *Terminological Logic Users Workshop, Proc.*, KIT-Report 95, TU, Berlin, 1991, p. 186.
17. Padgham, L. and Nebel, B.: Combining classification and nonmonotonic inheritance reasoning: A first step, in Z. W. Ras and J. Komorowski (eds), *Methodologies for Intelligent Systems (ISMIS'93)*, North-Holland, Amsterdam, 1993.
18. Peltason, C., v. Luck, K. and Kindermann, C. (Org.): *Terminological Logic Users Workshop, Proc.*, KIT Report 95, TU, Berlin, 1991.
19. Reiter, R.: A logic for default reasoning, *Artificial Intelligence* **13**(1–2) (1980), 81–132.
20. Reiter, R. and Criscuolo, G.: On interacting defaults, in *Proc. 7th Int. Joint Conf. on Artificial Intelligence*, 1981.