

Default Theory of Defeasible Entailment

Alexander Bochman

Computer Science Department,
Holon Institute of Technology, Israel
bochmana@hit.ac.il

Abstract

We suggest a new representation of defeasible entailment and specificity in the framework of default logic. The representation is based on augmenting the underlying classical language with the language of conditionals having its own (monotonic) internal logic. It is shown, in particular, that nonmonotonic inheritance reasoning can be naturally represented in this framework, and generalized to the full classical language.

The problem of nonmonotonic, defeasible inference can be seen as the main objective, as well as the main problem of the general theory of nonmonotonic reasoning. An impressive success has been achieved in our understanding of it, in realizing how complex it is, and, most importantly, how many different forms it may have. Many formalisms have been developed and implemented that capture significant aspects of defeasible inference, though a unified picture has not yet emerged. In this study we will suggest a new representation of defeasible inference in the framework of default logic. As the reader will see, however, the suggested representation will borrow the main insights of many previous approaches to this problem, hopefully without inheriting their shortcomings.

Preliminary version of this paper has appeared as (Bochman 2007), though in the present study we suggest a somewhat different, more transparent formalization.

The Problem of Defeasible Inference a very brief history

In this paper we will view the problem of defeasible inference as the problem of formalizing, or representing, commonsense rules of the form “*A normally implies B*” (they will be written as $A \rightarrow B$ in what follows). Viewed in this way, it can be safely argued that the problem of defeasible inference was born precisely with the birth of nonmonotonic reasoning. Indeed, one of the main objectives of the first nonmonotonic formalisms, namely circumscription, default logic and modal nonmonotonic logics, was a faithful formal representation of these commonsense rules (prominent examples being *Birds fly* versus *Penguins can't fly*).

Copyright © 2008, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The three initial formal nonmonotonic systems proposed a translational approach to this problem. Thus, Ray Reiter has suggested in (Reiter 1980) to represent $A \rightarrow B$ using ‘normal’ default rules

$$A : B/B.$$

The main effect of this representation consisted in securing two important properties of the corresponding commonsense rule $A \rightarrow B$: (i) the rule is applicable whenever A holds and $\neg B$ is not known to hold, yet (ii) a possible factual refutation $A \wedge \neg B$ does not create contradiction in the system.

A similar modal translation $A \wedge MB \supset B$ of the rules $A \rightarrow B$ was suggested in the modal nonmonotonic logics of (McDermott and Doyle 1980).

Initially, Reiter has mentioned in (Reiter 1980) that he knows of no naturally occurring default which cannot be represented in this form. His views have changed, however, already in (Reiter and Criscuolo 1981) where it has already been argued that normal default rules are insufficient to deal with interactions of different defaults. For example, in combining two defaults, *Birds fly* and *Penguins can't fly*, the specificity principle naturally suggests that the second, more specific, default should be preferred, so *Birds fly* shouldn't be applied to penguins. The authors argued that to capture this outcome in default logic, we need at least *semi-normal* defaults of the form $A : B \wedge C/C$, so *Birds fly* could now be represented roughly as “Birds normally fly, unless they are penguins”, encoded as $B : F \wedge \neg P/F$.

Unfortunately, this complication has significantly undermined the transparency and modularity of the initial default representation: according to the modified proposal, the very representation of the claim “*A normally implies B*” has become dependent on other defaults and constraints present in the description. Worse still, no systematic method of constructing the resulting default representation for a given set of rules has emerged since then. The problem of defeasible inference has shown its real complexity.

In contrast to the above representations, John McCarthy has suggested in (McCarthy 1980) a purely classical translation of normality conditionals $A \rightarrow B$ as implications of the form

$$A \wedge \neg ab \supset B,$$

where *ab* is a new ‘abnormality’ proposition serving to accumulate the conditions for violation of the source rule. Thus, *Birds fly* was translated into something like “Birds fly if they are not abnormal”.

In fact, viewed as a formalism for nonmonotonic reasoning, the central concept of McCarthy’s circumscriptive method was not circumscription itself, but his notion of an abnormality theory - a set of classical conditionals containing the abnormality predicate *ab* that provides a representation for defeasible information. The default character of commonsense rules $A \rightarrow B$ was captured in McCarthy’s theory by a circumscription policy that minimized abnormality (and thereby maximized the acceptance of the corresponding normality claims $\neg ab$). Since then, this representation of normality rules using auxiliary (ab)normality propositions has been employed both in applications of circumscription, and in many other theories of nonmonotonic inference in AI, sometimes in alternative logical frameworks. Some major examples are inheritance theories (Etherington and Reiter 1983), logic-based diagnosis (Reiter 1987), general representation of defaults in (Konolige and Myers 1989) and reasoning about time and action. Note also that naming of defaults (as in (Poole 1988)) can also be viewed as a species of this idea. In fact, the approach described later in this paper can also be seen as a development of this representation.

Abnormality theories have brought out, however, much the same problems of representing conflicting commonsense rules that plagued the default representation. A general approach to handle such problems in circumscription, suggested in (Lifschitz 1985) and endorsed in (McCarthy 1986), was to impose priorities among minimized predicates and abnormalities. The corresponding variant of circumscription has been called prioritized circumscription. The problem of defeasible inference was reduced in this way to the problem of finding the ‘right’ priority order for a given set of defaults. Unfortunately, in this case also no systematic understanding or principles behind constructing such a priority order have emerged.

A more specific problem of the abnormality representation (which will be especially relevant in what follows) has arisen from the need to cope with the possibility of having two independent commonsense rules $A \rightarrow B$ and $A \rightarrow C$ with the same antecedent, for example *Birds fly* and *Birds have wings*. In order to preserve one of these rules when the other is refuted, McCarthy was forced in (McCarthy 1986) to relativize the abnormality claims with respect to particular *aspects*, so that some aspects can be abnormal without affecting others. For our example, birds might be abnormal with respect to flying without being abnormal with respect to having wings, and vice versa.

In reading McCarthy’s papers on circumscription, one cannot help feeling uneasiness with which he adopted the abnormality predicates and their aspects into the language, since he thought that this compels us to introduce (aspects of) abnormalities as new entities into our ontology, the things that exist. As was noted by McCarthy himself in (McCarthy 1986), the aspects are abstract entities, and their unintuitiveness is somewhat a blemish on the theory.

The obvious difficulties encountered in representing defeasible inference in the first nonmonotonic formalisms have become a strong incentive for developing *direct*, non-translational theories of normality conditionals. One of the first approaches of this kind were based on an attempt of describing the very notion of defeasible inference, or derivation, that would take into account the impact and interaction of conflicting rules. Defeasible logic of Nute (see, e.g., (Nute 1994)) could be seen a representative example of this approach (see also (Wagner 1991) for a concise description of the underlying principles).

In fact, the same approach lied at the basis of the theory of nonmonotonic inheritance nets ((Touretzky 1986); see also an overview in (Horty 1994)). Reasoning in inheritance hierarchies was represented in this theory as a process of constructing acceptable derivations (‘paths’) that formed a basis of making conclusions. However, inheritance reasoning has dealt with a quite restricted class of conditionals built on literals. Nevertheless, in this restricted domain it has achieved a remarkably close correspondence between what is derived and what is expected intuitively. Accordingly, inheritance reasoning has emerged as an important test bed for adjudicating proposed theories. Mainly for this reason we have chosen to focus this study on a faithful generalization of inheritance reasoning to a full-fledged theory of defeasible inference.

The next stage in the development of the direct approach to defeasible inference was undoubtedly a theory of preferential entailment (Kraus, Lehmann, and Magidor 1990). In this theory the normality conditionals $A \vdash B$ have acquired a precise direct semantic interpretation in terms of a preference relation on worlds, and a complete syntactic characterization in a set of natural postulates for defeasible inference. Moreover, it has been shown that the resulting theory preserves the main intuitive features of normality claims and, in particular, properly handles the specificity principle.

Unfortunately, it was realized quite early that preferential entailment also cannot serve as the ultimate theory of defeasible inference. Preferential inference is severely sub-classical and does not allow us, for example, to infer *Red birds fly* from *Birds fly*. Clearly, there are good reasons for not accepting such a derivation as a *logical* rule for defeasible inference; otherwise *Birds fly* would imply also *Birds with broken wings fly* and even *Penguins fly*. Still, this should not prevent us from accepting *Red birds fly* on the basis of *Birds fly* as a reasonable *defeasible* conclusion, namely a conclusion made in the absence of information against it. By doing this, we would just follow the general strategy of nonmonotonic reasoning that involves making reasonable assumptions on the basis of available information. In other words, it was realized that preferential inference should be augmented at least with a mechanism of making nonmonotonic conclusions that goes beyond the latter. This has led to various theories of *defeasible entailment* built upon preferential inference.

Actually, the literature on nonmonotonic reasoning is abundant with theories of defeasible entailment. A history of studies on this subject could be briefly summarized as follows. Initial formal systems, namely Lehmann’s rational

closure (Lehmann and Magidor 1992) and Pearl's system Z (Pearl 1990), have turned out to be equivalent. This encouraging development has followed by a realization that both theories are still insufficient for representing defeasible inference, since they do not allow to make certain intended conclusions. Hence, they have been refined in a number of ways, giving such systems as lexicographic inference (Benferhat et al. 1993; Lehmann 1995), and similar modifications of Pearl's system. Unfortunately, these refined systems have encountered an opposite problem, namely, together with some desirable properties, they invariably produced some unwanted conclusions. All these systems have been based on a supposition that defeasible entailment should form a rational inference relation. A more general approach in the framework of preferential inference has been suggested in (Geffner 1992).

Conditional entailment (Geffner 1992) determines a systematic prioritization of conditional bases securing that $A \vdash B$ will hold for any default conditional $A \rightarrow B$ that belongs to the base. Then the intended models of a conditional base are defined as models that are generated by all such admissible priority orders.

It is interesting to note that, in the framework of conditional entailment, the conditionals $A \rightarrow B$ admit a representation much similar to their representation in McCarthy's circumscription¹. Namely, $A \rightarrow B$ can be represented as a classical implication

$$A \wedge \delta \supset B,$$

where δ is a new 'normality' proposition unique to each default conditional, plus a new, conceptually simpler conditional $A \rightarrow \delta$. Actually, this kind of reduction will also be used in our approach, described below.

Conditional entailment has shown itself as a serious candidate on the role of a general theory of defeasible entailment. Still, it does not capture another natural theory, namely inheritance reasoning. The first main difference between the two theories is that conditional entailment is based on establishing absolute priorities among defaults, while inheritance hierarchies determine such priorities in a context-dependent way, namely in presence of other defaults that provide a (preemption) link between conflicting defaults. As a result, in some cases (that we will discuss later) inheritance reasoning provides more adequate outcome. In fact, it has been shown in (Bochman 2001) that inheritance reasoning is representable by using *conditional* priority orders on defaults. But unfortunately, the corresponding construction could hardly be called simple or natural.

The second major difference between the two theories is that inheritance reasoning is based on viewing the normality conditionals $A \rightarrow B$ as inference *rules*, while conditional entailment preserves a strong connection between conditionals and corresponding material implications; as a by-product, it allows for a good deal of backward (contrapositive) reasoning. Geffner himself has shown in (Geffner 1992) that conditional entailment does not capture some im-

portant derivations that are based on a more directional, rule-based view of conditionals. Accordingly, in the last chapter of his book he suggested to extend conditional entailment with an explicit representation of causal, directional order among propositions. The resulting theory, however, has not received a widespread acceptance. On the final account, the problem of defeasible inference has remained a major problem of nonmonotonic reasoning.

The New Approach

It does not have to be argued that a most rational way of introducing a systematic theory of defeasible inference (or of anything else) would consist in stating right away its desiderata, and then proceeding to establishing the basic principles, or postulates that would both accord with our intuitions and fulfil these objectives. Unfortunately, there are two reasons why this way of dealing with the problem does not seem open in our present case, at least for now. First, we are not working from the scratch - many theories of defeasible inference have already been suggested, often with conflicting underlying principles. Second, it is well known today that our commonsense intuitions about defeasible inference are patently vague and fragmentary. They usually do not go beyond a few very simple instances, so they are often insufficient, taken by themselves, for adjudicating alternative theories.

In this study we are guided by a more modest approach. To begin with, we claim that human, commonsense defeasible inference, however vague and incomplete, plays an essential and extremely useful role in our everyday reasoning, so a formal representation of such a reasoning should be viewed as an important task of knowledge representation and AI. It should be noted, however, that, taken by itself, a formalization of commonsense reasoning is not the primary objective of Knowledge Representation and Reasoning. Rather, the aim of the latter consists in developing adequate and computationally efficient frameworks for a *rational* reasoning in general. In this sense, different formalizations of defeasible inference are possible, provided each of them produces solutions that work in resolving essential reasoning problems of AI. Of course, the degree of accord with commonsense reasoning is an important measure of success for such a formalization, but it is by no means the only measure.

In view of what was said above, the main objective of this study will consist not in suggesting a new general theory of defeasible inference, but rather in determining the basic principles and building blocks using which such a theory could be constructed. To this end, we will take below an existing theory of defeasible inheritance and try to find a most natural and minimal way this theory could be built. Our hope, however, is that on this way we will arrive at the principles and constructions that may have general significance.

On a most natural, commonsense understanding, a rule $A \rightarrow B$ represents a claim that A implies B , given some additional (unmentioned and even not fully known) assumptions that are presumed to hold in normal circumstances. Thus, a default causal rule $TurnKey \rightarrow CarStarts$ states

¹This alternative representation was actually employed in (Geffner and Pearl 1992).

that if I turn the key, the car will start given the normal conditions such as there is a fuel in the tank, the battery is ok, etc. etc.

Our approach below is based on the idea that the auxiliary default assumptions of a defeasible rule $A \rightarrow B$ provide a link (an information channel) that sanctions the inference from A to B . In other words, they jointly function as a conditional, that we will denote by A/B , that, once accepted, allows us to infer B from A . Accordingly, we may represent $A \rightarrow B$ as the classical implication

$$A \wedge (A/B) \supset B.$$

Now, as in the approach of Geffner, described above, the default character of the inference from A to B can be captured by requiring that A/B normally holds or, more cautiously, that A normally implies A/B , that is, $A \rightarrow (A/B)$. In contrast to the conditional entailment, however, we will achieve this effect by representing the latter rule simply as a normal default in default logic.

In many respects, our representation can be viewed as an ‘unfolding’ of the corresponding representations of McCarthy and Geffner. Namely, in place of unstructured (ab)normality propositions, we suggest to use more articulated conditional propositions that can be assigned a natural meaning that will facilitate their conscious and coherent use in defeasible inference. Note, for example, that already at this stage the representation suggests a natural explanation why the default assumptions required for the rules $A \rightarrow B$ and $A \rightarrow C$ can be presumed to be independent, so violation of $A \rightarrow B$ does not imply rejection of $A \rightarrow C$. It is this feature that has required introduction of aspects of abnormality in McCarthy’s circumscription. In fact, this feature goes also beyond the usual assumptions about normality made in preferential inference: according to the latter, a violation of $A \rightarrow B$ means that the situation at hand is abnormal with respect to A , so we are not entitled to infer anything that normally follows from A .

It is important to note that the change we have made so far to McCarthy’s or Geffner’s representation is purely terminological. This means, in particular, that the implementation of the theory of defeasible inference that will be described in the sequel can be made, in principle, by switching back to abnormality predicates and established formalisms of dealing with them. Our representation naturally suggests, however, that the (ab)normality claims might have its own internal logic. In fact, we will stipulate below that the conditionals A/B should satisfy at least the usual rules of Tarski consequence relations. It is this internal logic that will allow us to formulate purely logical principles that will govern a proper interaction of defeasible rules in cases of conflict.

The Language and Logic of Conditionals

Our basic language L will be a classical propositional language with the usual connectives and constants $\{\wedge, \vee, \neg, \supset, \text{t}, \text{f}\}$. $\models$ will stand for the classical entailment, while Th will denote the associated provability operator.

As a first step, we will augment the language L by adding new propositional atoms of the form A/B , where A and B are classical propositions of L . The conditionals A/B will

be viewed as propositional atoms of a new type, so nesting of conditionals will not be allowed. Still, the new propositional atoms will be freely combined with ordinary ones using the classical propositional connectives. We will denote the resulting language by L_c .

The essence and main functional role of our conditionals will be expressed by adopting the following ‘modus ponens’ axiom:

$$\text{MP} \quad (A \wedge A/B) \supset B.$$

In addition, conditionals will be viewed as ordinary inference rules that are ‘reified’ in the object language. Accordingly, we will require them to satisfy all the properties of a supraclassical Tarski consequence relation. The following postulates can be shown to be sufficient for this purpose:

$$\text{If } A \models B, \text{ then } A/B. \quad (\text{Dominance})$$

$$\frac{A/B \quad B/C}{A/C} \quad (\text{Transitivity})$$

$$\frac{A/B \quad A/C}{A/(B \wedge C)} \quad (\text{And})$$

Remark 1. The above conditional logic can be given a complete semantic interpretation. In fact, it coincides with the logic of *monotonic* consequence relations, sketched in (Kraus, Lehmann, and Magidor 1990), and the semantics of simple preferential models, described in that paper (namely, models with an empty preference relation) has been shown to be adequate for the latter. However, the semantic representation of our language will play no role in our subsequent constructions, so we will omit its detailed description here.

Remark 2. Our conditional logic has also many features in common with a (homogeneous) logic of information flow described in (Barwise, Gabbay, and Hartonas 1995), though the latter is far more general. An information network is a structure consisting of sites and (information) channels among them. Ordinary propositions provide description of sites (they characterize types of sites), while conditionals in our sense can be seen as a main kind of descriptions for channels. In a homogeneous network (as in our system) channels is part of sites, so both kinds of descriptions are mixed in the corresponding logic in a single formalism. The logic of information flow allows to describe more general kinds of interactions between sites and channels, and in this sense it is more general than our setting. On the other hand, in our more specific situation we can explore a much richer set of further properties relevant to our problem. One such possibility that is actually used in what follows is a possibility of rejecting conditionals.

So far, we haven’t exploited yet another useful aspect of our conditional language, namely the possibility of negating conditionals. Using this possibility, we can formulate the principles of rejection for conditionals that will play an important role in our final representation. One such principle, the principle of Commitment, will be given in the next section; it will hold only for default conditionals. Below we provide, however, a subsidiary logical principle that will determine how the rejection is propagated along chains of conditionals.

$$\frac{A/B \quad \neg(A/C)}{\neg(B/C)} \quad (\text{Forward Rejection})$$

Forward Rejection can be viewed as a specific partial contraposition of the basic Transitivity rule for conditionals. Note however that, since Transitivity was formulated as an inference *rule*, Forward Rejection cannot be logically derived from the latter. In fact, Forward Rejection provides a natural formalization of the principle of *forward chaining* adopted in many versions of the nonmonotonic inheritance theory.

The above logic describes the logical properties of arbitrary conditionals, not only default ones. The difference between the two will be reflected in the representation of defeasible conditionals in the framework of default logic, described next.

Defeasible Inference in Default Logic

We will describe now a modular representation of defeasible rules $A \rightarrow B$ in default logic. Due to space limitations, we will refrain from a detailed description of default logic itself, but only fix the notation.

A *default theory* is a pair (W, D) , where W is a set classical propositions (the axioms), and D is a set of default rules of the form $A : b \vdash C$, where A, C are propositions and b a finite set of propositions. A is called a *prerequisite* of the rule, b a set of its *justifications*, and C - its *conclusion*. The notion of an extension of a default theory is defined as usual (see, e.g., (Marek and Truszczyński 1993)): for a set s of propositions, let $\mathcal{D}(s)$ be the set of all propositions that are derivable from W using the classical entailment and the following inference rules:

$$\{A \vdash C \mid A : b \vdash C \in D \ \& \ \neg B \notin s, \text{ for any } B \in b\}.$$

Then s is an *extension* of the default theory if and only if $s = \mathcal{D}(s)$.

For the present study, we will suppose that our default theory is defined in the conditional language L_c and respects the corresponding logic of conditionals, that is, W includes MP and Dominance axioms², while D includes Transitivity, And and Forward Rejection as default rules without justifications.

Now, for each defeasible rule $A \rightarrow B$, we accept the normal default

$$A : A/B \vdash A/B.$$

This rule secures that a default conditional is acceptable whenever its antecedent holds and it is not rejected.

Finally, we have two natural options for representing strict (non-defeasible) rules in this framework. A more cautious understanding would lead to representing a strict rule $A \Rightarrow B$ as a strict default rule $A : \vdash B$ without justifications. A more ‘classical’ understanding would amount to representing $A \Rightarrow B$ as a material implication $A \supset B$; this could be achieved in our framework simply by adding A/B to the set W of axioms.

²Namely, W should include all A/B such that $A \models B$.

Specificity and Commitment

The default theory described so far is still insufficient, of course, for capturing defeasible inference, since it does not take into account the *principle of specificity*. In fact, it can be shown that such a default theory is practically indistinguishable from the default theory obtained by representing every defeasible rule $A \rightarrow B$ as a normal default rule $A : B \vdash B$.

Fortunately, we now have sufficient means for expressing the specificity principle in a simple and transparent way. The formulation of this principle, given below, can be seen as the main contribution of this study.

Taken literally, the specificity principle states that more specific default rules should override less specific rules in the case of conflict. In a simplest case, this pertains to the conflict between the rules $A \rightarrow C$ and $A \wedge B \rightarrow \neg C$, in which case the second rule should override the first. There are, however, less direct cases in which one of the conflicting rules is deemed more specific than the other, and, as we all know, the literature is abundant with the attempts to define a more general notion of specificity.

We claim that the specificity principle is a consequence of a more general principle that we will call the *principle of commitment*. According to the latter, by asserting a defeasible rule “If A , then normally B ”, we are also committing ourselves to the claim that no combination of accepted rules could allow us to infer $\neg B$ from A . A clear, though informal, expression of this principle can be found already in (Poole 1990):

“if ‘p’s are q’s’ is a default and if we know $p(c)$, then all of the objections that could be raised about $q(c)$ that follow from $p(c)$ have already been taken into account when building the knowledge base.”

The principle of commitment is related also to the principle of direct inference stated in (Geffner 1992). According to the latter, if $A \rightarrow B$ is accepted, then A should always imply B in the case A is the only evidence we have.

For our present purposes, the principle of commitment can be formally stated as follows:

Commitment If $A \rightarrow B$ is accepted, then

$$A : \vdash \neg(A/\neg B).$$

Thus, the Commitment principle states that if $A \rightarrow B$ is accepted, and A is known to hold, then the opposite conditional $A/\neg B$ should be rejected. Note, however, that $A/\neg B$ may be a consequence of a chain of conditionals that starts with A and ends with $\neg B$. In this case Forward Rejection would dictate, in effect, that the last conditional in any such chain should also be rejected. Indeed, the combination of Commitment and Forward Rejection implies that, if $A \rightarrow B$ is accepted, then, for any proposition C ,

$$A, A/C : \vdash \neg(C/\neg B).$$


In fact, the above rule has been used instead of Commitment and Forward Rejection in the previous formulation of our theory described in (Bochman 2007).

A general scheme of reasoning in the above framework could be described as follows. To begin with, the definition

of an extension requires that if the antecedent A of a default conditional A/B is accepted then the conditional itself should be accepted, unless it is rejected (i.e., its negation is accepted). Now, there are two basic ways in which a conditional can be rejected in an extension. First, it can be rejected due to the facts, since $A \wedge \neg B$ implies $\neg(A/B)$ by MP. Second, a conditional A/B can be rejected due to a commitment to the default rule $A \rightarrow \neg B$. The logic of conditionals supplies, however, important additional means of rejecting conditionals. For instance, if an already rejected conditional A/B is a consequence of two conditionals A/C and C/B , then acceptance of A/C leads to rejection of C/B due to Forward Rejection. Actually, it is this feature that allows us to reject a less specific conditional due to a commitment to a more specific one (see examples below).

The following couple of examples will show how the resulting system works and, in particular, how it handles specificity.

Example 1 (A generalized Penguin-Bird story).



Consider a generalized penguin-bird theory


$$\{P \rightarrow B, B \rightarrow F, P \rightarrow \neg F\}.$$

This theory differs from the usual example about penguins and birds only in that $P \rightarrow B$ is not a strict but defeasible rule. Of course, the same results will be obtained also in this original case.

As could be expected, given the only fact B , the corresponding default theory has a unique extension that contains B and F . However, given the fact P instead, the resulting default theory also has a unique extension that includes this time P, B and $\neg F$. This happens because, given P , the commitment principle for $P \rightarrow \neg F$ implies $\neg(P/F)$. Taken together with P/B , this implies rejection of B/F by Forward Rejection. This is the way a more specific rule $P \rightarrow \neg F$ cancels a less specific rule $B \rightarrow F$. That is why F cannot be derived in this informational situation, so the spurious extension containing F is blocked. Note, however, that the situation is not symmetric, since the commitment to the less specific default $B \rightarrow F$ does not allow us to reject more specific rule $P \rightarrow \neg F$. That is why, for instance, we would still have a unique extension containing $\neg F$ even if both P and B were given as facts.

The following example from (Dung and Son 2001) shows that the above representation deals correctly with the interplay of specificity and evidence, unlike the representations such as prioritized circumscription or Geffner's conditional entailment that are based on establishing context-independent priorities among default rules.

Example 2 (Married Students).



Let us consider the default theory

$$\{A \rightarrow M, S \rightarrow \neg M, S \rightarrow Y, Y \Rightarrow A\}$$

that represents, respectively, default assertions that adults are normally married, students are normally not married, students are normally young adults, and the strict rule “young adults are adults”. For the evidence S , the corresponding default theory will have a single extension containing $\neg M$. In this extension S/M is rejected due to the commitment to $S \rightarrow \neg M$. Moreover, the extension will contain also S/Y and Y/A (since nothing known could reject them), so it will include S/A (by Transitivity). As a result, the default conditional A/M will be rejected in this case by Forward Rejection from $\neg(S/M)$ and S/A .

In theories based on imposing preferences or priorities among defaults, the same effect is usually achieved by making $S \rightarrow \neg M$ preferred to $A \rightarrow M$. But since such a priority stipulation is context-independent, the latter preference should hold also for a more specific evidence $S \wedge \neg Y \wedge A$, which will result in an awkward conclusion that even non-young students are not married.

In contrast, the ‘priority’ of $S \rightarrow \neg M$ over $A \rightarrow M$ in our representation is not absolute, since it depends on the acceptance of S/A . Accordingly, given the more specific evidence $S \wedge \neg Y \wedge A$, the conditional S/A will no longer be acceptable (since the default S/Y is refuted by MP), so the resulting default theory will have two extensions, one containing M , another containing $\neg M$. In other words, as should be expected, the marital status of non-young students cannot be decided.

Our final example here is intended to show that the suggested theory provides reasonable solutions also for situations that involve compound logical descriptions.

Example 3. Consider the following default theory

$$\{A \rightarrow B, B \rightarrow C, B \rightarrow D, A \rightarrow \neg(C \wedge D)\}.$$

Suppose that A is a given fact. Note first that in this theory no pair of conditionals stands in a direct conflict with each other. Nevertheless, there is an obvious specificity conflict between the rule $A \rightarrow \neg(C \wedge D)$, on the one hand, and a combination of two rules $B \rightarrow C$ and $B \rightarrow D$, on the other hand. Intuitively, this conflict could be resolved if at least one of the latter conditionals could be rejected, though we have no reason to prefer one of them over another.

Now, in our representation of this conditional theory, the commitment to $A \rightarrow \neg(C \wedge D)$ implies $\neg(A/(C \wedge D))$, so when A/B is accepted, Forward Rejection will give us $\neg(B/(C \wedge D))$. Consequently, B/C and B/D cannot be

simultaneously accepted, since otherwise we would have $B/(C \wedge D)$ by the And rule. On the other hand, if we accept both A/B and $A/\neg(C \wedge D)$, we obtain B and $\neg(C \wedge D)$ (by MP), which logically implies $(B \wedge \neg C) \vee (B \wedge \neg D)$, and therefore $\neg(B/C) \vee \neg(B/D)$ by MP again. Consequently, acceptance of B/C will reject B/D and vice versa, acceptance of B/D will reject B/C . As a result, our theory will have precisely two extensions in which, respectively, either B/D or B/C will be rejected, while the rest of the conditionals will be accepted.

Defeasible Inheritance

As a matter of fact, our approach has also many features in common with known representations of defeasible inheritance in logic programming (You, Wang, and Yuan 1999) and in default logic (Dung and Son 2001). Moreover, similarly to that latter, we have an opportunity to provide a straightforward representation of defeasible inheritance in our framework. Due to space limitations, however, we can only be brief here.

Defeasible inheritance nets is a logical framework that has been originally designed to capture reasoning in taxonomic hierarchies that allowed to have exceptions. The theory of reasoning in such taxonomies has been called *nonmonotonic inheritance*. The guiding principle in resolving potential conflicts in such hierarchies was a *specificity principle* ((Poole 1985; Touretzky 1986)): more specific information should override more generic information in cases of conflict. Though obviously related to nonmonotonic reasoning, nonmonotonic inheritance relied more heavily on graph-based representations than on traditional logical tools. Nevertheless, it has managed to provide a plausible analysis of reasoning in this restricted context.

Let Γ be a consistent defeasible inheritance network. A credulous extension of Γ is defined as usual (see (Horty 1994) and the Appendix), with the only simplification that it is restricted to the set of paths from object nodes (as in (You, Wang, and Yuan 1999)).

For a propositional atom q , $\hat{q}$ will denote a corresponding literal, that is either q , or $\neg q$, while $\neg\hat{q}$ will denote the literal complementary to $\hat{q}$.

$D(\Gamma)$ will denote the default theory corresponding to Γ as follows³:

- any object link will correspond to an axiom $\hat{p} \in W$;
- every defeasible link will correspond to a defeasible rule of the form $p \rightarrow \hat{q}$, so $p : p/\hat{q} \vdash p/\hat{q}$ will be added to the default rules, as well as the corresponding Commitment rule:

$$p : \vdash \neg(p/\neg\hat{q}).$$

Then the following theorem shows, in effect, that the resulting default theory provides an exact formalization of defeasible nets. An abridged proof of this result is given in the Appendix below.

Theorem. *A set of paths Φ is a credulous extension of a consistent nonmonotonic inheritance network Γ if and only*

³In order to simplify the representation, we will consider only defeasible networks without strict links.

if there is an extension u of $D(\Gamma)$ such that Φ coincides with the set of paths constructed from the set of links $\{p \rightarrow \hat{q} \in \Gamma \mid p/\hat{q} \in u\}$.

The gain in simplicity and modularity provided by the above representation could be made vivid by comparing it with the much more complex translation of defeasible inheritance into default logic described in (Dung and Son 2001).

Conclusions

Nonmonotonic reasoning is not just a syntax plus nonmonotonic semantics. An account of the underlying *logic* behind our commonsense reasoning can provide immense improvement in the quality of representations. In our case, it has been shown that when the normality assumptions mediating defeasible rules are represented as conditionals having a relatively simple underlying logic, the resulting representation has allowed us to capture defeasible inheritance and specificity, generalized to the full classical language. Furthermore, in the suggested representation of defeasible conditionals we even did not use semi-normal defaults, as was advocated in (Reiter and Criscuolo 1981). All that was required to achieve a proper nonmonotonicity in our logical setting was the basic formalism of Reiter's normal defaults. The possibility of restriction to normal default rules was based, however, on the extensive use of the monotonic inference rules (without justifications). In fact, it has been shown in (Bochman 2008) that this possibility is quite general, since any default theory can be reduced ultimately to a simple default theory that contains only monotonic rules and plain default assumptions, aka supernormal defaults of the form $: A \vdash A$.

In this study we have restricted ourselves to the task of representing a particular system of defeasible inference, namely (a generalization of) nonmonotonic inheritance. It should be clear, however, that the approach itself is much more general, so hopefully it could provide a rigorous representation also for other forms of defeasible entailment, for example for (nonmonotonic) causal reasoning. However, this is a subject for further research.

References

- Barwise, J.; Gabbay, D.; and Hartonas, C. 1995. On the logic of information flow. *Bull. of the IGPL* 3:7–50.
- Benferhat, S.; Cayrol, C.; Dubois, D.; Lang, J.; and Prade, H. 1993. Inconsistency management and prioritized syntax-based entailment. In Bajcsy, R., ed., *Proceedings Int. Joint Conf. on Artificial Intelligence, IJCAI'93*, 640–645. Chambery, France: Morgan Kaufmann.
- Bochman, A. 2001. *A Logical Theory of Nonmonotonic Inference and Belief Change*. Springer.
- Bochman, A. 2007. On default representation of defeasible inference and specificity. In *Logical Formalizations of Commonsense Reasoning. Papers from the 2007 AAAI Spring Symposium*, Technical Report SS-07-05, 19–24. Stanford, Ca.: The AAAI Press.

- Bochman, A. 2008. Default logic generalized and simplified. In *The Tenth International Symposium on Artificial Intelligence and Mathematics, ISAIM 2008*.
- Dung, P. M., and Son, T. C. 2001. An argument-based approach to reasoning with specificity. *Artificial Intelligence* 133:35–85.
- Etherington, D., and Reiter, R. 1983. On inheritance hierarchies with exceptions. In *Proceedings of AAAI-83*, 104–108.
- Geffner, H., and Pearl, J. 1992. Conditional entailment: bridging two approaches to default reasoning. *Artificial Intelligence* 53(2-3):209–244.
- Geffner, H. 1992. *Default Reasoning. Causal and Conditional Theories*. MIT Press.
- Horty, J. F. 1994. Some direct theories of nonmonotonic inheritance. In Gabbay, D. M.; Hogger, C. J.; and Robinson, J. A., eds., *Handbook of Logic in Artificial Intelligence and Logic Programming 3: Nonmonotonic Reasoning and Uncertain Reasoning*. Oxford: Oxford University Press.
- Konolige, K., and Myers, K. L. 1989. Representing defaults with epistemic concepts. *Computational Intelligence* 5:32–44.
- Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44:167–207.
- Lehmann, D., and Magidor, M. 1992. What does a conditional knowledge base entail? *Artificial Intelligence* 55:1–60.
- Lehmann, D. 1995. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence* 15:61–82.
- Lifschitz, V. 1985. Computing circumscription. In *Proc. 9th Int. Joint Conf. on Artificial Intelligence, IJCAI-85*, 121–127. Morgan Kaufmann.
- Marek, W., and Truszczyński, M. 1993. *Nonmonotonic Logic, Context-Dependent Reasoning*. Springer.
- McCarthy, J. 1980. Circumscription – a form of nonmonotonic reasoning. *Artificial Intelligence* 13:27–39.
- McCarthy, J. 1986. Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence* 13:27–39.
- McDermott, D., and Doyle, J. 1980. Nonmonotonic logic. *Artificial Intelligence* 13:41–72.
- Nute, D. 1994. Defeasible logic. In Gabbay, D., and Hogger, C., eds., *Handbook of Logic for Artificial Intelligence and Logic Programming*, volume III. Oxford University Press. 353–395.
- Pearl, J. 1990. System Z: A natural ordering of defaults with tractable applications to default reasoning. In *Proceedings of the Third Conference on Theoretical Aspects of Reasoning About Knowledge (TARK'90)*, 121–135. San Mateo, CA: Morgan Kaufmann.
- Poole, D. 1985. On the comparison of theories: Preferring the most specific explanation. In *Proceedings Ninth International Joint Conference on Artificial Intelligence*, 144–147.
- Poole, D. 1988. A logical framework for default reasoning. *Artificial Intelligence* 36:27–47.
- Poole, D. 1990. Dialectics and specificity: conditioning in logic-based hypothetical reasoning (preliminary report). In *Proc. 8th biennial conf. of the Canadian Society for Computational Studies of Intelligence*, 69–76. San Francisco, CA: Morgan Kaufmann.
- Reiter, R., and Criscuolo, G. 1981. On interacting defaults. In *Proceedings of Int. Joint Conf. on Artificial Intelligence, IJCAI-81*, 270–276.
- Reiter, R. 1980. A logic for default reasoning. *Artificial Intelligence* 13:81–132.
- Reiter, R. 1987. A theory of diagnosis from first principles. *Artificial Intelligence* 32:57–95.
- Touretzky, D. S. 1986. *The Mathematics of Inheritance Systems*. Los Altos: Morgan Kaufmann.
- Wagner, G. 1991. Ex contradictione nihil sequitur. In *Proc. Int. Joint Conf. on Artificial Intelligence, IJCAI-91*, 538–543. Morgan Kaufmann.
- You, J.-H.; Wang, X.; and Yuan, L. Y. 1999. Compiling defeasible inheritance networks to general logic programs. *Artificial Intelligence* 113:247–268.

Appendix. Proof of the Theorem

A defeasible inheritance network Γ is a tuple (N, E) where N is a set of nodes and E a set of positive and negative links between nodes. Nodes are divided into two disjoint classes: object nodes, and property nodes. An object node can only be used as a root node. A link is called an object link if its root is an object node. A *path* is a sequence of links such that the head of a preceding link coincides with the root of the next link in the sequence, and all the links in the sequence, except possibly the last, are positive. A path is *positive*, if all its links are positive, otherwise it is a *negative* path. A positive path from p to q through a path σ will be denoted by $p\sigma q$. In addition, $p\sigma q \rightarrow r$ will denote a positive path with a prefix σ and the last link $q \rightarrow r$, while $p\sigma q \not\rightarrow r$ will denote a negative path with a prefix σ and the last link $q \not\rightarrow r$.

Path constructibility. Suppose that Φ is a path set of Γ . A path σ is *constructible* in Φ iff (i) it is an object link, or (ii) σ consists of a prefix $\tau \in \Phi$ and the last link that belongs to E .

Conflict. A positive (resp., negative) path σ is *conflicting* in Φ iff $\sigma \in \Phi$ and Φ contains a negative (resp., positive) path with the same beginning and end nodes.

Off-path preemption. A positive path $r\sigma u \rightarrow s$ (respectively, a negative path $r\sigma u \not\rightarrow s$) is preempted in Φ iff $\sigma \in \Phi$ and there is a node v such that

- There is a link $v \not\rightarrow s \in E$ (respectively, $v \rightarrow s \in E$) and
- either $r = v$ or there is a path $r\sigma_1 v\sigma_2 u \in \Phi$.

A path σ is *defeasibly inheritable* in Φ iff it is constructible, not conflicting and not preempted in Φ .

Definition. A set Φ of paths is a *credulous extension* of an inheritance net Γ if it coincides with the set of paths that are defeasibly inheritable in Φ .

A defeasible inheritance network is called *consistent* if it does not simultaneously contain a pair of conflicting links $p \rightarrow q$ and $p \not\rightarrow q$. For such inheritance networks we can prove the following

Theorem. A set of paths Φ is a credulous extension of a consistent defeasible inheritance network Γ if and only if there is an extension u of $D(\Gamma)$ such that Φ coincides with the set of paths constructed from the set of links

$$\{p \rightarrow q \in E \mid p/q \in u\} \cup \{p \not\rightarrow q \in E \mid p/\neg q \in u\}.$$

Proof. ($\Rightarrow$) If Φ is a credulous extension of Γ , let $l(\Phi)$ be the set of all non-object links appearing on paths of Φ . Also, let $R = \{p/\hat{q} \mid p \rightarrow \hat{q} \in l(\Phi)\}$, and define $Cl(R)$ to be the set of all conditionals A/B that are derivable from R by the rules Dominance, Transitivity and And. It can be easily shown that $p/\hat{q} \in Cl(R)$ if and only if Φ contains a path via p to $\hat{q}$ (cf. Lemma C3 in (Dung and Son 2001)).

Let u be the closure of the set $W \cup R$ with respect to the classical entailment and the strict rules of $D(\Gamma)$, i.e., Transitivity, And, Forward Rejection and Commitment.

Lemma 1. If q is a propositional atom, then $\hat{q} \in u$ iff Φ supports $\hat{q}$.

Proof. If $\hat{q}$ is supported by Φ , there is a path σ in Φ from an object link to $\hat{q}$. Since $p_i/p_{i+1} \in u$, for any link $p_i \rightarrow p_{i+1}$ on this path, we have $p_i \supset p_{i+1} \in u$ (by MP). Clearly then $\hat{q} \in u$. To prove the other direction, we define a propositional theory $v = \text{Th}(v_s \cup Cl(R) \cup N(R))$, where v_s is the set of ordinary literals that are supported in Φ , and $N(R) = \{\neg(A/B) \mid A/B \notin Cl(R)\}$. Note that, for any conditional atom A/B , we have either $A/B \in v$, or $\neg(A/B) \in v$. We will show that $u \subseteq v$.

We will demonstrate first that if $A/B \in Cl(R)$, then $v_s \models A \supset B$, by induction on the derivations of A/B . If $p/\hat{q} \in R$, then Φ contains a path that includes $p \rightarrow \hat{q}$, so $\hat{q}$ is supported, and consequently $v_s \models p \supset \hat{q}$. Now, if $A \models B$, then clearly $v_s \models A \supset B$. If $A/(B \wedge C)$ has been obtained from A/B and A/C by the rule And, then $v_s \models A \supset B$ and $v_s \models A \supset C$ by the inductive assumption, so $v_s \models A \supset B \wedge C$. The proof is similar if A/C has been obtained from A/B and B/C by Transitivity. Hence the claim holds.

Now we can show that the axiom MP belongs to v . If $A/B \notin v$, then $\neg(A/B) \in v$, so $A \wedge (A/B) \supset B \in v$. Assume that $A/B \in v$. Then by the preceding claim $v_s \models A \supset B$, and hence again $A \wedge (A/B) \supset B \in v$. Consequently, all the axioms of $D(\Gamma)$ are included in v . In addition, v is closed with respect to Transitivity and And (since it includes $Cl(R)$). In addition, suppose that $A/B \in v$, $\neg(A/C) \in v$, but $\neg(B/C) \notin v$. Then $B/C \in v$, and hence $A/C \in v$ by Transitivity - a contradiction. Thus, Forward Rejection is also satisfied. Finally, we will show that v is closed wrt the commitment rules. Suppose that $p \rightarrow \hat{q} \in \Gamma$, $p \in v$, but $\neg(p/\neg\hat{q}) \notin v$. Then $p/\neg\hat{q} \in v$, that is, $p/\neg\hat{q} \in Cl(R)$, and hence Φ should contain a path σ from an object link via p to $\neg\hat{q}$, which is impossible, since it is preempted by $p \rightarrow \hat{q}$. Thus, v includes W and is closed wrt all the strict rules of $D(\Gamma)$. Consequently, $u \subseteq v$.

Assume now that $\hat{q}$ is not supported in Φ . Clearly, v_s is exactly the set of ordinary literals in v , so $\hat{q} \notin v$, and therefore $\hat{q} \notin u$. This completes the proof of the lemma. $\square$

We will show now that u is an extension of $D(\Gamma)$, that is, $u = \mathcal{D}(u)$.

For the inclusion $u \subseteq \mathcal{D}(u)$, we show first that if $p \rightarrow \hat{q} \in l(\Phi)$, then $p \wedge (p/\hat{q}) \in \mathcal{D}(u)$. Now, $p \rightarrow \hat{q} \in l(\Phi)$ only if Φ contains a path σ that supports p . We will prove the claim by induction on the length of σ . If $p \rightarrow \hat{q}$ is the first non-object link, then $p \in W$. In addition, we can apply $p : p/\hat{q} \vdash p/\hat{q}$ to derive $p/\hat{q}$ (since $p/\hat{q} \in u$). Assume now that σ is a path of length n having $r \rightarrow p$ as the last link. By the inductive assumption, $r \wedge (r/p) \in \mathcal{D}(u)$, so $p \in \mathcal{D}(u)$ by MP: $r \wedge (r/p) \supset p$. Consequently, $p/\hat{q} \in \mathcal{D}(u)$ by the rule $p : p/\hat{q} \vdash p/\hat{q}$, and we are done.

The above claim implies $\{p/\hat{q} \mid p \rightarrow \hat{q} \in l(\Phi)\} \subseteq \mathcal{D}(u)$, so by the definition of u we immediately conclude $u \subseteq \mathcal{D}(u)$.

For the inclusion $\mathcal{D}(u) \subseteq u$, it is sufficient to show that u is closed with respect to all the rules of the default theory $D(\Gamma)$. First, u includes W . Suppose that $p \rightarrow \hat{q} \in \Gamma$ but the rule $p : p/\hat{q} \vdash p/\hat{q}$ does not hold in u , that is $p \in u$, $p/\hat{q} \notin u$ and $\neg(p/\hat{q}) \notin u$. Now, $p/\hat{q} \notin u$ implies $p \rightarrow \hat{q} \notin l(\Phi)$, which can happen only if it is either conflicted or preempted in Φ .

Suppose first that $p \rightarrow \hat{q}$ is conflicted in Φ . Then $\neg \hat{q}$ is supported by Φ , and hence $\neg \hat{q} \in u$. Consequently $\neg(p/\hat{q}) \in u$ by MP, contrary to our assumptions. Suppose now that $p \rightarrow \hat{q}$ is preempted in Φ . Then Γ contains a link $r \rightarrow \neg \hat{q}$ such that there is a positive path σ via r to p in Φ . Let us consider the sub-path of σ from r to p . Since $p_i/p_{i+1} \in u$, for any link $p_i \rightarrow p_{i+1}$ that belongs to this sub-path, we obtain $r/p \in u$ by transitivity. But by the Commitment rule $\neg(r/\hat{q}) \in u$, and hence $\neg(p/\hat{q})$ by Forward Rejection, again contrary to our assumptions. The obtained contradiction shows that u is closed with respect to all the rules of $D(\Gamma)$, and consequently $\mathcal{D}(u) \subseteq u$ holds.

Finally, we have to show that $l(\Phi)$ coincides with $\{p \rightarrow \hat{q} \in \Gamma \mid p/\hat{q} \in u\}$. Assume that $p \rightarrow \hat{q} \in \Gamma$ and $p/\hat{q} \in u$. By the construction of u , this can happen only if $p/\hat{q} \in Cl(R)$, and therefore there is a path in Φ of the form $\sigma p \tau \hat{q}$. Let us consider the path $\sigma_1 = \sigma p \rightarrow \hat{q}$. This path is constructible in Φ , and it is clearly neither conflicted, nor preempted in Φ . Since Φ is a credulous extension, we conclude $\sigma_1 \in \Phi$, and therefore $p \rightarrow \hat{q} \in l(\Phi)$.

($\Leftarrow$) Suppose that u is an extension of $D(\Gamma)$, and define Φ to be the set of paths constructed from the object links and the links $\{p \rightarrow \hat{q} \in \Gamma \mid p/\hat{q} \in u\}$. Let u_Φ denote the set of all literals that are supported by Φ . If $\hat{p} \in u_\Phi$, there is a path $\sigma \in \Phi$ from an object link to $\hat{p}$. Since $p_i/p_{i+1} \in u$, for any link $p_i \rightarrow p_{i+1}$ on this path, we have $p_i \supset p_{i+1} \in u$ (by MP). Clearly then $\hat{p} \in u$. In addition, we will need the following

Lemma 2. $A \in u$ only if $u_\Phi \models A$, and $A/B \in u$ only if $u_\Phi \models A \supset B$.

Proof. Since $u = \mathcal{D}(u)$, we will prove these two claims by a simultaneous induction on the derivations in $\mathcal{D}(u)$.

If $\hat{p} \in W$ since $\hat{p}$ corresponds to an object link, then clearly $\hat{p} \in u_\Phi$. Also, if $A/B \in W$ by Dominance, then $A \models B$, and hence $u_\Phi \models A \supset B$.

If $p/\hat{q}$ has been obtained by the default rule $p : p/\hat{q} \vdash p/\hat{q}$, then $p \rightarrow \hat{q} \in \Phi$ and $p \in u$. By the inductive assumption, $u_\Phi \models p$, so $p \in u_\Phi$, and hence p is supported by Φ . Therefore, $\hat{q}$ is also supported by Φ , and hence $u_\Phi \models p \supset \hat{q}$.

If $A/(B \wedge C)$ has been obtained from A/B and A/C by the rule And, then $u_\Phi \models A \supset B$ and $u_\Phi \models A \supset C$ by the inductive assumption, so $u_\Phi \models A \supset B \wedge C$. The proof is similar if A/C has been obtained from A/B and B/C by Transitivity.

Finally, the axiom $A \wedge A/B \supset B$ cannot be used for deriving new conditionals, but only for deriving $A \supset B$ when A/B has been proved. But in this case $u_\Phi \models A \supset B$ already by the inductive assumption, and we are done. $\square$

Now we will show that Φ is a credulous extension of Γ . We will prove first that Φ is defeasibly inheritable in Φ , that is, any path in Φ is constructible, conflict-free and not preempted in Φ . Now, any path in Φ is constructible by the definition. If there is a conflicted path in Φ , then there is an atom p , such that both p and $\neg p$ belong to u , which is impossible. Finally, assume that σ is a preempted path in Φ , and $p \rightarrow \hat{q}$ is the last link on σ . Then there exist a link $r \rightarrow \neg \hat{q}$ and a path τ via r to p that belongs to Φ . Now, $p/\hat{q} \in u$,

and $p_i/p_{i+1} \in u$, for any link $p_i \rightarrow p_{i+1}$ on the sub-path of τ from r to p . Therefore, $r/p \in u$ and hence $r/\hat{q} \in u$ by transitivity, contrary to the commitment to $r \rightarrow \neg \hat{q}$. Therefore, σ is not preempted in Φ .

Finally, we will show that any path that is defeasibly inheritable in Φ also belongs to Φ . Suppose that σ is defeasibly inheritable in Φ . If it is an object link, it is in Φ . Assume that σ is composed of a prefix $\tau \in \Phi$ and the last link $p \rightarrow \hat{q}$. Then p is supported in Φ , and therefore $p \in u$. Next we are going to show that $\neg(p/\hat{q}) \notin u$. Suppose that $\neg(p/\hat{q}) \in u$. Since $u = \mathcal{D}(u)$, we have that $\neg(p/\hat{q})$ should be derivable from W using the strict and active normal default rules of $D(\Gamma)$. As can be seen, this can happen only if $\neg(p/\hat{q})$ is obtained either (i) from the axiom MP when $p \wedge \neg \hat{q} \in u$, or (ii) by the commitment rule $p : \vdash \neg(p/\hat{q})$, given a link $p \rightarrow \neg \hat{q} \in \Gamma$, or (iii) by applying Forward Rejection. In the case (i) we have that $\neg \hat{q}$ is supported by Φ , and therefore σ is conflicted in Φ , contrary to our assumptions. The case (ii) is also impossible, since Γ is a consistent inheritance net. Hence $\neg(p/\hat{q})$ is obtained after one or more successive applications of Forward Rejection. Due to the form of these rules, this sequence should start with $\neg(A/\hat{q})$, where A is some proposition such that $A/p \in u$. Now since $\neg(A/\hat{q})$ is already not obtained by Forward Rejection, it should be provable either by MP, or by Commitment. But in the first case we would have again that $\neg \hat{q}$ is supported by Φ , and therefore σ is conflicted in Φ , contrary to our assumptions. In the second case we will have that A is an atom, and $A \rightarrow \neg \hat{q} \in \Gamma$, which contradicts our assumption that σ is not preempted in Φ . Thus, $\neg(p/\hat{q}) \notin u$, and hence we can apply the default rule $p : p/\hat{q} \vdash p/\hat{q}$ and conclude $p/\hat{q} \in u$. It then follows that $\sigma \in \Phi$. This completes the proof. $\square$



2008

This paper was published in the *Proceedings, Eleventh International Conference on Principles of Knowledge Representation and Reasoning*, edited by Gerhard Brewka and Jérôme Lang (Menlo Park, California: AAAI Press).

The 2008 KR conference was held in Sydney, Australia, 16–19 September, 2008.