



Deutsches
Forschungszentrum
für Künstliche
Intelligenz GmbH

**Research
Report**
RR-93-13

**A Semantics for Open Normal Defaults via a
Modified Preferential Approach**

Franz Baader, Karl Schlechta

April 1993

**Deutsches Forschungszentrum für Künstliche Intelligenz
GmbH**

Postfach 20 80
D-6750 Kaiserslautern, FRG
Tel.: (+49 631) 205-3211/13
Fax: (+49 631) 205-3210

Stuhlsatzenhausweg 3
D-6600 Saarbrücken 11, FRG
Tel.: (+49 681) 302-5252
Fax: (+49 681) 302-5341

Deutsches Forschungszentrum für Künstliche Intelligenz

The German Research Center for Artificial Intelligence (Deutsches Forschungszentrum für Künstliche Intelligenz, DFKI) with sites in Kaiserslautern and Saarbrücken is a non-profit organization which was founded in 1988. The shareholder companies are Atlas Elektronik, Daimler-Benz, Fraunhofer Gesellschaft, GMD, IBM, Insiders, Mannesmann-Kienzle, SEMA Group, Siemens and Siemens-Nixdorf. Research projects conducted at the DFKI are funded by the German Ministry for Research and Technology, by the shareholder companies, or by other industrial contracts.

The DFKI conducts application-oriented basic research in the field of artificial intelligence and other related subfields of computer science. The overall goal is to construct *systems with technical knowledge and common sense* which - by using AI methods - implement a problem solution for a selected application area. Currently, there are the following research areas at the DFKI:

- Intelligent Engineering Systems
- Intelligent User Interfaces
- Computer Linguistics
- Programming Systems
- Deduction and Multiagent Systems
- Document Analysis and Office Automation.

The DFKI strives at making its research results available to the scientific community. There exist many contacts to domestic and foreign research institutions, both in academy and industry. The DFKI hosts technology transfer workshops for shareholders and other interested groups in order to inform about the current state of research.

From its beginning, the DFKI has provided an attractive working environment for AI researchers from Germany and from all over the world. The goal is to have a staff of about 100 researchers at the end of the building-up phase.

Friedrich J. Wendl
Director

A Semantics for Open Normal Defaults via a Modified Preferential Approach

Franz Baader, Karl Schlechta

DFKI-RR-93-13

This work may not be copied or reproduced, in whole or in part, for any commercial purpose. Permission to copy in whole or in part without payment in fee is granted to individuals and institutions that are not for-profit, provided that all such whole or partial copies include the following: a notice that such copying is by permission of the German Research Community (DFG) and the Leibniz Institute for Formalized Knowledge (DFKI) of the Leibniz University Hannover; the name and address of the author and the title of the work; and appropriate acknowledgment of the source and technical contribution to the work. An appropriate notice of this copyright notice should appear in the copy. Copying for other than personal or research purposes is prohibited. A fee for copying should be paid to the Leibniz Institute for Formalized Knowledge.

This work has been supported by a grant from The Federal Ministry for Research and Technology (FKZ ITW-9201).

© Deutsches Forschungszentrum für Künstliche Intelligenz 1993

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Deutsches Forschungszentrum für Künstliche Intelligenz, Kaiserslautern, Federal Republic of Germany; an acknowledgement of the authors and individual contributors to the work; all applicable portions of this copyright notice. Copying, reproducing, or republishing for any other purpose shall require a licence with payment of fee to Deutsches Forschungszentrum für Künstliche Intelligenz.

A Semantics for Open Normal Defaults via a Modified Preferential Approach

Franz Baader and Karl Schlechta

German Research Center for AI (DFKI)

Stuhlsatzenhausweg 3, 6600 Saarbrücken 11, Germany

E-mail: baader@dfki.uni-sb.de, schle@dfki.uni-sb.de

Abstract

We present a new approach for handling open normal defaults that makes it possible

1. to derive existentially quantified formulae from other existentially quantified formulae by default,
2. to derive universally quantified formulae by default, and
3. to treat cardinality formulae analogously to other formulae.

This was not the case for previous approaches. Reiter uses Skolemization in his treatment of open defaults to achieve the first goal, but this has the unpleasant side-effect that logically equivalent facts may lead to different default consequences. In addition, Reiter's approach does not comply with our second requirement. Lifschitz's main motivation for his approach was to satisfy this second demand. However, to achieve this goal he has to violate the third requirement, and the first condition is also not observed.

Differing from these two previous approaches, we will not view open defaults as schemata for certain instantiated defaults. Instead they will be used to define a preference relation on models. But unlike the usual approaches to preferential semantics we shall not always take the minimal models to construct our semantics. Due to this new treatment of preference relations the resulting nonmonotonic consequence operator has "nice" proof-theoretic properties such as cumulativity.

Contents

1 Introduction	3
2 Problems of Reiter's and Lifschitz's Approaches	5
3 A Semantics for Open Normal Defaults	9
4 Examples	12
5 Proof-Theoretic Properties	19
6 Conclusion	23

Abstract

We present a new approach for handling open normal defaults that
is more expressive than the previous ones. It is based on a
1. to derive consistently qualified formulae from other
2. to derive consistently qualified formulae by default, and
3. to treat conditionally qualified formulae analogously to other formulae.
This was not the case for previous approaches. Reiter's approach
is the treatment of open defaults to achieve the first goal, but
the second requirement is not satisfied. In addition, Reiter's approach
lead to different default consequences. In addition, Lifschitz's main moti-
vation for his approach was to satisfy the second demand. However,
to achieve the goal he has to violate the third requirement, and the
first condition is also not observed.
- Differing from these two previous approaches, we will not view
open defaults as schemes for certain restricted details. Instead
they will be used to define a preference relation on models. But unlike
the usual approaches to preferential semantics we shall not always take
the minimal models to construct our semantics. Due to this new treat-
ment of preference relations the resulting nonmonotonic consequence
operator has "nice" proof-theoretic properties such as cancellativity.

1 Introduction

In many cases, default information expresses properties of (almost all) individuals of a given class. Thus “Birds normally can fly” speaks about a default property of the class of all birds, and not about the individual bird Tweety. On the technical level, in Reiter’s Default Logic [11], this means that most of the “naturally” occurring defaults are so-called *open defaults*, i.e., default rules with free variables. The standard birds-example is described by the rule

$$\frac{bird(x) : fly(x)}{fly(x)},$$

which contains the free variable x . On the other hand, Reiter’s definition of an extension—which describes the semantics of a default theory—makes sense only if considered for closed defaults.

To bridge this gap, Reiter assumes that open defaults stand for all their ground instances. The aim is to apply defaults not only to individuals that are explicitly given by ground terms, but also to those implicitly present because of the theory. Therefore, Reiter proposes to Skolemize the world description (i.e., the facts one starts with) and the consequents of defaults before building ground instances. For example, if we have a default that says that *doctors normally are rich*, and if we know that *Tom has a child that is a doctor*, we should like to conclude that *Tom has a rich child*, as long as we know nothing to the contrary. But to apply the default to Tom’s child, we need a ground term that stands for this child. This ground term is generated by Skolemization.

Unfortunately, it turns out that the Skolemization step has the very unpleasant effect that the consequences of a default theory may depend on the *syntactic* form of the world description one starts with (see [1] and Section 2 below). Another problem of Reiter’s treatment of open defaults has been pointed out by Lifschitz [8]. If one does not assume that all individuals of the universe are named by ground terms, one cannot deduce universally quantified formulae by default. For example, consider the default without prerequisite

$$\frac{: \neg\text{for-free}(x)}{\neg\text{for-free}(x)}$$

that expresses that, normally, nothing is for free. Together with a world description that does not imply the existence of a counterexample, one cannot conclude $\forall x. \neg\text{for-free}(x)$, even though $\neg\text{for-free}(t)$ follows by default for each ground term t .

To overcome this problem, Lifschitz proposes a modification of default logic in which free variables are not treated as metavariables for ground terms. In principle, he considers models of different cardinality separately, and fixes the universe for a given cardinality. Then he introduces new object constants representing all elements of the universe, and instantiates the defaults with these new constants. Since the new constants are assumed to represent *all* elements of the universe, one can now deduce some universally quantified formulae by default. However, as we shall show in Section 2, this way of proceeding also has some strange effects. In addition, it does not handle defaults with prerequisites in the same way as defaults without. From the standard birds-default, together with a world description that does not imply the existence of a counterexample, one still cannot conclude $\forall x. \text{bird}(x) \rightarrow \text{fly}(x)$ by default. For similar reasons, the problem of applying defaults to implicit individuals is not solved, i.e., in the doctors-example from above one cannot deduce that Tom has a rich child.

In the present paper we shall propose a new treatment of open defaults that tries to avoid the above mentioned problems of previous approaches. We shall restrict our attention to normal defaults, i.e., defaults where the justification and the consequent are identical. Open normal defaults will not be instantiated by ground terms or new object constants, but will be used to define a preference relation on the models of the world description. To move to a “better” (or preferred) model we will sometimes have to change the cardinality of the universe of the model, which means that we do not consider models of different cardinality separately. Unlike the usual approaches to preferential semantics we cannot always take the *best* models to construct our semantics. The reason is that we may get infinite chains of models that become better and better. Instead of dismissing such a chain because it contains no optimal model we will consider something like the limit (i.e., end segments) of the chain.

The definition of this new approach is given in Section 3, and illustrated by examples in Section 4. In Section 5 we shall show that the nonmonotonic consequence operator we get has “nice” proof-theoretic properties such as cumulativity.

2 Problems of Reiter's and Lifschitz's Approaches

For the convenience of the reader, we repeat the basic definitions of Reiter's default logic. A *default rule* is any expression of the form

$$\frac{\alpha : \beta_1, \dots, \beta_n}{\gamma},$$

where α , β_i , γ are first-order formulae. Here α is called the *prerequisite* of the rule, β_1, \dots, β_n are its *justifications*, and γ is its *consequent*. If the prerequisite of a default $\alpha : \beta_1, \dots, \beta_n / \gamma$ is trivially true, i.e., if the universal closure of α is valid, we often write it as $\beta_1, \dots, \beta_n / \gamma$, and call it a default without prerequisite. A default rule is *closed* iff α , β_1, \dots, β_n , γ do not contain free variables. It is *normal* if it has only one justification, and if this justification coincides with the consequent.

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, i.e., 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*, which is a set of closed first-order formulae defined by the following fixed point construction.

Definition 2.1 (Reiter extensions) Let $(\mathcal{W}, \mathcal{D})$ be a closed default theory. We define an operator Γ from sets of closed formulae to sets of closed formulae as follows. For a set \mathcal{F} of closed formulae, let $\Gamma(\mathcal{F})$ be the smallest set \mathcal{F}' of closed formulae satisfying the conditions

1. $\mathcal{W} \subseteq \mathcal{F}'$,
2. \mathcal{F}' is deductively closed, and
3. for any default $\alpha : \beta_1, \dots, \beta_n / \gamma \in \mathcal{D}$, if $\alpha \in \mathcal{F}'$ and $\neg\beta_1, \dots, \neg\beta_n \notin \mathcal{F}'$ then $\gamma \in \mathcal{F}'$.

Then \mathcal{F} is an extension of $(\mathcal{W}, \mathcal{D})$ iff $\mathcal{F} = \Gamma(\mathcal{F})$.

In general, a closed default theory may have more than one extension, or even no extension. Depending on whether one wants to employ skeptical or credulous reasoning, a closed formula δ is a *consequence of a closed default theory* iff it is in all extensions or if it is in at least one extension of the theory. In the following we shall restrict our attention to skeptical reasoning.

Reiter defines *extensions of arbitrary default theories* $(\mathcal{W}, \mathcal{D})$, i.e., default theories with open defaults, as follows. First, the formulae of \mathcal{W} and the consequents of the defaults are Skolemized (see [11], Section 7). Second, a set \mathcal{D}' of closed default rules is generated by taking all ground instances (over the initial signature together with the newly introduced Skolem functions) of defaults of \mathcal{D} (with the consequents in Skolemized form). Now \mathcal{E} is an *extension* of $(\mathcal{W}, \mathcal{D})$ iff \mathcal{E} is an extension of the closed default theory $(\mathcal{W}', \mathcal{D}')$, where \mathcal{W}' is the Skolemized form of \mathcal{W} . Let us first give an example that motivates the Skolemization step.

Example 2.2 The world description \mathcal{W} consists of the formula $\exists x. P(x)$, \mathcal{D} consists of the default $P(x) : Q(x)/Q(x)$, and the language contains only one constant symbol c and no other function symbols.

If we do not Skolemize before instantiating the default, there is only one instance of the default, and its prerequisite $P(c)$ cannot be deduced. Thus we have only one extension, $Th(\mathcal{W})$, and we cannot deduce $\exists x. Q(x)$ by default.

However, if the world description is Skolemized, we get a new constant b and the fact $P(b)$. Now our default can fire for this constant, we can deduce $Q(b)$, and thus also $\exists x. Q(x)$, by default. \square

In the above example, Skolemization is necessary to get the intuitive consequence. But in general, the Skolemization step is problematic. The next example shows that, due to Skolemization, the consequences of a default theory depend on the *syntactic* form of the world description, i.e., for identical sets of open defaults, logically equivalent world descriptions may lead to different results.

Example 2.3 We consider the two world descriptions $\mathcal{W}_1 := \{\exists x. P(x) \wedge Q(x)\}$ and $\mathcal{W}_2 := \mathcal{W}_1 \cup \{\exists x. P(x)\}$. Obviously, these two sets of formulae are logically equivalent. When Skolemizing \mathcal{W}_1 , we get a single new Skolem constant b and the fact $P(b) \wedge Q(b)$, whereas when Skolemizing \mathcal{W}_2 we get two Skolem constants c and d , and the facts $P(c) \wedge Q(c)$ and $P(d)$.

Now consider the (open) default $P(x) : \neg Q(x)/\neg Q(x)$. For the Skolemized version of \mathcal{W}_1 , this default is instantiated with b , whereas for the Skolemized

version of \mathcal{W}_2 it is instantiated with c and d . Obviously, the default rule *cannot* fire for b and c , because their being in P and Q is inconsistent with its justification. On the other hand, this default rule *can* be applied to d , because being in P is consistent with being in $\neg Q$. For this reason, d is put into $\neg Q$, which shows that the Skolemized version of \mathcal{W}_2 has $\exists x. \neg Q(x)$ as a default consequence, whereas this fact cannot be deduced by default from the Skolemized version of \mathcal{W}_1 . \square

Lifschitz's proposal [8] for a semantics for open defaults was motivated by another problem of Reiter's approach: It is not possible to derive universally quantified formulae by default.

Example 2.4 Assume that we have the single default $: Q(x)/Q(x)$. As long as there is no information to the contrary (i.e., one does not know that there is an object that is not in Q) one would like to conclude that all objects are in Q . In particular, this means that with the empty world description one would like to derive the formula $\forall x. Q(x)$.

With Reiter's approach, one can only get $Q(t)$ for all ground terms t . Since there is no domain-closure assumption this does not imply the desired universal formula $\forall x. Q(x)$. \square

To overcome this problem, Lifschitz does not instantiate defaults by ground terms, but by new "parameters," of which he assumes that they cover the whole universe, considering models of different cardinality separately. For a given cardinality, he defines the notion of an extension by a fixed point definition, which is a model-theoretic variant of Reiter's proof-theoretic definition.

More formally, let $(\mathcal{W}, \mathcal{D})$ be a (possibly open) default theory, and let \mathcal{U} be a nonempty set (the universe). Lifschitz extends the language by new object constants representing all elements of \mathcal{U} . For a set \mathcal{V} of models of \mathcal{W} with universe \mathcal{U} , $Th^*(\mathcal{V})$ denotes the closed formulae (of the extended language) that are true in all elements of \mathcal{V} . The restriction of $Th^*(\mathcal{V})$ to the original language is then denoted by $Th(\mathcal{V})$.

Definition 2.5 (\mathcal{U} -extensions) Let \mathcal{D}' be the set of closed defaults (over the extended language) that is obtained by instantiating the open defaults of \mathcal{D} with the new object constants. We define an operator Δ from sets of models of \mathcal{W} with universe \mathcal{U} to sets of models of \mathcal{W} with universe \mathcal{U} as follows. For

a set \mathcal{V} of models of \mathcal{W} with universe \mathcal{U} . let $\Delta(\mathcal{V})$ be the largest set \mathcal{V}' of models of \mathcal{W} with universe \mathcal{U} satisfying

- For all defaults $\alpha : \beta_1, \dots, \beta_n / \gamma \in \mathcal{D}'$, if $\alpha \in Th^*(\mathcal{V}')$ and $\neg\beta_1, \dots, \neg\beta_n \notin Th^*(\mathcal{V})$ then $\gamma \in Th^*(\mathcal{V}')$.

Then $Th(\mathcal{V})$ is an \mathcal{U} -extension of $(\mathcal{W}, \mathcal{D})$ iff $\mathcal{V} = \Delta(\mathcal{V})$.

Default consequences are then defined in a skeptical way by taking the intersection of all \mathcal{U} -extensions for all nonempty sets \mathcal{U} . Obviously, it is enough to consider one set \mathcal{U} for each cardinality.

In Example 2.4, Lifschitz's approach yields the formula $\forall x. Q(x)$ as a default consequence since the default can be applied for all new constants, and it is assumed that these constants cover the whole universe \mathcal{U} . If the world description contains a counterexample—e.g., $\mathcal{W} = \{\neg Q(a)\}$ —the approach still allows one to get universally quantified formulae as default consequences—here $\forall x. (x = a \vee Q(x))$.

Lifschitz shows that for normal defaults without prerequisites there is a close connection between the approach presented above and circumscription. We shall now argue that for defaults with prerequisites Lifschitz's approach is still not totally satisfactory.

One problem is that in Example 2.2, where we considered the single default $P(x) : Q(x) / Q(x)$ and the world description $\mathcal{W} = \{\exists x. P(x)\}$, one still cannot conclude $\exists x. Q(x)$ by default. The reason is that—because of the existential formula—in all models of \mathcal{W} with universe \mathcal{U} one of the new constants is in P , but for different models this may be the case for different constants. This means that we do not get $P(c) \in Th^*(\mathcal{W})$ for any of the new constants, and thus $\Delta(\mathcal{W}) = \mathcal{W}$.

This argument also shows that one does not get the universal formula $\forall x. P(x) \rightarrow Q(x)$ by default, even though the world description does not imply the existence of a counterexample. Thus the treatment of defaults without prerequisites (where $\forall x. True(x) \rightarrow Q(x)$ is a default consequence in the corresponding situation) and of defaults with prerequisites is not symmetric.

Another—sometimes unpleasant—behaviour of this approach comes from the fact that different cardinalities are considered separately. The effect is that cardinality formulae are treated differently from other formulae.

Example 2.6 Assume that \mathcal{D} consists of the (closed) normal prerequisite-free default $\text{ : } \exists_{\leq 1} / \exists_{\leq 1}$, where $\exists_{\leq 1}$ stands for the formula $\forall x, y. x = y$ expressing that the universe has only 1 element.

For the empty world description, one might expect that one can conclude $\exists_{\leq 1}$ by default. This is the case if we use Reiter's approach for closed default theories. In fact, $\exists_{\leq 1}$ can consistently be assumed since there is a model (of the empty theory) of cardinality 1. Thus the default can be applied. However, with Lifschitz's approach the default is not applicable for universes \mathcal{U} of cardinality greater than 1, and since one takes the intersection of all \mathcal{U} -extensions, $\exists_{\leq 1}$ is not concluded by default.

On the other hand, if one has a default $\text{ : } \exists_{\leq 1} P / \exists_{\leq 1} P$, where $\exists_{\leq 1} P$ stands for the formula $\forall x, y. (P(x) \wedge P(y) \rightarrow x = y)$, then this default would be applied in Lifschitz's approach. This shows that the treatment of defaults concerning the cardinality of the whole universe is not symmetric to the treatment of formulae concerning the cardinality of the extension of predicates in the language. \square

3 A Semantics for Open Normal Defaults

The restriction to normal defaults is justified by the fact that these are the most natural types of default rules. The main reason for using non-normal defaults is that they can sometimes encode priorities between (otherwise normal) defaults [12]. As pointed out by Brewka [4], it is preferable to treat priorities directly (see, e.g., [3, 6, 2]), and not by an ad hoc encoding. In the present paper we do not take priorities between open defaults into account. This will be a subject of further research.

The aim of the modified preferential semantics presented below is to overcome the problems of previous approaches that we have pointed out above. Thus the new approach should be able to handle defaults with prerequisites such that it is possible

1. to derive existentially quantified formulae from other existentially quantified formulae by default,
2. to derive universally quantified formulae by default, and
3. to treat cardinality formulae analogously to other formulae.

As we have seen in the previous section, an approach that instantiates defaults by new object constants or ground terms, without Skolemizing existential quantifiers, is problematic with respect to the first requirement. For this reason we will not consider defaults as schemata for certain instantiated defaults. Instead they will be used to define a preference relation on the models of the world description. Intuitively speaking, a model M_2 will be better than M_1 with respect to a given default d if M_2 contains “less counterexamples” to d than M_1 . To satisfy the third requirement from above, we cannot assume that the models we compare have the same universe. Nevertheless, we assume that all universes are subsets of a given class, and that universes of different models are not necessarily disjoint. In the following, we are not interested in set-theoretic niceties. Thus we shall frequently talk about sets of models even though these will most probably only be classes.

Before we can give a formal definition of the preference relation, we have to introduce some notation. An open normal default $d = \alpha : \beta / \beta$ will be written as $\alpha(\underline{x}) : \beta(\underline{x}) / \beta(\underline{x})$, where the tuple $\underline{x} = (x_1, \dots, x_n)$ consists of the free variables occurring in α, β . Now let $\underline{x} = (x_1, \dots, x_n)$ be a tuple of variables, and let ϕ be a formula having some of these variables as its free variables. Assume that M is a model with universe \mathcal{U}_M , and ν is a valuation that replaces the x_i by the elements u_i of \mathcal{U}_M . We say that $\phi(u_1, \dots, u_n)$ holds in M iff $\nu(\phi)$ is true in M .

Definition 3.1 (positive and negative examples) *Let $(\mathcal{W}, \mathcal{D})$ be a default theory, and let $d = \alpha(\underline{x}) : \beta(\underline{x}) / \beta(\underline{x}) \in \mathcal{D}$ be an open default containing n free variables. For a model M of \mathcal{W} with universe \mathcal{U}_M , a tuple $\underline{u} \in \mathcal{U}_M^n$ is a positive example for d in M iff $\alpha(\underline{u})$ and $\beta(\underline{u})$ both hold in M . It is a negative example for d in M iff $\alpha(\underline{u})$ holds in M , but $\beta(\underline{u})$ does not hold in M . The set of positive examples will be denoted by $\text{Pos}(d, M)$ and the set of negative examples by $\text{Neg}(d, M)$.*

To move to a preferred model we require that we do not lose any positive example, and that at least one negative example is changed into a positive one. However, these two conditions are not sufficient since it could still happen that, while changing one negative example into a positive one, we introduce a lot of new negative examples, by changing the universe or the interpretation of α . This would not capture the intuition behind the notion of a preferred model. For this reason we introduce a third condition that disallows the introduction of new negative examples.

Definition 3.2 (*d*-preferred models) Let $(\mathcal{W}, \mathcal{D})$ be a default theory, d be a default in \mathcal{D} , and M_1, M_2 be models of \mathcal{W} . We say that M_2 is *d*-preferred to M_1 ($M_1 \succ_d M_2$) iff the following conditions hold:

1. The set of positive examples for d in M_1 is a subset of the set of positive examples for d in M_2 (i.e., $\text{Pos}(d, M_1) \subseteq \text{Pos}(d, M_2)$).
2. At least one of the negative examples for d in M_1 is a positive example for d in M_2 (i.e., $\text{Neg}(d, M_1) \cap \text{Pos}(d, M_2) \neq \emptyset$).
3. M_2 does not contain new negative examples for d (i.e., $\text{Neg}(d, M_2) \subseteq \text{Neg}(d, M_1)$).

The fact that one may first apply a default d_1 , then a default d_2 , etc., is captured by going from a model M_1 of \mathcal{W} to a d_1 -preferred model M_2 , from there to a d_2 -preferred model M_3 , etc.

Definition 3.3 (preferred models) For a default theory $(\mathcal{W}, \mathcal{D})$, the preference relation \succ between models of \mathcal{W} is defined to be the transitive closure of the union of all relations \succ_d for $d \in \mathcal{D}$. As usual, the reflexive closure of \succ is denoted by \succeq .

In general, this preference relation will not be antisymmetric, and even if it is antisymmetric there may be infinitely decreasing chains. Thus there may exist models of \mathcal{W} that are not above a minimal model, and it may even be the case that there are no minimal models.

The usual treatment of preference relations (see, e.g., [14, 7]) is to consider just the minimal models. Thus a chain that does not end in a minimal model is completely disregarded by this kind of semantics, and the non-existence of minimal models is treated as an inconsistency. In Section 4 we shall give examples that demonstrate that this is not the adequate way to treat our preference relation. For this reason we shall define a modified preferential semantics that does not disregard infinitely decreasing chains. Instead it takes something like their limit by considering all end segments.

Definition 3.4 (dense sets) Let $>$ be a transitive relation on a set \mathcal{S} , and let \geq denote its reflexive closure. Let \mathcal{A} be a subset of \mathcal{S} .

1. \mathcal{A} is $>$ -complete in \mathcal{S} iff for all $s \in \mathcal{S}$ there exists $a \in \mathcal{A}$ such that $s \geq a$.

2. \mathcal{A} is \succ -closed in \mathcal{S} iff for all $a \in \mathcal{A}$ and $s \in \mathcal{S}$, $a \succ s$ implies $s \in \mathcal{A}$.

3. \mathcal{A} is \succ -dense in \mathcal{S} iff it is both \succ -complete and \succ -closed in \mathcal{S} .

If it is clear from the context we will often omit the prefix \succ .

For example, if \mathcal{S} is just an infinitely decreasing chain with respect to \succ , then dense sets are exactly the end segments of the chain.

The rôle of minimal models in normal preferential semantics is now taken over by dense sets.

Definition 3.5 (modified preferential semantics) *Let $(\mathcal{W}, \mathcal{D})$ be a normal default theory, and let \succ be the preference relation on models of \mathcal{W} defined by this theory. We say that a closed formula ϕ is a default consequence of $(\mathcal{W}, \mathcal{D})$ iff there is a \succ -dense set in the set of all models of \mathcal{W} such that ϕ is true in all elements of this dense set.*

If every model of \mathcal{W} lies above some minimal model¹ then the set of minimal models is dense. Moreover, any dense set must contain all minimal models. This shows that in this case our modified preferential semantics coincides with the usual approach of taking minimal models.

The definition of a preferential semantics based on dense sets makes sure that the nonmonotonic consequence relation has nice proof-theoretic properties. For example, one can only deduce contradictory information if \mathcal{W} was already inconsistent. This and other properties of the consequence relation will be considered in Section 5. But before proving these abstract properties we shall give concrete examples in the next section.

4 Examples

The purpose of this section is to give examples that motivate the way our preferential semantics was defined. In particular, they will demonstrate that the three requirements stated at the beginning of the previous section are met by our approach. In the following, P, Q will always be unary predicate symbols, x, y variable symbols, and b, c constant symbols.

The first example shows that our approach can be considered as a skeptical one: from contradictory defaults we do not draw default conclusions.

¹This property is called smooth in [7] or stoppered in [10].

Example 4.1 [contradictory defaults]

defaults: $d_1 = : P(x)/P(x)$, $d_2 = : \neg P(x)/\neg P(x)$.

world description: empty.

The only consequences we get are the classical consequences of the (empty) world description, i.e., the valid closed formulae. The reason is that the only dense set is the set of all models of $\mathcal{W} = \emptyset$ itself. This is an easy consequence of the definition of dense sets and the fact that $M_1 \succ M_2$ implies $M_2 \succ M_1$. To show this fact it is enough to prove that $M_1 \succ_{d_1} M_2$ implies $M_2 \succ M_1$ (the claim follows by symmetry and induction). Thus assume that $M_1 \succ_{d_1} M_2$.

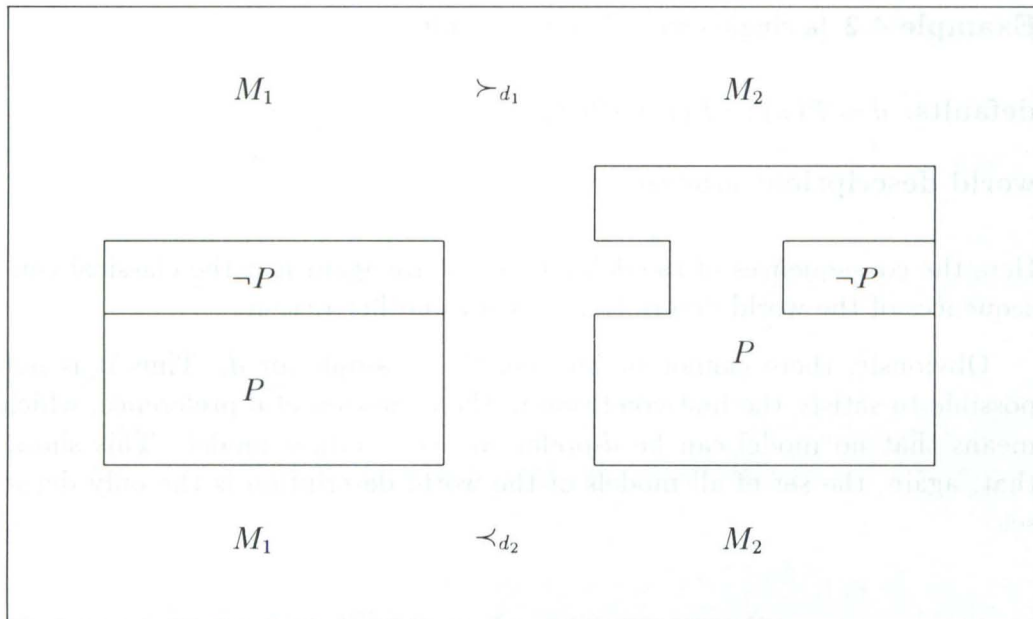


Figure 1: $M_1 \succ_{d_1} M_2$ implies $M_2 \succ M_1$

As illustrated in Figure 1, this means that at least one of the negative examples for d_1 in M_1 (i.e., elements of $\neg P$) is a positive example for d_1 in M_2 (i.e., an element of P). Some of the negative examples for d_1 in M_1 may no longer be in the universe of M_2 , but we cannot lose positive examples. In addition, the universe of M_2 may contain new elements. These new elements must not be negative examples for d_1 in M_2 , which means that they all have to be elements of P .

Now it is easy to see that $M_2 \succ_{d_2} M_1$. In fact, the elements of $\neg P$ are positive examples for d_2 in M_2 , and the elements of P are negative examples

for d_2 in M_2 . One can regain M_1 from M_2 by keeping the positive examples for d_2 , making some of the negative examples positive for d_2 , adding some new positive examples for d_2 , and discarding some negative examples for d_2 . Because in M_2 the predicate P contains at least one element that was in $\neg P$ in M_1 , we can be sure that at least one negative example for d_2 in M_2 becomes positive in M_1 . Thus we have seen that $M_2 \succ_{d_2} M_1$, which yields $M_2 \succ M_1$. \square

Something very similar happens when the contradictory defaults have a prerequisite $R(x)$ that is not trivially true. The next examples considers a single contradictory default.

Example 4.2 [a single contradictory default]

defaults: $d = P(x) : \neg P(x) / \neg P(x)$.

world description: arbitrary.

Here the consequences of the default theory are again just the classical consequences of the world description, but for another reason.

Obviously, there cannot be any positive example for d . Thus it is not possible to satisfy the first condition in the definition of d -preference, which means that no model can be d -preferred over another model. This shows that, again, the set of all models of the world description is the only dense set. \square

The next four examples consider a default that says “normally all individuals are in P ” with varying world descriptions. The first of these examples demonstrates that, as in Lifschitz’s approach, one can derive universally quantified formulae by default.

Example 4.3 [universally quantified formulae by default]

defaults: $d = : P(x) / P(x)$.

world description: empty.

A model where all elements of the universe are in P is minimal with respect to \succ_d because it contains no negative examples. In addition, any model where $\neg P$ is not empty can be modified to a d -preferred model where all elements of the universe are in P . Thus the set of minimal models is dense, and any dense set must contain all minimal models. Obviously, the formula $\forall x. P(x)$ holds in all minimal models. \square

If the world description implies that there is a counterexample for the default, one can no longer derive that all elements are in P . But as in Lifschitz's approach one can deduce that all elements different from the counterexample are in P .

Example 4.4 [at least one counterexample]

defaults: $d = : P(x)/P(x)$.

world description: $\exists x. \neg P(x)$.

First, we show that models where all but one element of the universe are in P are \succ_d -minimal models. Assume to the contrary that M_1 is such a model, and that $M_1 \succ_d M_2$. Let $u \in \mathcal{U}_{M_1}$ be the only element of $\neg P$ in M_1 . Since this is the only negative example for d in M_1 , it has to be a positive example in M_2 . In addition, all the other elements of \mathcal{U}_{M_1} are positive examples for d in M_1 . Thus they also have to be positive examples for d in M_2 . Because M_2 is a model of $\exists x. \neg P(x)$, it thus contains a new negative example. This means that the third condition in the definition of d -preference is not satisfied.

Second, it is easy to see that any model where $\neg P$ has more than one element can be modified to a d -preferred model with exactly one element in $\neg P$.

Finally, the formula

$$\exists x. (\neg P(x) \wedge \forall y. (x \neq y \rightarrow P(y)))$$

holds in all minimal models. \square

This example demonstrates why the third condition in the definition of d -preference is necessary. Assume that d -preference was defined just using the first and the second condition. Then one could move from a model having exactly one element in $\neg P$ to a better one by first putting this one element

into P , and then augmenting the universe by a new element that is put into $\neg P$. It is easy to see that this way one would deduce by default all formulae saying that the universe contains at least n elements (for each cardinality n).

While these formulae are no desirable default consequences in the above example, one should like to deduce them (for finite cardinalities) in the next example.

Example 4.5 [P becomes infinite by default]

defaults: $d = : P(x)/P(x)$.

world description: For each finite cardinality n a formula saying that the existence of n elements in P implies the existence of $n + 1$ elements in the universe.

Obviously, a model of the world description where all elements of the universe are in P must be infinite. Since such a model contains no negative examples, it is \succ_d -minimal. In addition, any model in which $\neg P$ is not empty can be modified to an infinite d -preferred model where $\neg P$ is empty.

All the minimal models satisfy $\forall x. P(x)$ and, for all finite cardinalities n , all formulae saying that the universe contains at least n elements. \square

If the world description implies the existence of infinitely many elements of $\neg P$, a default that says that normally all elements are in P should allow one to derive that P is also infinite. The next example shows that our approach satisfies this requirement.

Example 4.6 [infinitely many counterexample]

defaults: $d = : P(x)/P(x)$.

world description: For each finite cardinality n a formula saying that $\neg P$ contains at least n elements.

For all finite cardinalities n , the formula saying that P has at least n elements is a default consequence of this theory. To prove this it is enough to show that the set of all models of the world description for which P is infinite is a dense set.

This set is closed because when moving to a d -preferred model one must not lose positive examples (i.e., elements of P). To see that it is complete, assume that M is a model of the world description that interprets P as a finite set. By definition of the world description, M contains infinitely many negative examples (i.e., elements of $\neg P$). Obviously, one can change infinitely many of these negative examples to positive ones, while still keeping infinitely many elements in $\neg P$. This shows that one can get from M to a d -preferred model that interprets P by an infinite set. \square

In the above example, no model can be minimal with respect to \succ_d . In fact, any model of the world description must contain infinitely many negative examples. One can always change one of these to a positive example, thus obtaining a d -preferred model. This shows that a minimal-model approach would not be appropriate in our context.

The next example demonstrates that our approach is able to deduce existentially quantified formulae by default.

Example 4.7 [existentially quantified formulae by default]

defaults: $d = P(x) : Q(x)/Q(x)$.

world description: $\exists x. P(x)$.

By an argument very similar to the one used in Example 4.3 one can show that $\forall x. (P(x) \rightarrow Q(x))$ is a consequence of this theory. Together with $\exists x. P(x)$ this yields $\exists x. Q(x)$ as a default consequence.

Recall that with Lifschitz's approach these consequences could not be obtained. \square

In the presence of a counterexample, e.g., if the world description consists of the formula $\exists x. (P(x) \wedge \neg Q(x))$, one can still deduce an appropriate universally quantified formula, namely

$$\exists x. (P(x) \wedge \neg Q(x) \wedge \forall y. (x \neq y \rightarrow (P(y) \rightarrow Q(y))))).$$

But the existential formula $\exists x. Q(x)$ can no longer be obtained. This is reasonable since P could have just one element, which then must be in $\neg Q$. However, if the world description contains an additional formula saying that P has more than one element, we again get $\exists x. Q(x)$.

It should be noted that in Example 4.7 one would get the same result if the formula $\exists x. P(x)$ was replaced by $P(b)$ for a constant b . In particular, in the presence of a counterexample, one would not conclude $Q(b)$ because b could be this counterexample. This shows that our approach does not assume that named individuals (i.e., individuals described by ground terms) are more normal than other individuals. This assumption is, however, made by approaches that instantiate defaults with ground terms.

To show that our approach does not treat cardinality formulae differently from other types of formulae, we reconsider Example 2.6.

Example 4.8 [“at most one” by default]

defaults: $d = \exists_{\leq 1} / \exists_{\leq 1}$, where $\exists_{\leq 1}$ stands for the formula $\forall x, y. x = y$ expressing that the universe has only 1 element.

world description: empty.

The default d is a closed default. Thus there is only one example, the 0-tuple,² which is either positive or negative, depending on whether the universe has cardinality 1 or greater cardinality. For this reason, any model of cardinality 1 is d -preferred to any model of cardinality greater than 1, and the models of cardinality 1 are incomparable with each other. This shows that $\exists_{\leq 1}$ is a default consequence. \square

Our final example demonstrates that not only information on the cardinality of models can be deduced by default, but also new equalities between constants. The reason is that the interpretation of constant or function symbols can be changed when moving to a preferred model.

Example 4.9 [new equalities by default]

defaults: $\exists P(x) / P(x)$.

world description: $\neg P(b) \wedge \neg P(c)$.

By an argument as in Example 4.4 one can show that the formula

$$\exists x. (\neg P(x) \wedge \forall y. (x \neq y \rightarrow P(y)))$$

is a default consequence of this theory. Together with the world description this formula implies that $b = c$. \square

²The unique function from the empty set into the universe.

This consequence seems to be reasonable since b and c are counterexamples for the default. To keep the number of counterexamples as small as possible, the best one can do is to identify b with c (as long as there is no information to the contrary), and thus have only one counterexample instead of two.

5 Proof-Theoretic Properties

The modified preferential approach we have presented in Section 3 satisfies most of the proof-theoretic properties considered by Gabbay [5], Makinson [9], and Kraus, Lehmann and Magidor [7]. It should be noted that this does not really depend on the specific preference relation \succ defined in Section 3. The main reason for the properties to hold is our use of \succ -dense sets for defining the nonmonotonic consequence relation. This provides use with all the properties satisfied by the usual preferential approaches. In addition, our approach satisfies cautious monotony, and thus cumulativity (without an additional smoothness condition) because minimal models are replaced by dense sets.

The properties are usually formulated as rules for a nonmonotonic consequence relation \vdash . In our context, this relation is defined as follows.

Definition 5.1 *Assume that we have a fixed set \mathcal{D} of default rules. For closed formulae α, β , the expression $\alpha \vdash \beta$ means that β is a consequence of the default theory $(\mathcal{D}, \{\alpha\})$.*

We shall consider the seven properties defined below:

Right weakening: If $\alpha \rightarrow \beta$ is valid then $\gamma \vdash \alpha$ implies $\gamma \vdash \beta$.

Reflexivity: $\alpha \vdash \alpha$.

Right and: $\alpha \vdash \beta$ and $\alpha \vdash \gamma$ implies $\alpha \vdash \beta \wedge \gamma$.

Left or: $\alpha \vdash \gamma$ and $\beta \vdash \gamma$ implies $\alpha \vee \beta \vdash \gamma$.

Left logical equivalence: If $\alpha \leftrightarrow \beta$ is valid then $\beta \vdash \gamma$ implies $\alpha \vdash \gamma$.

Cautious monotony: $\alpha \vdash \beta$ and $\alpha \vdash \gamma$ implies $\alpha \wedge \beta \vdash \gamma$.

Rational monotony: $\alpha \vdash \beta$ and $\alpha \not\vdash \neg\gamma$ implies $\alpha \wedge \gamma \vdash \beta$.

If the first six properties are satisfied then the nonmonotonic consequence relation is also cumulative,³ i.e., the first six properties imply the rule

Cumulativity: If $\alpha \sim \beta$ then we have $\alpha \sim \gamma$ if, and only if, $\alpha \wedge \beta \sim \gamma$.

This is a property that is not satisfied for Reiter's approach, even if all defaults are closed.

Theorem 5.2 *The nonmonotonic consequence relation defined in Section 3 satisfies the properties "right weakening," "reflexivity," "right and," "left or," "left logical equivalence," and "cautious monotony." It does not satisfy "rational monotony."*

The first property, *right weakening*, is obvious. In fact, $\gamma \sim \alpha$ means that there is a dense set \mathcal{A} in the set of all models of γ such that all elements of \mathcal{A} satisfy α . Since $\alpha \rightarrow \beta$ is valid, this formula holds in all elements of \mathcal{A} . Thus β also holds in all elements of the dense set \mathcal{A} .

The second property, *reflexivity* follows from the fact that the set of all models of α is dense in itself.

To show *right and*, we need the following property of dense sets.

Lemma 5.3 *Let $>$ be a transitive relation on a set \mathcal{S} , and let \mathcal{A}, \mathcal{B} be subsets of \mathcal{S} . If \mathcal{A} and \mathcal{B} are $>$ -dense in \mathcal{S} then so is their intersection $\mathcal{A} \cap \mathcal{B}$.*

Proof. First, we show that $\mathcal{A} \cap \mathcal{B}$ is $>$ -complete in \mathcal{S} . Thus assume that $s \in \mathcal{S}$. Because \mathcal{A} is $>$ -complete there exists $a \in \mathcal{A}$ such that $s \geq a$. Now completeness of \mathcal{B} implies that there is $b \in \mathcal{B}$ such that $a \geq b$. Transitivity of $>$ implies $s \geq b$, and since \mathcal{A} is closed we also know $b \in \mathcal{A}$.

Second, $\mathcal{A} \cap \mathcal{B}$ is obviously $>$ -closed. In fact $c \in \mathcal{A} \cap \mathcal{B}$ and $c > d$ implies $d \in \mathcal{A}$ and $d \in \mathcal{B}$ since \mathcal{A} and \mathcal{B} are $>$ -closed. \square

Now $\alpha \sim \beta$ means that there is a $>$ -dense set \mathcal{A} in the set of all models of α such that β is satisfied in all elements of \mathcal{A} . Analogously, $\alpha \sim \gamma$ means that there is a $>$ -dense set \mathcal{B} in the set of all models of α such that γ is satisfied in all elements of \mathcal{B} . But then $\beta \wedge \gamma$ is satisfied in all elements of $\mathcal{A} \cap \mathcal{B}$, and this set is $>$ -dense by the lemma. Thus we have shown that our nonmonotonic consequence relation satisfies the "right and" property.

³See [7], Lemma 19 and 20 for a proof.

This property implies that one can only conclude contradictory information by default if the precondition is already inconsistent. In fact, assume that β and γ are contradictory, i.e., $\beta \wedge \gamma$ is inconsistent. From $\alpha \vdash \beta$ and $\alpha \vdash \gamma$ one can conclude $\alpha \vdash \beta \wedge \gamma$, by “right and.” But this means that there exists a dense set \mathcal{A} in the set \mathcal{S} of all models of α such that the (inconsistent) formula $\beta \wedge \gamma$ holds in all elements of \mathcal{A} . Obviously, this is only possible if \mathcal{A} is empty. But then \mathcal{S} must be empty as well since otherwise the empty set \mathcal{A} could not be complete in \mathcal{S} .

The *left or* property is an easy consequence of the next lemma.

Lemma 5.4 *Let $>$ be a transitive relation on a set \mathcal{S} , and let $\mathcal{S}_1, \mathcal{S}_2$ be subsets of \mathcal{S} . If \mathcal{A}_1 is $>$ -dense in \mathcal{S}_1 and \mathcal{A}_2 is $>$ -dense in \mathcal{S}_2 then there exists a subset \mathcal{B} of $\mathcal{A}_1 \cup \mathcal{A}_2$ such that \mathcal{B} is dense in $\mathcal{S}_1 \cup \mathcal{S}_2$.*

Proof. We define $\mathcal{B} := \mathcal{A}'_1 \cup \mathcal{A}'_2$ where the sets \mathcal{A}'_i ($i = 1, 2$) are defined as

$$\mathcal{A}'_i := \mathcal{A}_i \setminus \{a \in \mathcal{A}_i \mid \text{there is } b \in \mathcal{S}_j \setminus \mathcal{A}_j \text{ (} j \neq i \text{) with } a > b\}.$$

First, we show that \mathcal{B} is $>$ -complete in $\mathcal{S}_1 \cup \mathcal{S}_2$. Thus assume that $a \in \mathcal{S}_1 \cup \mathcal{S}_2$. We consider the case where $a \in \mathcal{S}_1$. (The other case is symmetric.) Since \mathcal{A}_1 is complete in \mathcal{S}_1 there exists an element a_1 of \mathcal{A}_1 such that $a \geq a_1$. If $a_1 \in \mathcal{A}'_1$ then we are done.

Otherwise, there exists $b \in \mathcal{S}_2 \setminus \mathcal{A}_2$ such that $a_1 > b$. Since \mathcal{A}_2 is complete in \mathcal{S}_2 we know that there is an element b_2 of \mathcal{A}_2 with $b \geq b_2$. Because $>$ is transitive we get $a > b_2$. It remains to be shown that b_2 is in \mathcal{A}'_2 .

Assume to the contrary that $b_2 \in \mathcal{A}_2 \setminus \mathcal{A}'_2$. This means that there exists $a'_1 \in \mathcal{S}_1 \setminus \mathcal{A}_1$ with $b_2 > a'_1$. But we know that $a_1 > b_2 > a'_1$, and $a_1 \in \mathcal{A}_1$. Since \mathcal{A}_1 was assumed to be $>$ -closed in \mathcal{S}_1 , we get $a'_1 \in \mathcal{A}_1$, which is in contradiction with $a'_1 \in \mathcal{S}_1 \setminus \mathcal{A}_1$.

Second, let us prove that \mathcal{B} is $>$ -closed in $\mathcal{S}_1 \cup \mathcal{S}_2$. Thus assume that $b > a$ for $b \in \mathcal{B}$ and $a \in \mathcal{S}_1 \cup \mathcal{S}_2$. We have to show that $a \in \mathcal{B}$. We consider the case where $b \in \mathcal{A}'_1$. (The case $b \in \mathcal{A}'_2$ is symmetric.) Since $a \in \mathcal{S}_1 \cup \mathcal{S}_2$ we have to distinguish two cases.

First assume that $a \in \mathcal{S}_1$. Because \mathcal{A}_1 is $>$ -complete in \mathcal{S}_1 and $b \in \mathcal{A}'_1 \subseteq \mathcal{A}_1$, $b > a$ implies that $a \in \mathcal{A}_1$. Assume that $a \notin \mathcal{A}'_1$. But then there exists $b' \in \mathcal{S}_2 \setminus \mathcal{A}_2$ such that $a > b'$. By transitivity of $>$ we get $b > b'$. This, together with $b' \in \mathcal{S}_2 \setminus \mathcal{A}_2$ contradicts our assumption that $b \in \mathcal{A}'_1$.

Finally, assume that $a \in \mathcal{S}_2$. We know that $a \in \mathcal{A}_2$ since otherwise $b > a$ would imply that $b \notin \mathcal{A}'_1$. It remains to be shown that $a \in \mathcal{A}'_2$. Assume to the contrary that $a \in \mathcal{A}_2 \setminus \mathcal{A}'_2$. This means that there exists $b' \in \mathcal{S}_1 \setminus \mathcal{A}_1$ with $a > b'$. But then $b > b'$ and $b \in \mathcal{A}_1$ is in contradiction with our assumption that \mathcal{A}_1 is $>$ -closed in \mathcal{S}_1 . \square

To show the “left or” property, assume that $\alpha \vdash \gamma$ and $\beta \vdash \gamma$. By \mathcal{S}_1 we denote the set of all models of α , and by \mathcal{S}_2 the set of all models of β . Obviously, $\mathcal{S}_1 \cup \mathcal{S}_2$ is the set of all models of $\alpha \vee \beta$.

Now $\alpha \vdash \gamma$ ($\beta \vdash \gamma$) means that there exists a \succ -dense set \mathcal{A}_1 in \mathcal{S}_1 (\mathcal{A}_2 in \mathcal{S}_2) such that γ holds in all elements of \mathcal{A}_1 (\mathcal{A}_2). The lemma yields a set $\mathcal{B} \subseteq \mathcal{A}_1 \cup \mathcal{A}_2$ that is \succ -dense in $\mathcal{S}_1 \cup \mathcal{S}_2$, the set of all models of $\alpha \vee \beta$. Obviously, γ holds in all elements of \mathcal{B} . This completes the proof that the “left or” property is satisfied by our nonmonotonic consequence relation.

Left logical equivalence is trivially satisfied. In fact, if α and β are logically equivalent then they have the same set of models.

To show *cautious monotony*, we need another simple property of dense sets.

Lemma 5.5 *Let $>$ be a transitive relation on a set \mathcal{S} , and let \mathcal{A}, \mathcal{B} be subsets of \mathcal{S} such that $\mathcal{A} \subseteq \mathcal{B}$. If \mathcal{A} is $>$ -dense in \mathcal{S} then \mathcal{A} is $>$ -dense in \mathcal{B} .*

Proof. Assume that \mathcal{A} is not $>$ -dense in \mathcal{B} . If \mathcal{A} is not $>$ -complete in \mathcal{B} then there is an element b of \mathcal{B} such that there does not exist $a \in \mathcal{A}$ with $b \geq a$. Since $\mathcal{B} \subseteq \mathcal{S}$ this implies that \mathcal{A} is not $>$ -complete in \mathcal{S} .

If \mathcal{A} is not $>$ -closed in \mathcal{B} then there exists $a \in \mathcal{A}$ and $b \in \mathcal{B} \setminus \mathcal{A}$ with $a > b$. Since $\mathcal{B} \subseteq \mathcal{S}$ we get $b \in \mathcal{S} \setminus \mathcal{A}$, which shows that \mathcal{A} is not $>$ -closed in \mathcal{S} . \square

Now assume that $\alpha \vdash \beta$ and $\alpha \vdash \gamma$. Let \mathcal{S} be the set of all models of α , and let \mathcal{B} be the set of all models of $\alpha \wedge \beta$. Now $\alpha \vdash \beta$ ($\alpha \vdash \gamma$) yields a \succ -dense set \mathcal{A}_1 (\mathcal{A}_2) in \mathcal{S} such that β (γ) holds in all elements of \mathcal{A}_1 (\mathcal{A}_2).

By Lemma 5.3 we know that $\mathcal{A}_1 \cap \mathcal{A}_2$ is \succ -dense in \mathcal{S} . Obviously, $\alpha \wedge \beta$ holds in all elements of $\mathcal{A}_1 \cap \mathcal{A}_2$. Thus $\mathcal{A}_1 \cap \mathcal{A}_2 \subseteq \mathcal{B}$, and by Lemma 5.5 $\mathcal{A}_1 \cap \mathcal{A}_2$ is dense in \mathcal{B} . Since γ holds in all elements of $\mathcal{A}_1 \cap \mathcal{A}_2$, we get $\alpha \wedge \beta \vdash \gamma$. This completes the proof that “cautious monotony” holds for our nonmonotonic consequence relation.

The next example demonstrates that *rational monotony* is not always satisfied (as was to be expected).

Example 5.6 We consider the fixed set of defaults $\mathcal{D} = \{ : P(x)/P(x) \}$. The closed formulae α, β, γ are defined as

$$\begin{aligned}\alpha &\equiv \exists x. \neg P(x) \\ \beta &\equiv \exists x. (\neg P(x) \wedge \forall y. (x \neq y \rightarrow P(y))), \\ \gamma &\equiv \forall x. \neg P(x).\end{aligned}$$

We have seen in Example 4.4 that β is a consequence of the default theory $(\mathcal{D}, \{\alpha\})$. Thus we know that $\alpha \sim \beta$. In addition, the argument in Example 4.4 shows that $\neg\gamma$ is not a default consequence of $(\mathcal{D}, \{\alpha\})$, i.e., $\alpha \not\sim \neg\gamma$. In fact, all models of α of cardinality 1 must be elements of a \succ -dense set, and in such a model γ holds.

However, β is not a consequence of the theory $(\mathcal{D}, \{\alpha \wedge \gamma\})$, i.e., $\alpha \wedge \gamma \not\sim \beta$ does not hold. Obviously, models of γ cannot contain positive examples for the default $d = : P(x)/P(x)$. For this reason we know that the preference relation \succ_d is empty on the set of all models of $\alpha \wedge \gamma$. This implies that the set of all models of $\alpha \wedge \gamma$ is the only dense set. But β does not hold in all models of $\alpha \wedge \gamma$. \square

6 Conclusion

In the approach for handling open normal defaults presented in this paper the defaults of a given default theory induce a preference relation \succ between models of its world description. We have shown by examples that this avoids some of the drawbacks of approaches that view open defaults as schemata for certain instantiations.

Differing from the usual preferential approaches, our nonmonotonic consequence relation is not defined with reference to \succ -minimal models. Instead we have introduced the notion of \succ -dense sets. In Section 5 it was shown that, because of this modified preferential approach, our nonmonotonic consequence relation has most of the “nice” proof-theoretic properties mentioned by Gabbay [5], Makinson [9], and Kraus, Lehmann and Magidor [7]. An interesting problem for further research in this direction is to establish an exact correspondence between proof-theoretic properties and our new preferential approach, similarly to the results of Kraus, Lehmann, and Magidor [7] and Schlechta [13] for the usual preferential approaches.

In the present paper, we did not take priorities among defaults into account. A possible solution to this problem could be to restrict the preference


relation \succ_d for a default d in the following way: If d' is of higher priority than d then going from a model M_1 to a d -preferred model M_2 must not delete positive examples or introduce new negative examples for d' .

In Example 4.1, where we considered contradictory defaults d_1, d_2 , this approach would solve the conflict if one gives one of the defaults higher priority. In fact, one would get the same consequences as if only the default with higher priority was present. However, it is a subject of further research to examine this treatment of priorities among open defaults more closely.

References

- [1] F. Baader and B. Hollunder. Embedding defaults into terminological knowledge representation formalisms. In *Proceedings of the 3rd International Conference on Knowledge Representation and Reasoning*, Cambridge, Mass., 1992.
- [2] F. Baader and B. Hollunder. How to prefer more specific defaults in terminological default logic. To appear in *Proceedings of the 13th International Joint Conference on Artificial Intelligence*, Chambery, France, 1993.
- [3] G. Brewka. Preferred subtheories: An extended logical framework for default reasoning. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 1043–1048, Detroit, Mich., 1989.
- [4] G. Brewka. *Nonmonotonic Reasoning: Logical Foundations of Commonsense*. Cambridge University Press, Cambridge, 1991.
- [5] D. Gabbay. Theoretical foundations for non-monotonic reasoning in expert systems. In K.R. Apt, editor, *Proceedings of the NATO Advanced Studies Institute on Logics and Models of Concurrent Systems*, pages 439–457, La Colle-sur-Loup, France, Springer Verlag, 1985.
- [6] U. Junker and G. Brewka. Handling partially ordered defaults in TMS. In *Proceedings of the 1st European Conference on Symbolic and Quantitative Approaches for Uncertainty*, Marseilles, France, 1991.
- [7] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1-2):137–207, 1990.

- [8] V. Lifschitz. On open defaults. In *Proceedings of the Symposium on Computational Logics*, Brüssel, Belgium, 1990.
- [9] D. Makinson. General theory of cumulative inferences. In M. Reinfrank, J. de Kleer, M.L. Ginsberg, editors, *Proceedings of the Second International Workshop on Non-Monotonic Reasoning*, pages 1–18, Springer Verlag, Berlin, 1989.
- [10] D. Makinson. General patterns in nonmonotonic reasoning. To appear in D. Gabbay and C. Hogger, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, vol. II: Nonmonotonic and Uncertain Reasoning, Oxford.
- [11] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1-2):81–132, 1980.
- [12] R. Reiter and G. Criscuolo. On interacting defaults. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, 1981.
- [13] K. Schlechta. Some results on classical preferential models. To appear in *J. of Logic and Computation*, 2(6), 1992.
- [14] Y. Shoham. Nonmonotonic logics: Meaning and utility. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, pages 388–393, Milan, Italy, 1987.



Deutsches
Forschungszentrum
für Künstliche
Intelligenz GmbH

DFKI
-Bibliothek-
PF 2080
D-6750 Kaiserslautern
FRG

DFKI Publikationen

Die folgenden DFKI Veröffentlichungen sowie die aktuelle Liste von allen bisher erschienenen Publikationen können von der oben angegebenen Adresse bezogen werden. Die Berichte werden, wenn nicht anders gekennzeichnet, kostenlos abgegeben.

DFKI Publications

The following DFKI publications or the list of all published papers so far can be ordered from the above address. The reports are distributed free of charge except if otherwise indicated.

DFKI Research Reports

RR-92-15

Winfried Graf: Constraint-Based Graphical Layout of Multimodal Presentations
23 pages

RR-92-16

Jochen Heinsohn, Daniel Kudenko, Bernhard Nebel, Hans-Jürgen Profitlich: An Empirical Analysis of Terminological Representation Systems
38 pages

RR-92-17

Hassan Ait-Kaci, Andreas Podelski, Gert Smolka: A Feature-based Constraint System for Logic Programming with Entailment
23 pages

RR-92-18

John Nerbonne: Constraint-Based Semantics
21 pages

RR-92-19

Ralf Legleitner, Ansgar Bernardi, Christoph Klauck: PIM: Planning In Manufacturing using Skeletal Plans and Features
17 pages

RR-92-20

John Nerbonne: Representing Grammar, Meaning and Knowledge
18 pages

RR-92-21

Jörg-Peter Mohren, Jürgen Müller: Representing Spatial Relations (Part II) -The Geometrical Approach
25 pages

RR-92-22

Jörg Würtz: Unifying Cycles
24 pages

RR-92-23

Gert Smolka, Ralf Treinen: Records for Logic Programming
38 pages

RR-92-24

Gabriele Schmidt: Knowledge Acquisition from Text in a Complex Domain
20 pages

RR-92-25

Franz Schmalhofer, Ralf Bergmann, Otto Kühn, Gabriele Schmidt: Using integrated knowledge acquisition to prepare sophisticated expert plans for their re-use in novel situations
12 pages

RR-92-26

Franz Schmalhofer, Thomas Reinartz, Bidjan Tschaitchian: Intelligent documentation as a catalyst for developing cooperative knowledge-based systems
16 pages

RR-92-27

Franz Schmalhofer, Jörg Thoben: The model-based construction of a case-oriented expert system
18 pages

RR-92-29

Zhaohui Wu, Ansgar Bernardi, Christoph Klauck: Skeletal Plans Reuse: A Restricted Conceptual Graph Classification Approach
13 pages

RR-92-30

Rolf Backofen, Gert Smolka: A Complete and Recursive Feature Theory
32 pages

RR-92-31

Wolfgang Wahlster: Automatic Design of Multimodal Presentations
17 pages

DFKI Documents**D-92-09**

Gernod P. Laufkötter: Implementierungsmöglichkeiten der integrativen Wissensakquisitionsmethode des ARC-TEC-Projektes
86 Seiten

D-92-10

Jakob Mauss: Ein heuristisch gesteuerter Chart-Parser für attributierte Graph-Grammatiken
87 Seiten

D-92-11

Kerstin Becker: Möglichkeiten der Wissensmodellierung für technische Diagnose-Expertensysteme
92 Seiten

D-92-12

Otto Kühn, Franz Schmalhofer, Gabriele Schmidt: Integrated Knowledge Acquisition for Lathe Production Planning: a Picture Gallery (Integrierte Wissensakquisition zur Fertigungsplanung für Drehteile: eine Bildergalerie)
27 pages

D-92-13

Holger Peine: An Investigation of the Applicability of Terminological Reasoning to Application-Independent Software-Analysis
55 pages

D-92-14

Johannes Schwagereit: Integration von Graph-Grammatiken und Taxonomien zur Repräsentation von Features in CIM
98 Seiten

D-92-15

DFKI Wissenschaftlich-Technischer Jahresbericht 1991
130 Seiten

D-92-16

Judith Engelkamp (Hrsg.): Verzeichnis von Softwarekomponenten für natürlichsprachliche Systeme
189 Seiten

D-92-17

Elisabeth André, Robin Cohen, Winfried Graf, Bob Kass, Cécile Paris, Wolfgang Wahlster (Eds.): UM92: Third International Workshop on User Modeling, Proceedings
254 pages

Note: This document is available only for a nominal charge of 25 DM (or 15 US-\$).

D-92-18

Klaus Becker: Verfahren der automatisierten Diagnose technischer Systeme
109 Seiten

D-92-19

Stefan Dittrich, Rainer Hoch: Automatische, Deskriptor-basierte Unterstützung der Dokumentanalyse zur Fokussierung und Klassifizierung von Geschäftsbriefen
107 Seiten

D-92-21

Anne Schauder: Incremental Syntactic Generation of Natural Language with Tree Adjoining Grammars
57 pages

D-92-22

Werner Stein: Indexing Principles for Relational Languages Applied to PROLOG Code Generation
80 pages

D-92-23

Michael Herfert: Parsen und Generieren der Prolog-artigen Syntax von RELFUN
51 Seiten

D-92-24

Jürgen Müller, Donald Steiner (Hrsg.): Kooperierende Agenten
78 Seiten

D-92-25

Martin Buchheit: Klassische Kommunikations- und Koordinationsmodelle
31 Seiten

D-92-26

Enno Tolzmann: Realisierung eines Werkzeugauswahlmoduls mit Hilfe des Constraint-Systems CONTAX
28 Seiten

D-92-27

Martin Harm, Knut Hinkelmann, Thomas Labisch: Integrating Top-down and Bottom-up Reasoning in COLAB
40 pages

D-92-28

Klaus-Peter Gores, Rainer Bleisinger: Ein Modell zur Repräsentation von Nachrichtentypen
56 Seiten

D-93-01

Philipp Hanschke, Thom Frühwirth: Terminological Reasoning with Constraint Handling Rules
12 pages

D-93-02

Gabriele Schmidt, Frank Peters, Gernod Laufkötter: User Manual of COKAM+
23 pages

D-93-03

Stephan Busemann, Karin Harbusch (Eds.): DFKI Workshop on Natural Language Systems: Reusability and Modularity - Proceedings
74 pages

RR-92-59

Karl Schlechta and David Makinson: On Principles and Problems of Defeasible Inheritance
13 pages

RR-92-60

Karl Schlechta: Defaults, Preorder Semantics and Circumscription
19 pages

RR-93-02

Wolfgang Wahlster, Elisabeth André, Wolfgang Finkler, Hans-Jürgen Profülich, Thomas Rist: Plan-based Integration of Natural Language and Graphics Generation
50 pages

RR-93-03

Franz Baader, Bernhard Hollunder, Bernhard Nebel, Hans-Jürgen Profülich, Enrico Franconi: An Empirical Analysis of Optimization Techniques for Terminological Representation Systems
28 pages

RR-93-04

Christoph Klauck, Johannes Schwagereit: GGD: Graph Grammar Developer for features in CAD/CAM
13 pages

RR-93-05

Franz Baader, Klaus Schulz: Combination Techniques and Decision Problems for Disunification
29 pages

RR-93-08

Harold Boley, Philipp Hanschke, Knut Hinkelmann, Manfred Meyer: COLAB: A Hybrid Knowledge Representation and Compilation Laboratory
64 pages

RR-93-09

Philipp Hanschke, Jörg Würtz: Satisfiability of the Smallest Binary Program
8 Seiten

RR-93-11

Bernhard Nebel, Hans-Juergen Buerckert: Reasoning about Temporal Relations: A Maximal Tractable Subclass of Allen's Interval Algebra
28 pages

RR-93-12

Pierre Sablayrolles: A Two-Level Semantics for French Expressions of Motion
51 pages

RR-93-13

Franz Baader, Karl Schlechta: A Semantics for Open Normal Defaults via a Modified Preferential Approach
25 pages

DFKI Technical Memos**TM-91-12**

Klaus Becker, Christoph Klauck, Johannes Schwagereit: FEAT-PATR: Eine Erweiterung des D-PATR zur Feature-Erkennung in CAD/CAM
33 Seiten

TM-91-13

Knut Hinkelmann: Forward Logic Evaluation: Developing a Compiler from a Partially Evaluated Meta Interpreter
16 pages

TM-91-14

Rainer Bleisinger, Rainer Hoch, Andreas Dengel: ODA-based modeling for document analysis
14 pages

TM-91-15

Stefan Busemann: Prototypical Concept Formation An Alternative Approach to Knowledge Representation
28 pages

TM-92-01

Lijuan Zhang: Entwurf und Implementierung eines Compilers zur Transformation von Werkstückrepräsentationen
34 Seiten

TM-92-02

Achim Schupeta: Organizing Communication and Introspection in a Multi-Agent Blocksworld
32 pages

TM-92-03

Mona Singh: A Cognitive Analysis of Event Structure
21 pages

TM-92-04

Jürgen Müller, Jörg Müller, Markus Pischel, Ralf Scheidhauer: On the Representation of Temporal Knowledge
61 pages

TM-92-05

Franz Schmalhofer, Christoph Globig, Jörg Thoben: The refitting of plans by a human expert
10 pages

TM-92-06

Otto Kühn, Franz Schmalhofer: Hierarchical skeletal plan refinement: Task- and inference structures
14 pages

TM-92-08

Anne Kilger: Realization of Tree Adjoining Grammars with Unification
27 pages

RR-92-33

Franz Baader: Unification Theory
22 pages

RR-92-34

Philipp Hanschke: Terminological Reasoning and Partial Inductive Definitions
23 pages

RR-92-35

Manfred Meyer:
Using Hierarchical Constraint Satisfaction for Lathe-Tool Selection in a CIM Environment
18 pages

RR-92-36

Franz Baader, Philipp Hanschke:
Extensions of Concept Languages for a Mechanical Engineering Application
15 pages

RR-92-37

Philipp Hanschke: Specifying Role Interaction in Concept Languages
26 pages

RR-92-38

Philipp Hanschke, Manfred Meyer:
An Alternative to Θ -Subsumption Based on Terminological Reasoning
9 pages

RR-92-40

Philipp Hanschke, Knut Hinkelmann: Combining Terminological and Rule-based Reasoning for Abstraction Processes
17 pages

RR-92-41

Andreas Lux: A Multi-Agent Approach towards Group Scheduling
32 pages

RR-92-42

John Nerbonne:
A Feature-Based Syntax/Semantics Interface
19 pages

RR-92-43

Christoph Klauck, Jakob Mauss: A Heuristic driven Parser for Attributed Node Labeled Graph Grammars and its Application to Feature Recognition in CIM
17 pages

RR-92-44

Thomas Rist, Elisabeth André: Incorporating Graphics Design and Realization into the Multimodal Presentation System WIP
15 pages

RR-92-45

Elisabeth André, Thomas Rist: The Design of Illustrated Documents as a Planning Task
21 pages

RR-92-46

Elisabeth André, Wolfgang Finkler, Winfried Graf, Thomas Rist, Anne Schauder, Wolfgang Wahlster: WIP: The Automatic Synthesis of Multimodal Presentations
19 pages

RR-92-47

Frank Bomarius: A Multi-Agent Approach towards Modeling Urban Traffic Scenarios
24 pages

RR-92-48

Bernhard Nebel, Jana Koehler:
Plan Modifications versus Plan Generation: A Complexity-Theoretic Perspective
15 pages

RR-92-49

Christoph Klauck, Ralf Legleitner, Ansgar Bernardi:
Heuristic Classification for Automated CAPP
15 pages

RR-92-50

Stephan Busemann:
Generierung natürlicher Sprache
61 Seiten

RR-92-51

Hans-Jürgen Bürckert, Werner Nutt:
On Abduction and Answer Generation through Constrained Resolution
20 pages

RR-92-52

Mathias Bauer, Susanne Biundo, Dietmar Dengler, Jana Koehler, Gabriele Paul: PHI - A Logic-Based Tool for Intelligent Help Systems
14 pages

RR-92-54

Harold Boley: A Direkt Semantic Characterization of RELFUN
30 pages

RR-92-55


John Nerbonne, Joachim Laubsch, Abdel Kader Diagne, Stephan Oepen: Natural Language Semantics and Compiler Technology
17 pages

RR-92-56

Armin Laux: Integrating a Modal Logic of Knowledge into Terminological Logics
34 pages

RR-92-58

Franz Baader, Bernhard Hollunder:
How to Prefer More Specific Defaults in Terminological Default Logic
31 pages



A Semantics for Open Normal Defaults via a Modified Preferential Approach

Franz Baader, Karl Schlechta

RR-93-13

Research Report